



AIDDATA

A Research Lab at William & Mary

WORKING PAPER 125

May 2023

Landmine Clearance and Economic Development: Evidence from Nighttime Lights, Multispectral Satellite Imagery, and Conflict Events in Afghanistan

Christian Baehr
Princeton University

Ariel BenYishay
AidData, William & Mary

Rachel Sayers
AidData, William & Mary

Kunwar Singh
AidData, William & Mary

Madeleine Walker
AidData, William & Mary

Abstract

The widespread prevalence of unexploded landmines in many countries leads to thousands of deaths and injuries annually. In addition, landmines may impede the flow of goods and people; this not only restricts trade and labor flows, but also hinders access to schools and medical care. Moreover, landmines can block the productive use of contaminated land. We study the clearance of more than 15,000 hazardous areas in Afghanistan carried out since 1992, obtaining precise geographic boundaries for these areas. We identify a window during which a policy shift in clearance targeting created plausibly exogenous variation in the timing of clearance, which we study using a two-way fixed effects panel design. The clearance of hazardous areas leads to changes in land use observed using multispectral, moderate-resolution Landsat-series imagery, as well as to increases in economic activity reflected in nighttime lights data. Our precise clearance data and satellite imagery allows us to observe the shifts occurring even in the small towns and villages that comprise most of our sample. We also find reductions in conflict risks due to the clearance of hazardous areas. Landmine action efforts thus appear to have substantial economic effects even in areas with ongoing conflict risks.

Author Information

Christian Baehr

Princeton University

Ariel BenYishay

AidData, William & Mary

Rachel Sayers

AidData, William & Mary

Kunwar Singh

AidData, William & Mary

Madeleine Walker

AidData, William & Mary

The views expressed in AidData Working Papers are those of the authors and should not be attributed to AidData or funders of AidData's work, nor do they necessarily reflect the views of any of the many institutions or individuals acknowledged in the AidData Working Paper Series.

Acknowledgments

The authors are grateful to Anne-Li Naucler, Tom Gillhespy, and George Bowles for facilitating access to the mine clearance geospatial data and providing excellent comments, as well as to seminar participants at VCU, William & Mary, and the Center for Environmental Economics - Montpellier.

1 Introduction

The burgeoning geospatial data on conflict events has powered a new spate of research over the past decade, uncovering insights on the interplay between active conflicts and economic development (Sundberg and Melander, 2013; Fjelde and Hultman, 2014; Michalopoulos and Papaioannou, 2016; Uexkull et al., 2016; Wagner et al., 2018). Yet even after active conflict subsides, unexploded ordinances like landmines remain major dangers throughout much of the world. Because landmines cost relatively little to produce (typically \$3-\$30), they have been heavily used throughout the developing world, and more than 100 million landmines remained buried across 60 countries as of 2021 (ICBL, 2021). The detonation of unexploded landmines leads to thousands of deaths and injuries annually (Frost et al., 2017). In addition to being public health threats, landmines may also severely limit economic development. When located near paths or roads, they can impede the flow of goods and people, restricting not only trade and labor flows, but also access to schools and medical care. Moreover, landmines can restrict the productive use of contaminated and surrounding land, preventing fields from being sown and commercial activities from being expanded.

Yet, because precise geospatial data on landmines has been very limited, relatively little is known about the extent of the limitations they impose on development, or, conversely, about the gains in economic development afforded by landmine clearance. To date, only four studies have examined these development impacts. Three of these studies use cross-sectional instrumental variable models to compare populations based on their proximity to hazardous areas, finding worse child nutrition, greater poverty, and lower educational attainment in higher risk areas (Arcand, Rodella-Boitreaud and Rieger, 2015; Merrouche, 2008, 2011). More recently, Chiovelli, Michalopoulos and Papaioannou (2018) obtain data on both the location and timing of clearance activities in Mozambique, allowing them to adopt a panel design. These authors find large benefits from clearance of transport corridors and major hubs, but relatively limited effects in rural areas. This may be in part because they rely on nighttime lights emissions as their primary outcome measurement, and the lights data may not capture meaningful changes in unelectrified areas.

While these studies offer promising findings, major gaps in our knowledge remain. First, these studies focus on post-conflict settings, but many landmine clearance efforts take place in areas with ongoing conflict. Are the findings from post-conflict clearance representative of the gains from clearance when ongoing conflict risks might dampen broader economic changes? This may be particularly challenging to assess because the timing of clearance may depend in part on the ebb and flow of conflict, potentially confounding the roles of clearance and conflict dynamics. Secondly, although landmine clearance enables repurposing land, none of the aforementioned studies have examined land use as an outcome. Finally, there is still no clear evidence on whether landmine clearance leads to economic development in both rural and urban areas.

The latter question is critical to current policy debates, which focus on how to target the relatively limited landmine clearance efforts on an efficient and equitable basis. Landmine clearance organizations receive a total of \$650M annually worldwide, but with removal costs often exceeding \$1000 per mine, recent rates of removal have only reached a maximum of 200,000 mines annually (ICBL, 2021). Even ignoring the >2M new mines planted annually, full worldwide clearance would take >1,000 years at this rate. These constraints lead to fre-

quent targeting and prioritization debates, many of them centered on the relative importance of clearance in rural or agricultural land.

We provide several key contributions. First, we obtain a comprehensive, long-term, temporally and spatially precise dataset on hazardous area clearance in Afghanistan. Our spatial data includes the boundaries of each hazardous area, recorded during the actual clearance activities by each field team. This affords us greater statistical precision than previous studies, which relied on aggregated administrative units or geographically perturbed survey cluster locations. We combine the clearance data not only with nighttime lights data, but also with much finer land use changes in and around the cleared sites derived from multispectral (MS) satellite imagery. Our high-quality landmine clearance boundary data and use of MS imagery provides us with improved ability to detect land change over time in rural as well as urban and peri-urban areas. Unlike [Chiovelli, Michalopoulos and Papaioannou \(2018\)](#), we find some of the largest effects occur outside of urban areas. In fact, in villages, we observe building construction and expansion near newly cleared areas. At the same time, we do not find increased conversion of cleared land to farming, likely because water and fertilizer availability are the most constraining factors to agricultural production.

Moreover, by also incorporating time-varying conflict conditions, we show that landmine clearance can actually reduce conflict risks. These conflict reductions appear independent of economic improvements due to clearance, showing much more nuanced impacts from landmine clearance than previously understood.

Our paper proceeds as follows: in Section 2, we lay out the context of landmine clearance efforts in Afghanistan over the past two decades. In Section 3, we describe the geospatial data we obtain on the clearance activities, as well as the satellite and survey data we use. Section 4 provides our empirical strategy for identifying the impacts of clearance, and Section 5 lays out our results, each separated into subsections on nighttime lights (NL), land use, and conflict. We offer conclusions in Section 6.

2 Landmine Deployment and Clearance in Afghanistan

2.1 Overall context

The extent of landmine contamination in Afghanistan—as well as the scale and longevity of Afghanistan’s demining program—present a uniquely important context for understanding the impacts of demining on economic development in the face of ongoing conflict risks.

The Soviet Army was the first to plant landmines in Afghanistan during their occupation of the country from 1978-1989. In the four decades of war that followed, most major players, including the Taliban, Mujahideens, and many tribal groups, have used landmines in a variety of settings ([Chawla, 2011](#)).¹ Since 1978, over 40,000 casualties have resulted from landmines and other unexploded ordinances ([UNMAS, 2021](#)). Every province of Afghanistan has been contaminated with landmines, and despite the major progress in mine clearance that we detail below, nearly every province continued to have landmine contamination in 2021 (affecting over 1,500 communities in 253 districts) ([UNMAS, 2021](#)).

¹The United States is known to have used only one landmine during its presence in Afghanistan, in a Special Operations mission in 2002 ([Crowley and Ismay, 2022](#)).

While many landmines remain buried, Afghanistan has nonetheless made major progress in landmine clearance over the past 30 years. Between 1989 and 2021, 81.3% of known hazardous areas in the country were cleared or determined non-hazardous (UNMAS, 2022). Demining initiatives began when the United Nations established the Mine Action Program of Afghanistan (MAPA) in 1989. Between 1989 and 2008, the UN and a consortium of international NGOs managed mine clearance operations in the country. Landmine clearance in this period was restricted by a lack of strategic planning and the destruction of demining equipment in the conflict between the US-backed Northern Alliance and the Taliban in the wake of the 2001 World Trade Center attacks. Before 2008, at most 40 square km of minefields were ever cleared in a single year. In 2010, the number had increased to 65 square km. (Paterson, 2012).

Afghanistan gained official authority over national landmine clearance operations in January 2008, when mine action management transitioned to the Department of Mine Clearance (DMAC) in the Afghanistan National Management Authority (ANDMA), under the President's office. Concurrent with this transfer of oversight, MAPA went through a number of decentralizing reforms. Local staff were prioritized over international staff, and an operations reform restructured demining units into smaller units (Paterson, 2012). MAPA also encouraged local and international NGOs to adopt a cluster approach to clearance tasks. The new approach emphasized clearing multiple hazard sites in one area, instead of just high-priority tasks, intended to utilize equipment and personnel more efficiently. Operators also increased security by strengthening their engagement and local relationships (Monitor, 2009).

These reforms led to major changes in the overall rate of mine clearance, as well as in the distribution of which sites were cleared. Between 1989 and 2007, a total of 3,084 hazardous areas were cleared across the entire country (an average of approximately 171 each year), with these areas concentrated around provincial capitals, towns, and main transportation corridors. Between 2008 and 2020, 17,648 hazardous areas were cleared, with these areas spanning a much broader set of locations.

The clearance of mines has taken place despite ebbs and flows of conflict risks, potentially dampening the extent to which landmine clearance can benefit economic growth. Taliban attacks on civilians rose steadily in the years leading up to the US withdrawal from Afghanistan, increasing to 10,469 in 2021 (SIGAR, 2021). Despite high conflict levels, 1,470 mines were cleared per year on average between 2008 and 2020. Continuing conflict between the Taliban and resistance groups will likely put civilians at high risk of experiencing violence into the future (ACLEDD, 2022).


2.2 Mine Clearance Data

We obtain the Mine Action Geo Data from DMAC, corresponding to all hazardous area clearance activities carried out throughout the country between 1992 and 2020. The spatial data include the polygon boundaries of each confirmed or suspected hazardous area, recorded during the survey and clearance conducted by each ground team, as well as the precise dates of these activities. There are 20,749 areas reflected in this data, of which more than 85% (17,913) were cleared since 2008.

For illustration, Figure 1 maps the precise boundaries of cleared hazardous polygons in


one area, showing the mix of polygon types and intended land uses. The data we obtain also contains information collected by the field teams on the type of blockage caused by the landmines (with multiple blockage types frequently cited). The vast majority of these areas include blockages to grazing (67%) and agricultural (27%) uses, as well as housing (12%), road (8%), water (3%), and other infrastructure types.

Figure 1: Hazardous Areas cleared in one location



We use this data to assess the expansion of mine clearance after the 2008 reforms. Figure 2 maps the hazardous areas cleared over time for Kabul province, depicting the substantial variation in timing of clearance even for relatively nearby hazardous areas in one province.

Figure 2: Hazardous Areas Cleared in Kabul Province



As described below, we use the precise boundaries of each cleared hazardous area to assess both the changes to broader economic activity reflected in 1 km² NL measures as well as to the specific land uses of the cleared areas observable in 30m resolution MS imagery.

3 Outcome Data

3.1 Satellite-based Data

As our primary outcome measures for economic development, we require data that have been consistently collected over the past two decades with sufficiently dense geographic coverage, thereby ruling out household surveys and other ground-based measures. Instead, we use several satellite-based sources, including NL and MS daytime imagery.

3.1.1 Nighttime Lights

Our source for NL emission data is the Defense Meteorological Satellite Program (DMSP), which reflects the NL emission of 1km x 1km grid-cells worldwide during the years 1992-2013.² The DMSP NL have been extensively used as proxies of economic activity, and have been shown to correlate with a variety of economic indicators at national and subnational

²The DMSP-based NL series has since been replaced with the VIIRS series; while the two have been recently inter-calibrated, the VIIRS data are collected later at night and may be even less able to detect economic changes in areas with little electrification.

scales (Chen and Nordhaus, 2010; Henderson, Storeygard and Weil, 2012; Mellander et al., 2015). We join the hazardous area (HA) polygons to the DMSR grid and limit our sample to those grid-cells that were partially or fully classified as confirmed hazardous areas (CHAs) or suspected hazardous areas (SHAs) with landmines for at least one year during 1992-2013. We thus obtain a cell x year dataset for all cells with at least some hazardous areas in our time window.

To determine treatment dates for each grid cell, we match grid cells with intersecting hazardous areas. Our primary analysis sample includes only grid cells intersecting with hazardous areas that were ever cleared in our sample window (i.e. only grid cells cleared 2008-2013). As the area of each grid cell intersecting with hazardous areas varies, we weight the analysis by the percent of each grid cell covered by hazardous areas at baseline.

To construct our time-varying treatment measure, we assign each grid cell a treatment date corresponding to the latest clearance of all intersecting hazardous areas. That is, a grid cell is considered treated when all hazardous areas in that grid-cell have been cleared.

3.1.2 Multispectral Imagery


While widely used, NL are known to face important limitations in their ability to reflect conditions in rural parts of developing countries. Dugoua, Kennedy and Urpelainen (2018) show that NL closely track rural electrification, which may be driven in part by government infrastructure planning rather than local economic conditions. Moreover, Jean et al. (2016) document that lights do not correlate well with (survey-based) wealth measures at the bottom of the wealth distribution (presumably those in unlit areas), and Bluhm and McCord (2022) show weak correlations between NL and agricultural GDP. We thus supplement our NL data with economic measures derived from MS imagery. We utilize the visible and near-infrared portions of imagery from Landsat 5 and 8 sensors, available at 30-m resolution, to create a time series of land-use and land-cover (LULC) conditions at five-year intervals from 1995 to 2020.³ The 30-m resolution Landsat imagery has been shown to more robustly detect wealth differences than NL alone (Yeh et al., 2020). The accuracy of LULC classification depends on developing accurate training and testing data from geographies with similar landscape and climate characteristics, so we limit our sample to the three contiguous provinces with the greatest number of clearance activities between 1993 and 2020: Baghlan, Parwan, and Kabul.

The utilized LULC classification scheme with seven class types follows those defined by Thompson and Hubbard (2014) and FAO (2016). These classes are: (i) built-up, (ii) farmland/cropland, (iii) grassland, (iv) forests/shrubs, (v) bare areas, (vi) snow, and (vii) water. Following clearance, we expect conversion of landmine contaminated land to built-up and farmland uses, and potentially from grassland or bare area classes. The mosaic of Landsat imagery for each selected year was between June and September months with little to no precipitation and <5% cloud coverage. To create these LULC measures, we first manually collect training and testing data (i.e., polygon shapefiles). Figure 3 offers illustrative examples of polygons in our training data. To create these training data, polygons for each class are digitized by capturing a homogenous group of pixels on Google Earth Pro in a

³We opt for five-year intervals because of the computational costs of downloading and processing Landsat images over these scales.

Keyhole Markup Language (KML) format. To ensure the training data can be consistently applied across time steps, we create separate training data for 1995-2014 (training set 1) and 2015-2020 (training set 2), and visually inspect each polygon by relating it to available imagery for previous years on Google Earth to ensure that LULC remained unchanged. For example, if a farmland training polygon for the training set 1 is created using 2014 imagery, we cross-check the digitized homogeneous group of pixels in every available imagery between 1995 and 2014. If the digitized area changed at any time step, we discard that training polygon.

Figure 3: Examples of Polygons in Training Sets for Land Use Classification



We use the same approach to create the testing data containing 200 sites for each LULC class to assess the performance of mapping outcomes. Table 1 shows the number of polygons and points we collect for each class.

To classify the Landsat imagery into LULC classes, we apply a random forest (RF) algorithm.⁴ The RF algorithm may be suitable due to its accurate performance on an array of input variables (spectral bands), and its ability to determine significant variables in the classification. The RF operates by growing many individual decision trees from randomized subsets of training samples to maximum size without pruning and then selecting only the best split among a random subset at each node. The optimal classification is then determined

⁴We also applied a maximum likelihood (ML) algorithm, finding that it performed slightly worse than the RF algorithm. ML analyzes pixels across multiple bands and assigning class membership based on statistical probability using a set of training data (i.e., spectrally similar pixels representing a class of interest). ML estimates a probability density function for each class based on mean and covariance statistics generated from training data. Then, the ML algorithm assigns each pixel to its most probable class by comparing image pixel values with the class-level statistics (Pal and Mather, 2003).

Table 1: Training Datasets for LULC Classification

Class	Train. Set 1995-2014	Train. Set 2015-20	Testing data
Bare areas	11	14	200
Built-up	72	45	200
Farmland/cropland	58	51	200
Forest and shrubs	8	33	200
Grassland	88	157	200
Snow	28	30	200
Water	54	32	200
Total	319	362	1400

by selecting the most common classification results at each node within the group of multiple decision trees (Breiman, 2001). RF may be more efficient due to the use of both categorical and continuous variables, automatic handling of missing values and outliers, use of rule-based approach instead of distance calculation, use of non-linear parameters, and little to no impact by data noise. This algorithm is particularly useful for time-series, large extent land use assessments (e.g., the entire province or a region within a country). After applying the RF algorithm, we conduct accuracy assessments using the testing data we collect. We quantify and compare accuracy using three metrics: producer’s, user’s, and overall accuracy.⁵ These are shown in Table 2.

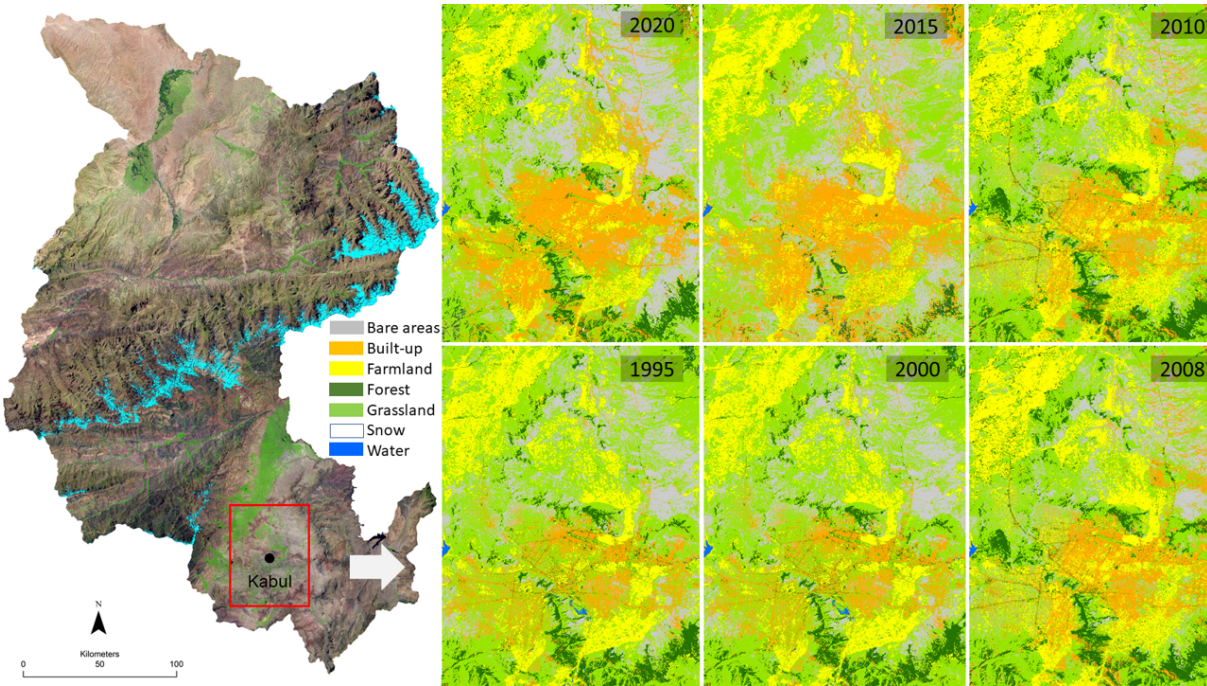
Table 2: LULC Prediction Accuracy Assessments

LULC class	2020		2015		2010		2008		2000		1995	
	User	Prod	U	P	U	P	U	P	U	P	U	P
Bare areas	94	91	95	89	94	75	90	85	90	84	81	78
Built-up	92	92	90	94	71	87	81	78	82	79	71	74
Farmland	93	96	87	87	86	86	95	93	92	89	97	94
Forest	90	86	86	84	71	70	99	93	98	92	89	89
Grassland	81	81	77	80	80	71	70	72	62	69	68	69
Snow	97	99	100	100	81	97	81	88	72	73	100	100
Water	98	99	99	99	98	100	93	97	92	99	99	100

The final LULC classification included a (1) classified image in which pixels are labeled by the class number used in the training data (Figure 4), and (2) class probability image. We extract the values at 30-m resolution for all hazardous areas for each year in our sample, then create shares of each area that fall into each class. Our final LULC analysis data includes 32,365 hazardous area (HA) by year observations.

⁵Overall accuracy for a land-use product represents the percentage of training points correctly classified when compared to testing points (ground reference). User and producer accuracies represent probabilities.

Figure 4: Land Use Classification Visualization



3.2 Conflict

We source conflict event data from the Uppsala Conflict Data Program (UCDP), which includes geocoded data on individual events of organised violence. Each event is categorised as (i) state-based armed conflict, (ii) non-state conflict, or (iii) one-sided violence. Because our interest is in the effects of mine clearance on nearby conflict, we focus on conflict events within a relatively close range of the hazardous areas. To do so, we construct a grid of 10km cells, and aggregate the total number of conflict events in each grid cell and each year. We then use this grid in several ways: first, to analyze clearance effects on conflict outcomes, we merge the conflict data with our HA perimeters by assigning the 10km grid cell into which each HA falls.⁶ We then determine the year in which all HAs within the grid cell have been cleared, and use this as our primary treatment measure in analyzing impacts on conflict. The unit of analysis is thus the grid cell by year.

We also use the conflict data in conjunction with our analysis of NL and LULC outcomes. To do so, we merge the 10km grid cells into our NL dataset, creating annual measures of nearby conflict for each NL grid cell and year. Finally, we also merge the conflict data at the 10km grid cell level into the HA by year level dataset we use for our LULC analysis.

⁶In cases where a HA perimeter overlaps the boundary between multiple 10km cells, we assign the values from the grid cell with the largest share of the perimeter.

4 Empirical Strategy

To identify the impacts of clearance, we utilize a panel difference-in-difference approach with two-way fixed effects (FEs). This approach capitalizes on the staggered nature of clearance activities, as well as the geospatial variation in hazardous areas being cleared. In this approach, hazardous areas are part of both the control and the treatment group, depending on the time of observation relative to clearance. Depending on the outcome, we utilize either grid cells that intersect with hazardous areas or the full hazardous area polygons as the spatial units of analysis.

4.1 Nighttime lights

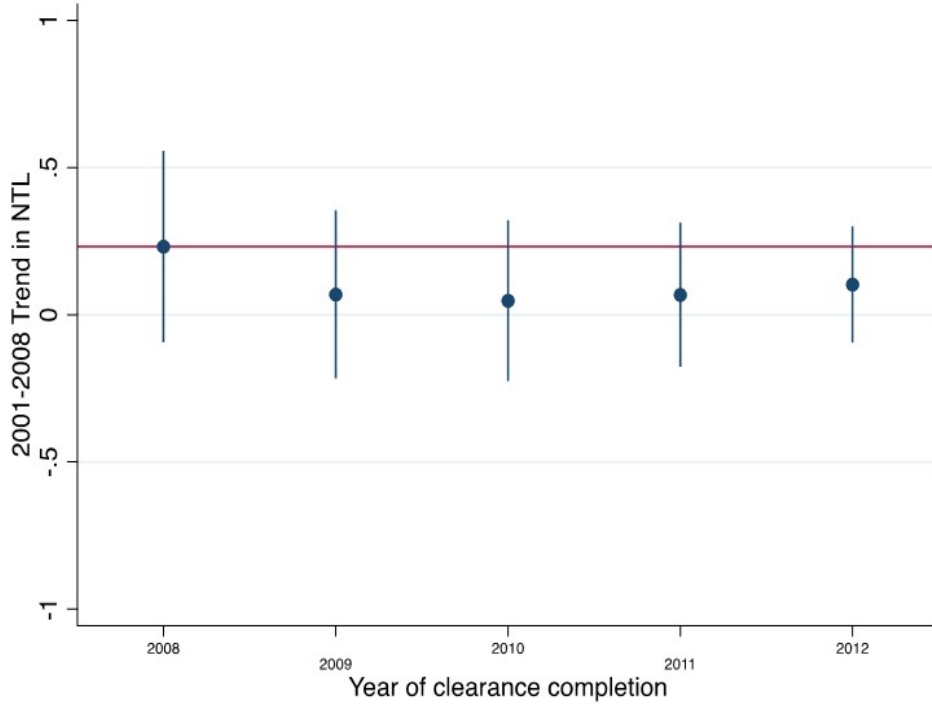
We estimate the effect of a grid cell being entirely cleared on NL using ordinary least squares (OLS) regressions with two-way FEs at the grid cell and province-by-year levels. The estimating equation for this analysis is:

$$NL_{jpt} = \alpha + \beta Cleared_{jpt} + \delta_j + \delta_{pt} + \epsilon_{jpt}$$

where NL_{jpt} represents the intensity of nighttime light emission from grid cell j in province p and year t . $Cleared_{jpt}$ is a binary variable taking a value of 1 if grid cell j is entirely cleared of landmines by year t , and 0 otherwise. δ_j is a vector of grid-cell fixed effects, which capture time-invariant local characteristics. δ_{pt} is a vector of year-by-province fixed effects, which capture geographically invariant characteristics within each year-by-province combination. We use two-way clustering of standard errors at the district and year level following [Cameron, Gelbach and Miller \(2011\)](#). Because we rely only on the within-grid cell variation in timing of clearance, we limit our sample to only grid cells cleared between 1992 and 2013 (i.e., we exclude cells cleared after 2013, when our NL data series ends).

This empirical strategy relies on the parallel trends assumption: individual grid-cells cleared earlier in our analysis window would have experienced similar trends to the not-yet-treated grid-cells had they not been cleared. As our FEs adjust for province-by-year-specific means, such differential trends would have to exist for earlier-cleared grid-cells compared to not-yet-cleared cells within the same province. To validate this assumption, we first use the long time-series of NL data preceding treatment to confirm that there were not differential pre-trends for earlier-cleared grid-cells. We estimate a linear model over the 2001-2008 time window with continuous time trends specific to each distinct year of clearance completion (along with the grid-cell and province-by-year fixed effects used in our main specification). The results, plotted in Figure 5, show little evidence of systematically different pre-trends based on the year of clearance completion.

Figure 5: Pre-trends in Nighttime Lights



We further validate our empirical strategy using recently developed approaches accounting for potential biases in panel models with two-way FEs. These biases arise when the treatment effects vary over time or over other characteristics, and can cause the standard two-way FE models to diverge substantially from the intended interpretation as the average effects comparing already treated to not-yet-treated units. We adopt the approaches of both [Chaisemartin and d’Haultfoeuille \(2020\)](#) and [Callaway and Sant’Anna \(2021\)](#) to account for these biases. The event study results produced by these approaches (shown in Section 5) also confirm the lack of differential pre-trends, giving further credence to our identifying assumption.

4.2 Land use

To estimate the impact of clearance on land use, we utilize our data on LULC categories derived from RF algorithms based on MS for five-year time intervals from 1995 to 2020. We generate this data for 30 m x 30 m grid-cells. In order to assign land use classifications to entire hazardous areas, we calculate the percentage of each hazardous area devoted to distinct land use categories. As such, the unit of observation is the hazardous-area-polygon-by-year level.

The estimating equation for this analysis is:

$$LandUse_{uip} = \alpha_u + \beta_u Cleared_{ipt} + \delta_{ui} + \delta_{upt} + \epsilon_{uip}$$


where $LandUse_{uip}$ represents the share of the hazardous area polygon i in province p

classified as in use for land use u in year t . We conduct the analysis separately for each land use category u , so obtain specific parameters for each category. β_u is our parameter of interest, denoting the treatment effect associated with $Cleared_{jpt}$, a binary variable taking a value of 1 if hazardous area i is cleared of landmines at year t , and 0 otherwise. δ_{ui} is a vector of hazardous area FEs, which capture time-invariant local characteristics. δ_{upt} is a vector of province-by-year FEs, which capture geographically invariant characteristics within each year-by-province combination. We cluster standard errors by district and year.

Our land use analysis leverages the finer units of analysis of the specific hazardous area polygons, as well as the extended time window of our analysis sample, which now goes up to 2020. This allows us to capture both longer tailed impacts as well as those occurring in unelectrified rural areas. Moreover, our treatment measure now reflects the clearance of the full hazardous area observation (rather than the final clearance of a grid cell, where intervening clearance of some areas may have already occurred).


As with our NL analysis, our land use analysis is predicated on the parallel trends assumption. We thus also document that there are no differential changes in land use categories between 2000 and 2008 that correlate with the timing of each hazardous areas' clearance. Figure 6 shows the pre-trends in built-up land use, while Figure 7 shows the pre-trends in farmland land use.⁷

Figure 6: Pre-trends in Built-Up Land



⁷Pre-trends for additional land use classifications can be found in the appendix.

Figure 7: Pre-trends in Farmland



The main challenge of the land use analysis is that the computational demands of processing all of the Landsat imagery limit us to conducting our analysis at five year rather than annual intervals. As such, event study-based approaches such as those we adopt in our NL analysis are not feasible for our land use analysis. However, given the apparent robustness of our standard two-way FE model in our NL estimations, we expect our land use results to be similarly valid.

4.3 Conflict

We similarly leverage the roll-out of clearance activities to estimate the impact of clearance on subsequent conflict. Our data here are at more aggregate 10km grid cell scales, but our time window now stretches until 2020 at annual frequency.

The estimating equation for this analysis is:

$$Conflict_{kpt} = \alpha + \beta Cleared_{kpt} + \delta_k + \delta_{pt} + \epsilon_{kpt}$$

where $Conflict_{kpt}$ is a count of conflict events in grid-cell k in province p during year t . $Cleared_{kpt}$ is a binary variable taking a value of 1 if grid cell k is entirely cleared of landmines at year t , and 0 otherwise. δ_k is a vector of grid-cell FEs, which capture time-invariant local characteristics. δ_{pt} is a vector of province-by-year FEs, which capture geographically invariant characteristics within each year-by-province combination. As the area of each grid-cell intersecting with hazardous areas varies, this analysis is weighted by the percent of grid cell k covered by hazardous areas.

5 Results

5.1 Nighttime Lights

We begin by estimating the effect of landmine clearance on economic development among hazardous areas cleared during or after 2008, using NL. First we examine the event-study results using the [Chaisemartin and d’Haultfoeuille \(2020\)](#) dynamic estimator, shown in Figure 8. Confirming our primary assumption, we see little evidence of meaningful changes in NL prior to the full clearance of a grid cell. However, NL output seems to increase consistently with each additional year after clearance. These effects are quite large, with increases of more than the pre-treatment mean NL emissions after only 3 years, suggesting that clearance of hazardous areas can rapidly alter the trajectory of economic output and conditions.

Figure 8: Event Study of Clearance Impacts on NL
Dynamic Estimator ([Chaisemartin and d’Haultfoeuille, 2020](#))





Figure 8 also indicates dynamic treatment effects that mount over time (rather than a one-time change in NL levels post-treatment that remains constant over time thereafter). This finding also validates our use of the dynamic estimator, as we would expect the treatment effects to be otherwise attenuated by comparisons of changes in treatment among late-treated cells with means among already-treated cells (whose NL levels would be continuing to rise).

Multiple approaches to addressing the potential bias from dynamic and heterogeneous treatment effects are now available. We thus also utilize the [Callaway and Sant’Anna \(2021\)](#) estimator, which makes a related but distinct set of identifying assumptions to recover unbiased event study results in the presence of dynamic treatment effects. The results, shown below in Figure 9, are nearly identical to our prior findings, with little evidence of upward trends in NL preceding clearance and mounting effects beginning shortly after clearance.

Figure 9: Event Study of Clearance Impacts on NL
Dynamic Estimator (Callaway and Sant’Anna, 2021)



Nonetheless, we also use (non-dynamic) two-way FE estimation to explore heterogeneity in the effects of clearance on NL for the sample cleared between 2008 and 2013. Table 3 first shows that we continue to find significant impacts of clearance on NL even in this specification (Column 1). We then examine how these effects vary by road access⁸ and population density.⁹ The impacts of clearance are strongest among the cells located closest to a road (Column 2), and diminish to approximately zero at more than 4.5 km away from a road. In Column 3, we also find that more populated areas experience larger NL impacts from clearance. While the effects are largest in areas with the greatest initial population levels, the estimates imply relatively large effects even in low population areas. For example, these estimates imply that areas with population densities of only 200 inhabitants/km² (classified as “rural” in most schema) see NL gains equal to 100% of the control mean (0.459), effectively doubling their NL due to clearance.

⁸We construct a *Distance to Roads* measure that reflects the distance from the grid cell to the nearest road on the United Nations Office for the Coordination of Humanitarian Affairs (OCHA) road network map.

⁹To reflect the baseline population density for the broader area in which each grid cell is located, we use the district mean population density in the year 2000 from the CIESIN Gridded Population of the World Version 4.

Table 3: Clearance Impacts on NL

DV=Nighttime Lights	(1)	(2)	(3)
Cleared of Landmines	.699*** (0.209)	1.047*** (0.296)	0.0801 (0.162)
Cleared * Dist. to Road		-0.230** (0.0825)	
Cleared * Baseline Pop.			0.00220*** (0.000538)
Observations	121,968	121,968	121,968
R-squared	0.627	0.629	0.646
Control Mean	0.459	0.459	0.459

*** p<0.01, ** p<0.05, * p<0.1. All models include grid cell and province-year fixed effects, and are weighted by percent cell covered by hazardous area. Standard errors clustered by district and year in parentheses.

Taken together, these NL impacts suggest important gains in economic conditions around a hazardous area once it is cleared. Since a hazardous area is at most $\sqrt{2}\text{km} \cong 1\text{mi}$ away from any point within the same grid cell, these effects suggest material economic changes even at reasonably small distances from cleared hazardous areas. Finding important nearby gains from clearance of hazards even in sparsely populated areas also contrasts with [Chiovelli, Michalopoulos and Papaioannou \(2018\)](#), who find primarily broader effects and ones primarily from clearing major roadways.

5.2 Land Use

We next consider the extent to which clearance of the hazardous areas altered the land use of these specific polygons using classes derived from MS imagery. In Table 4, we show impacts on built-up land use (reflected as the share of the polygon classified as built-up). We find significant increases in built-up land use after clearance (Column 1), equal to an increase of approximately 2 percentage point (over a sample mean of $\approx 11\%$). These are meaningful changes, given that construction of new built-up uses or expansion of existing ones can reflect larger investment in a context of challenging economic and security conditions.

Another advantage of the land use classification at this scale is that it enables us to unpack these effects by the types of blockages in each hazardous area. In Columns 2-7 of Table 4, we show separate estimations for sub-samples of polygon blockage types as designated by the clearance teams and recorded in the DMAC database (grazing, agriculture, road, housing infrastructure and water blockages). We find the gains are concentrated in polygons facing grazing and agriculture blockages, which comprise the vast majority of the sample. We find relatively few gains in built up use in polygons facing road, infrastructure, or housing blockages. At first glance, this may seem counter-intuitive, as grazing and agricultural blockages occur in areas that are not as built-up on average as other blockage types (such as housing and infrastructure). However, clearing landmine hazards from these areas could

allow the expansion of built-up uses to a greater degree than increasing density of built-up uses in already heavily built-up areas. The blockages initially recorded by the clearance teams were predictions, but clearing grazing and agricultural areas could have led to a broader set of uses that were not originally envisioned (in this case, particularly built-up uses). These effects appear particularly important in villages and small towns where grazing and agricultural blockages most frequently occur (and which comprise the vast majority of our sample).

Table 4: Impacts on Built-up Land Use

DV = Pct. Built-Up	(1) All	(2) Grazing	(3) Agriculture	(4) Road	(5) Housing	(6) Infrastructure	(7) Water
Cleared	0.0227** (0.00728)	0.0212** (0.00756)	0.0326 (0.0171)	0.0114 (0.0122)	0.00983 (0.0202)	-0.00398 (0)	0.0794** (0.0293)
Observations	33,972	26,874	6,366	2,604	2,130	720	582
R-squared	0.656	0.664	0.641	0.649	0.684	0.702	0.740
Hazard FEs	Y	Y	Y	Y	Y	Y	Y
Year*Prov. FEs	Y	Y	Y	Y	Y	Y	Y
Pre-clearance builtup	.11	.104	.113	.101	.16	.308	.149

Column headings reflect blockage type subsamples
Standard errors in parentheses clustered by district and year
*** p<0.01, ** p<0.05, * p<0.1

As a corollary to the increasing prevalence of built-up uses on newly cleared lands originally classified as blocked for grazing and agricultural uses, we investigate whether these areas exhibit increasing farmland uses. In Table 5, we show impacts of clearance on the share of each polygon classified as a farmland use. In Column 1, we find no material effects on the full sample of polygons. In Columns 2 and 3, we repeat this estimation for subsamples of polygons categorized as grazing and agricultural blockages, again finding no evidence of impacts of clearance on the reclassification of these areas into farmland uses. It is possible that the changes in farming on these lands were not dramatic enough to be observable as changes in the land use classification. For example, in some cases, farmers may have already attempted to sow and harvest wheat even in hazardous areas, and intensified their planting after clearance (and thus be classified as farmland both pre- and post-clearance). However, such examples are likely outliers, as landmine hazards posed major risks and many of the polygons designated as agricultural blockages are not classified as farmland uses at baseline. A more likely explanation may be that water availability and access to fertilizer and other inputs may be the binding constraints to farming expansion in this context.

Table 5: Clearance Impacts on Farmland Uses

DV=Pct. Farmland	(1) All	(2) Grazing	(3) Agriculture
Cleared	-0.00488 (0.00553)	-0.00236 (0.00347)	0.00306 (0.0119)
Observations	33,972	22,332	3,534
R-squared	0.657	0.605	0.702
Hazard FEs	Y	Y	Y
Year*Prov. FEs	Y	Y	Y


Column headings reflect blockage type subsamples
Standard errors in parentheses clustered by district and year
*** p<0.01, ** p<0.05, * p<0.1

Taken together, these results indicate important shifts in investments on already built-up lands where landmines restricted their use. Again, these occur even in small towns and villages throughout Afghanistan, and are not restricted to major roadways. At the same time, we do not observe farmers returning to use newly cleared agricultural land, suggesting other constraints may have limited these responses. In the next section, we turn to the potential role of conflict risks as one such constraint.

5.3 Conflict


We begin by examining the impacts of clearance on the frequency of conflict events in the 10km grid cell into which each hazardous area falls. In Figure 10, we observe no evidence of significant trends in conflict leading up to the clearance year, further validating our causal identification assumption. After clearance, however, we see large, significant drops in the frequency of conflict events. These reductions are statistically distinguishable two years after the clearance, and appear to generally grow in magnitude over time.

Figure 10: Event Study of Clearance Impacts on Conflict Dynamic Estimator (Chaisemartin and d'Haultfoeuille, 2020)



These results are quite consistent with those for NL, exhibiting mounting improvements after clearance that grow with time. We also consider the alternative approach of Callaway and Sant'Anna (2021) in Figure 11, finding very similar results.

Figure 11: Event Study of Clearance Impacts on Conflict
Dynamic Estimator (Callaway and Sant’Anna, 2021)



We next consider how conflict intersects with NL and land use conditions. We first return to our NL grid cell sample and incorporate both contemporaneous and baseline conflict conditions in the 5km grid cell into which each 1km NL grid cell falls. We split our sample into high and low baseline conflict subsamples (i.e., above/below median baseline conflict), with results for the former in Table 6 and latter in Table 7. In Columns 1 of each respective table, we find quite similar estimates of clearance impacts on NL (if anything, effects are slightly larger in the “high baseline conflict” subsample). These suggest that landmine clearance can be an effective development strategy even in areas that have experienced substantial conflict risks.

We next investigate whether the clearance impacts on satellite-based outcomes are correlated with the impacts on conflict risks. This can be framed as asking whether the same locations experience impacts on both dimensions, or whether these improvements are independently distributed across locations. To do so, we include nearby contemporaneous conflict as a covariate (recognizing that these are themselves endogenous outcomes of the clearance measure). If the impacts on both dimensions are geographically clustered (i.e., the same locations largely experience both gains in NL and conflict reductions), we would expect the coefficient on clearance to be attenuated toward zero.

However, when we add contemporaneous conflict as a covariate in both models in Columns 2 of Tables 6 and 7, we see no material difference in the coefficient on landmine clearance from the base models. At the same time, we do find significant negative correlations between contemporaneous conflict and NL outcomes. Taken together, these suggest that the

geographic distribution of NL gains from clearance is likely to be independent of the distribution of conflict reductions. This also implies that in this context, clearing hazardous areas has distinct, direct effects on both economic activity and conflict (rather than one of these outcomes serving as pathway for effects on the other).

Finally, we also consider whether contemporaneous conflict could dampen the impacts on economic outcomes (proxied by NL). Although we find in Columns 1 that baseline conflict does not dampen such impacts, this may be because the baseline measures reflect conditions occurring long before a later observation. We thus also interact clearance status with contemporaneous conflict measures in the third columns of the aforementioned tables. In high baseline conflict settings, we find significant interactions between contemporaneous conflict and clearance, suggesting that the largest gains from clearance in our sample actually occur in areas with the worst conflict conditions (both current and prior conflict).

Table 6: Impacts on NL Among High Baseline Conflict

	(1)	(2)	(3)
	NTL	NTL	NTL
Cleared of Landmines	1.210*** (0.416)	1.195*** (0.400)	0.508 (0.450)
Conflict (5km)		-0.143** (0.0607)	-0.168*** (0.0469)
Cleared * Conflict(5km)			0.330** (0.104)
Observations	34,738	34,738	34,738
R-squared	0.734	0.741	0.745
Year FEs	N	N	N
Grid Cell FEs	Y	Y	Y
Year*Prov. FEs	Y	Y	Y

Standard errors in parentheses clustered by district and year

Weighted by percent cell covered by hazardous area.

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

Table 7: Impacts on NL Among Low Baseline Conflict

	(1)	(2)	(3)
	NTL	NTL	NTL
Cleared of Landmines	.458*** (0.158)	.465*** (0.159)	.446*** (0.157)
Conflict (5km)		-0.0763** (0.0351)	-0.138** (0.0664)
Cleared * Conflict (5km)			0.0960 (0.0783)
Observations	87,230	87,230	87,230
R-squared	0.572	0.572	0.573
Year FEs	N	N	N
Grid Cell FEs	Y	Y	Y
Year*Prov. FEs	Y	Y	Y

Standard errors in parentheses clustered by district and year

Weighted by percent cell covered by hazardous area.

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

Finally, we also find similar interactions in effects on built up land uses, shown in Table 8. In the high baseline conflict sub-sample, clearance raises the incidence of built-up land uses, especially in those locations with the greatest conflict risks (results among low conflict locations are consistent but omitted for brevity).

Table 8: Clearance Impacts on Land Use Among High Conflict

	(1)	(2)
	Built-Up	Farmland
Cleared	0.0317 (0.0220)	-0.000944 (0.00892)
Conflict(2km)	0.0116*** (0.000363)	.000147 (0.00210)
Cleared * Conflict(2km)	0.0126** (0.00310)	0.00399 (0.00361)
Observations	4,175	4,175
R-squared	0.620	0.788

Standard errors in parentheses clustered by district and year

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

Taken together, these results suggest clearing hazardous areas where conflict risks are most pronounced can both improve economic conditions in and near these areas. This

clearance can also reduce the likelihood of further conflict in these areas. Both dimensions of improvements are important from a welfare perspective. Although it is possible that conflict reductions due to clearance could have knock-on effects on subsequent economic activity, we find no evidence of clustering of these outcomes, suggesting improvements in one dimension are not necessarily dependent on improvements in the other.

6 Conclusions

We offer the first evidence that clearing areas potentially contaminated with landmines can reshape economic activity even in the midst of ongoing conflict. The limited work on landmine clearance impacts to date has focused on post-conflict settings, finding the greatest gains in clearing transport routes linking major markets. Using spatially precise clearance data and satellite imagery, we find that clearing areas in rural settings can also spur new development. The largest changes appear to be the improvement and expansion of buildings and small-scale infrastructure in small towns and villages. These shifts are particularly pronounced in locations with the greatest conflict risks.

At the same time, landmine clearance also appears to reduce the risks of further conflict directly. This may be because landmine explosions can trigger subsequent violence in response, or because they weaken the broader sense of safety and rule of law in an area. Whatever the pathway, landmine clearance helps to avert further risks beyond those due to the landmine themselves.

Clearing a hazardous area of potential landmines remains a painstaking and dangerous task, and the sector remains underfunded given the geographic scope of areas still considered hazardous (even ignoring newly laid landmines). Although a full benefit-cost analysis is beyond the scope of our study, our results suggest there are important economic gains beyond those represented by the averted direct harm of the mines to the people they injure or kill. Our work also suggests that more equitably targeting rural and urban areas for landmine clearance is warranted. Given the aforementioned budget constraints, more urban areas and major transport corridors are often given precedence both because their clearance is believed to induce the largest gains and because they are often easier to reach. Our results suggest a rebalancing may be at least partly in order, as there are important gains from clearance even in small towns and villages.

Finally, our findings point to some important areas for future research. First, although we do not observe changes in whether an area is farmed, there may be finer changes in the intensity and productivity of farming activities that could be detectable with MS-based data. Second, as climate change induces new pressures on economic activity—particularly in conflict-affected settings—the gains from landmine clearance may be magnified or dampened. A further stream of research may therefore tackle the intersection of clearance impacts and climate conditions.

References

- ACLED.** 2022. “10 Conflicts to Worry About in 2022: Afghanistan.”
- Arcand, Jean-Louis, Aude-Sophie Rodella-Boitreaud, and Matthias Rieger.** 2015. “The impact of land mines on child health: evidence from Angola.” *Economic Development and Cultural Change*, 63(2): 249–279.
- Bluhm, Richard, and Gordon C. McCord.** 2022. “What can we learn from nighttime lights for small geographies? measurement errors and heterogeneous elasticities.” *Remote Sensing*, 14(5): 1190.
- Breiman, Leo.** 2001. “Random Forests.” *Journal of Machine Learning*, 45: 5–32.
- Callaway, Brantly, and Pedro HC Sant’Anna.** 2021. “Difference-in-differences with multiple time periods.” *Journal of Econometrics*, 225(2): 200–230.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller.** 2011. “Robust inference with multiway clustering.” *Journal of Business Economic Statistics*, 29(2): 238–249.
- Chaisemartin, Clément De, and Xavier d’Haultfoeuille.** 2020. “Two-way fixed effects estimators with heterogeneous treatment effects.” *American Economic Review*, 110(9): 2964–2996.
- Chawla, Shailini.** 2011. “Diffusion of Landmines in Afghanistan.” *Strategic Analysis*, , (Vol. XXIV No.3).
- Chen, Xi, and William D. Nordhaus.** 2010. “No title.” *The value of luminosity data as a proxy foreconomic statistics*.
- Chiovelli, Giorgio, Stelios Michalopoulos, and Elias Papaioannou.** 2018. “Landmines and spatial development.”
- Crowley, Michael, and John Ismay.** 2022. “Biden Bans Most Antipersonnel Land Mine Use, Reversing Trump-Era Policy.”
- Dugoua, Eugenie, Ryan Kennedy, and Johannes Urpelainen.** 2018. “Satellite data for the social sciences: measuring rural electrification with night-time lights.” *International Journal of Remote Sensing*, 39(9): 2690–2701.
- FAO, Afghanistan, Food- Organization.** 2016. ““The Islamic Republic of Afghanistan: Land Cover Atlas”.”
- Fjelde, Hanne, and Lisa Hultman.** 2014. “Weakening the enemy: A disaggregated study of violence against civilians in Africa.” *Journal of Conflict Resolution*, 58(7): 1230–1257.

- Frost, Alexandra, Peter Boyle, Philippe Autier, Colin King, Wim Zwijnenburg, David Hewitson, and Richard Sullivan.** 2017. “The effect of explosive remnants of war on global public health: a systematic mixed-studies review using narrative synthesis.” *The Lancet Public Health*, 2(6): e286–e296.
- Henderson, J. Vernon, Adam Storeygard, and David N. Weil.** 2012. “Measuring economic growth from outer space.” *American economic review*, 102(2): 994–1028.
- ICBL, (International Campaign to Ban Landmines).** 2021. “LANDMINE MONITOR 2021 23 RD ANNUAL EDITION.”
- Jean, Neal, Marshall Burke, Michael Xie, W. Matthew Davis, David B. Lobell, and Stefano Ermon.** 2016. “Combining satellite imagery and machine learning to predict poverty.” *Science*, 353(6301): 790–794.
- Mellander, Charlotta, José Lobo, Kevin Stolarick, and Zara Matheson.** 2015. “Night-time light data: A good proxy measure for economic activity?” *PloS one*, 10(10): e0139779.
- Merrouche, Ouarda.** 2008. “Landmines and poverty: IV evidence from Mozambique.” *Peace Economics, Peace Science and Public Policy*, 14(1): 23–38.
- Merrouche, Ouarda.** 2011. “The long term educational cost of war: evidence from landmine contamination in Cambodia.” *The Journal of Development Studies*, 47(3): 399–416.
- Michalopoulos, Stelios, and Elias Papaioannou.** 2016. “The long-run effects of the scramble for Africa.” *American Economic Review*, 106(7): 1802–1848.
- Monitor, Landmine Cluster Muniton.** 2009. “Afghanistan.”
- Pal, Mahesh, and Paul M. Mather.** 2003. “An assessment of the effectiveness of decision tree methods for land cover classification.” *Remote Sensing of Environment*, 86: 554–565.
- Paterson, Ted.** 2012. “Transitioning Mine Action Programmes to National Ownership.”
- SIGAR.** 2021. “Quarterly Report to the United States Congress.” Special Inspector General for Afghanistan Reconstruction.
- Sundberg, Ralph, and Erik Melander.** 2013. “Introducing the UCDP georeferenced event dataset.” *Journal of Peace Research*, 50(4): 523–532.
- Thompson, Allyson L., and Bernard E. Hubbard.** 2014. “A comprehensive population dataset for Afghanistan constructed using GIS-based dasymetric mapping methods: U.S. Geological Survey Scientific Investigations Report.”
- Uexkull, Nina Von, Mihai Croicu, Hanne Fjelde, and Halvard Buhaug.** 2016. “Civil conflict sensitivity to growing-season drought.” *Proceedings of the National Academy of Sciences*, 113(44): 12391–12396.

UNMAS. 2021. “30 Years of Impact: An Evaluation of the Mine Action Programme of Afghanistan.”

UNMAS. 2022. “Afghanistan.”

Wagner, Zachary, Sam Heft-Neal, Zulfiqar A. Bhutta, Robert E. Black, Marshall Burke, and Eran Bendavid. 2018. “Armed conflict and child mortality in Africa: a geospatial analysis.” *The Lancet*, 392(10150): 857–865.

Yeh, Christopher, Anthony Perez, Anne Driscoll, George Azzari, Zhongyi Tang, David Lobell, Stefano Ermon, and Marshall Burke. 2020. “Using publicly available satellite imagery and deep learning to understand economic well-being in Africa.” *Nature communications*, 11(1): 1–11.