Online Appendix

Illuminating the World Cup Effect: Night Lights Evidence from South Africa

Gregor Pfeifer *
University of Hohenheim

Fabian Wahl †
University of Hohenheim

Martyna Marczak ‡
University of Hohenheim

June 28, 2018

*Corresponding author: University of Hohenheim, Department of Economics (520B), D-70593 Stuttgart, Germany; E-mail: g.pfeifer@uni-hohenheim.de; Phone: +49 711 459 22193.

†University of Hohenheim, Department of Economics (520J), D-70593 Stuttgart, Germany; E-mail: fabian.wahl@uni-hohenheim.de; Phone: +49 711 459 24405.

‡University of Hohenheim, Department of Economics (520G), D-70593 Stuttgart, Germany; E-mail: marczak@uni-hohenheim.de; Phone: +49 711 459 23823.

A Data and Setting Details

A.1 Details on Luminosity Data

Luminosity data

The data are made available by the National Geophysical Data Center (NGDC) of the National Oceanic and Atmospheric Administration of the U.S., and originate from images taken by satellites of the Defense Meteorological Satellite Program (DMSP) of the U.S. Department of Defense. We use shapefiles containing the average visible, stable nighttime lights and cloud-free coverage, where ephemeral events (fires, etc.) as well as background noise are removed and only light from sites with persistent lightning is included. A common problem with these data is that the light of gas flares is included and could be mistaken for lights of settlements. However, in the area studied in this paper, no gas flares exist.

Night light intensity is measured by an integer ranging from 0 (unlit) to 63. Night light intensity data are available on pixel (grid cell) level, with each pixel corresponding to 30x30 arc seconds, i.e., one value represents the average night light intensity of an area of 0.86 square kilometer (on the equator). This high level of spatial resolution allows to study economic effects of highly localized investments like those related to urban infrastructure (Gonzalez-Navarro and Turner, 2016), natural disasters (Elliot et al., 2015) or natural amenities like surfing spots (McGregor and Wills, 2017).

In what follows, we describe problems coming along with the use of luminosity data and explain how we addressed them.

Too small wards

From the entire $4,277$ wards, we are able to use $4,222$ throughout the empirical analysis as some wards are too small to calculate exact luminosity or elevation values.

Luminosity measured by different satellites over time

When using the luminosity series pooled over time, a potential issue is that a portion of the temporal variation can be due to the fact that these data are collected by different satellites, which are not calibrated on a common level. To make luminosity values more comparable over satellites, the data can be inter-calibrated manually following, e.g., a procedure suggested by Elvidge et al. (2009). However, Chen and Nordhaus (2011) find that results only marginally change when using inter-calibrated luminosity. In line with this, after inter-calibrating our data set based on the values given by Elvidge et al. (2014), we find a correlation with the original data of 0.991. In fact, corresponding results would even suggest a slightly stronger effect as compared to our baseline results provided by Figure 3 in Section 4 of the paper; see Figure C.6 in Online Appendix C.1.

Pixels with zero luminosity

Especially within the municipalities located in the western South African wastelands, there exist pixels which are zero throughout the entire sampling period. In one of the robustness checks in Online Appendix C, we follow Elliot et al. (2015) and remove such pixels prior to the aggregation on municipal level. Corresponding estimation results stay almost exactly identical to our baseline ones including the zero pixels.

Alternative luminosity measures

In our analysis, we use untransformed levels of luminosity as dependent variable. Alternatively, we could draw on luminosity per capita. However, population figures on municipality or ward level are only available for a very few years throughout our observation period, which would lead to distinctly reduced number of observations. The same problem occurs with alternative luminosity data—the so-called 'radiance calibrated at night data' (available at: `http://ngdc.noaa.gov/eog/dmsp/download_radcal.html`, last accessed on January 19, 2017), that have been used by, e.g., Gonzalez-Navarro and Turner (2016). These data do not feature the cap at 63, but comprise seven cross sections, out of which only three correspond to the analyzed post-treatment time span 2004–2013 (2004, 2006, 2011).

A.2 Description of Variables

Area

Area of a municipality in square kilometers. Calculated with ArcGIS Pro "Calculate Geometry" tool according to the shapefile of municipality territories from the DIVA-GIS database. It can be downloaded at: `http://biogeo.ucdavis.edu/data/diva/adm/ZAF_adm.zip` (last accessed on January 19, 2017).

Distance to railway

Geodesic distance from the centroid (mid-point) of each municipality to the next railway in kilometers. The course of the railways in South Africa is taken from a shapefile provided by the DIVA-GIS database of UC Davies. It can be downloaded at: `http://biogeo.ucdavis.edu/data/diva/rrd/ZAF_rrd.zip` (last accessed on January 22, 2017). The centroids are calculated using the ArcGIS Pro "Calculate Centroids" tool.

Elevation

Elevation data are taken from the NASA SRTM (Shuttle Radar Topographic Mission) data set (Version 4.0). The spatial resolution of the data is 3 arc seconds (approx. 90m on the equator). The data is provided by the Consortium for Spatial Information (CGIAR-CSI). A detailed description can be found at: `http://srtm.csi.cgiar.org/` and here `http://www.cgiar-csi.`

org/data/srtm-90m-digital-elevation-database-v4-1#introduction. It can be downloaded from this webpage: http://srtm.csi.cgiar.org/SELECTION/inputCoord.asp (last accessed on January 21, 2017).

Share of indigenous people
Share of a municipality's population with indigenous background. Data come from the project "Spatial Aspects of Unemployment in South Africa 1991–2011", conducted by the Human Sciences Research Council (HSRC). The data set covers the years 1991, 1996, 2001, and 2011. It merges information from the official South African Censuses in 1991, 1996, 2001, and 2011 and the Community Survey in 2007. The geographic units are defined according to the 2005 municipal boundaries. It can be downloaded (after free registration) at: http://curation.hsrc.ac.za/Dataset-342-datafiles.phtml (last accessed on January 22, 2017).

GDP per capita
GDP per capita (in Rand) on provincial level is calculated based on the official regional GDP estimates provided by Statistics South Africa (e.g., here http://www.statssa.gov.za/publications/P0441/P04413rdQuarter2009.pdf, last accessed on January 8, 2017). We take GDP estimates in current Rand and deflate them in prices of 2005 using annual inflation rates from the website: http://www.inflation.eu/inflation-rates/south-africa/historic-inflation/cpi-inflation-south-africa.aspx (last accessed on January 7, 2017). Finally, to obtain provincial GDP per capita, the deflated GDP is divided by provincial population. Provincial population figures come from the South African Census in 2011 (http://www.statssa.gov.za/census/census_2011/census_products/Provinces%20at%20a%20glance%2016%20Nov%202012%20corrected.pdf, last accessed on January 7, 2017).

Income per capita
Income per capita for the year 2007 originates from the HSRC data set "Spatial Aspects of Unemployment in South Africa 1991–2011" (see *GDP per capita*).

Share of tertiary education
Share of a municipality's population with a degree from a university or another educational institution in the tertiary sector. Data originate from the HSRC data set "Spatial Aspects of Unemployment in South Africa 1991–2011" (see *GDP per capita*).

Soil quality
Data on soil quality are taken from the Zabel et al. (2014) data set. The data set is described further at: http://geoportal-glues.ufz.de/stories/globalsuitability.html (last accessed on January 22, 2017). The measure used in the paper is the average agricultural suit-

ability over the period 1961–1990. To measure suitability, Zabel et al. (2014) consider climate (temperature, precipitation, solar radiation), soil (pH, texture, salinity, organic carbon content, etc.), and topography (elevation and slope) of an area. They consider rain-fed conditions and irrigation. To construct the suitability measure, they contrast these factors with growing requirements of 16 plants (Barley, Cassava, Groundnut, Maize, Millet, Oilpalm, Potato, Rapeseed, Rice, Rye, Sorghum, Soy, Sugarcane, Sunflower, Summer wheat, Winter wheat).

Unemployment rate
Share of a municipality's work force that is unemployed. Data originate from the HSRC data set "Spatial Aspects of Unemployment in South Africa 1991–2011" (see *GDP per capita*).

The variables of the ward level data set used in Section 4.2 of the paper originate from the same sources and are defined analogously to those of the municipality level data set. A descriptive overview of the data on municipality and district level (including luminosity) is provided in Tables A.1 and A.2, respectively.

Table A.1: Descriptive Overview of the Muncipality Level Data Set

Variable	Obs.	Mean	Std. Dev.	Min.	Max.
Luminosity	5,192	3.117	6.432	0	58.107
Area	5,192	5,203.189	5,395.449	251.434	36,127.950
Distance to Railway	5,192	2.045	6.854	0	44.842
Elevation	5,192	986.556	477.231	53.303	1,922.893
GDP per capita	702	26,896.73	9,828.787	12,986.54	56,768.83
Income per capita	237	15,781.66	11,986.76	1,113.61	82,477.9
Share Indigenous People	472	0.008	0.024	0	0.206
Share Tertiary Education	708	0.02	0.019	0	0.125
Soil Quality	5,192	36.214	13.419	0.023	64.255
Unemployment Rate	944	33.714	16.356	6.100	84.07

Table A.2: Descriptive Overview of the District Level Data Set

Variable	Obs.	Mean	Std. Dev.	Min.	Max.
Luminosity	946	2.788	5.145	0.043	32.469
Area	946	26298.92	26692.81	2531.568	126863.6
Distance to Railway	946	0	0	0	0
Elevation	946	1024.919	439.643	122.293	1673.966
Share Indigenous People	946	0.028	0.093	0	0.659
Share Tertiary Education	946	0.008	0.055	0	0.788
Soil Quality	946	35.167	12.564	0.667	51.476
Unemployment Rate	946	4.797	13.435	0	68.869

B Converting Night Lights into Standard Economic Outcomes

As has been mentioned in the main text, our study is the first which offers a conversion of night lights into the unemployment rate. As for GDP and income per capita, we could have drawn on conversion factors from the existing literature. However, for our analysis to be consistent, we have decided to use South African data for these two variables too. The results will show that our conversion factors are in line with those found in previous studies; see also the discussion in Section 3.2 of the paper. In the following, we provide regression equations allowing for the conversion of night lights into GDP and income per capita, as discussed for the unemployment rate in Section 3.2 of the main text.

GDP per capita is available for the years 1996, 2001, and 2007, and only on the province level. Thus, the respective conversion factor is computed based on the estimation of a specification similar to that for the unemployment rate (see eq. (4) in the main text), but adjusted accordingly since the unemployment rate was available at the municipality level. Consequently, GDP per capita and luminosity vary at the province level, while the control variables are still given at the municipality level:

$$GDP_{i,t} = \alpha + \beta Nightlights_{i,t} + \gamma' X_{ik} + \lambda_t + \epsilon_{ik,t}, \tag{B.1}$$

where $GDP_{i,t}$ and $Nightlights_{i,t}$ are GDP per capita and luminosity, respectively—in province i in year t. X_{ik} denotes the controls in province i and municipality k. Control variables include the area of a municipality in square kilometers, its distance to the closest railway, its average elevation, and soil quality. λ_t are year fixed effects introduced to account for the business cycle, and $\epsilon_{ik,t}$ is the error term.

Income per capita is available at the municipality level, but solely for the year 2007. Consequently, we cannot include year or municipality fixed effects when estimating the conversion factor from luminosity to income per capita. Thus, we estimate the following equation:

$$Income_{k,2007} = \alpha + \beta Nightlights_{k,2007} + \gamma' X_k + \theta_j + \epsilon_{k,2007}, \tag{B.2}$$

where $Income_{k,2007}$ is income per capita in municipality k, and in the year 2007. $Nightlights_{k,2007}$ denotes the corresponding night lights. θ_j are district fixed effects. X_k are, similarly as in eq. (B.1), control variables at the municipality level, and $\epsilon_{k,2007}$ is the error term.

The results of the OLS estimations for GDP/income per capita are summarized in Table B.1. For completeness, we also show the results for the unemployment rate.

Table B.1: Relationship between Night Lights and Economic Outcomes[a)]

Characteristics/	Dependent Variable		
Regressors	Unemployment Rate	GDP per capita	Income per capita
	(1)	(2)	(3)
Level of Variation	Municipality	Province	Municipality
Years in Sample	1996, 2001, 2007, 2011	1996, 2001, 2007	2007
Luminosity	−0.723**	1,400.699***	911.555***
	(0.282)	(38.543)	(270.819)
Conversion Factor[b)]	**−0.723**	**191.690**	**124.749**
District Dummies	No	No	Yes
Municipality Dummies	Yes	No	No
Year Dummies	Yes	Yes	No
Controls	No	Yes	Yes
Observations	936	702	234
R^2	0.819	0.555	0.509

a) Robust standard errors are reported in parentheses. Coefficient is statistically different from zero at the ***1%, **5%, and *10% level. Unit of observations in all three columns is a municipality. However, the dependent variable and luminosity do only vary on provincial level in Column (2). The set of controls includes a municipality's area in square kilometers, its distance to the closest railway, its elevation, and soil quality. Each regression includes a constant not reported.
b) Conversion factors related to GDP per capita and income per capita are obtained by translating the corresponding luminosity effects expressed in the national currency 'Rand' into U.S. dollars ($) using the average exchange rate of 2010.

C Robustness Checks

C.1 Analysis on the Municipality Level

Excluding the three largest World Cup Venues

We examine the robustness of our results by conducting several sensitivity checks. First, the average municipality result from Section 4.1 of the paper does not change if we compute our unit of interest as being the World Cup venue average of six municipalities instead of nine. The excluded municipalities are those with the highest luminosity level: Johannesburg Metropolitan Municipality, eThekwini Metropolitan Municipality, and Cape Town Metropolitan Municipality. Figure C.1 provides corresponding results, assuring that our findings are not driven by only a few locations which feature a high level of luminosity anyway.

Including venues with training facilities

Second, we try the opposite, namely including more 'treated' municipalities in addition to the nine in which the football stadiums were located at. These additional seven municipalities affected by World Cup measures are listed (along with the corresponding World Cup measures) in Table 1 of Section 2 in the companion document available at: `http://www.martynamarczak.com/static/pdf/SA_AllProjects.pdf`. We adjust our donor pool by excluding the additional neighboring municipalities according to the same rules as in the main analysis in Section 4 of the paper. Figure C.2 provides corresponding SCM results, showing that our findings stay robust and are in fact even stronger. Alternatively, we also find our results to be confirmed if we additionally include the border regions of Lesotho into the original donor pool an re-run the baseline SCM analysis.

'Leave-one-out' check

Third, we conduct a version of so-called 'leave-one-out' checks (Abadie et al., 2015), where a control unit that received positive weight during the original SCM optimization procedure is excluded from the donor pool. In our case, we pick the uMhlathuze Local Municipality which received a weight of over 80% when synthesizing the average World Cup municipality. In uMhlathuze, Africa's largest coal export facility (Richards Bay Coal Terminal) was constructed parallel and independent of the World Cup, a fact that could disqualify this region as a valid control unit. When removing this municipality from the donor pool, SCM still delivers robust results, which becomes evident when looking at Figure C.3. Now, the synthetic average World Cup venue is comprised of Msunduzi and Govan Mbeki, with corresponding w-weights of 84.5% and 15.5%, respectively.

Alternative predictor sets

Fourth, we try altering the predictor set. Particularly, we restrict the predictors to the same set as we had to work with on ward level (data set with time-invariant predictors), i.e., without the share of people with indigenous heritage, the share of people with tertiary education, and the percentage of people unemployed. We find our results to be confirmed. Adding more predictors, e.g., the respective distance to coast, does not affect our findings either. The same holds true if we additionally add population density in 1996 to the set of predictors. Results turn out to be almost identical. Hence, the estimated effects do not simply reflect the correlation between luminosity and agglomeration but can really be interpreted as being due to differences in wealth. Since the results corresponding to each of the three alternative predictor sets are nearly coincident with our baseline results, for the sake of completeness we provide only the plots of the synthetic unit compared to that of the treated unit (for all three alternative predictor sets), see Figure C.4.

Removing pixels with zero luminosity

Fifth, we follow Elliot et al. (2015) and remove all pixels from the data set that were unlit (0) over the entire sampling period, as one could assume that in the corresponding areas there was no economic activity at all. The fact that most of these pixels are located in sparsely populated municipalities in the western South African wastelands could lead to an underestimation of averaged night lights in the actually settled areas. However, when we re-run our SCM framework using the alternative luminosity measure where all such pixels were removed before the aggregation on municipality level, our results remain virtually identical to our baseline results. Similarly as in the case of alternative predictor sets, for the sake of completeness we provide only the comparison of the plot for the synthetic unit and the plot for the treated unit, see Figure C.5.

Inter-calibrated luminosity values

Lastly, we conduct SCM analysis with data after they have been inter-calibrated manually. For information on the reasons of such a procedure, see Section A.1. The results presented in Figure C.6 would even suggest a slightly stronger effect as compared to our baseline results provided by Figure 3 in Section 4 of the paper.

Taken together, all of the above-mentioned sensitivity checks support the credibility of our findings in Section 4.1 of the paper, revealing positive, short-run effects for the average World Cup municipality.

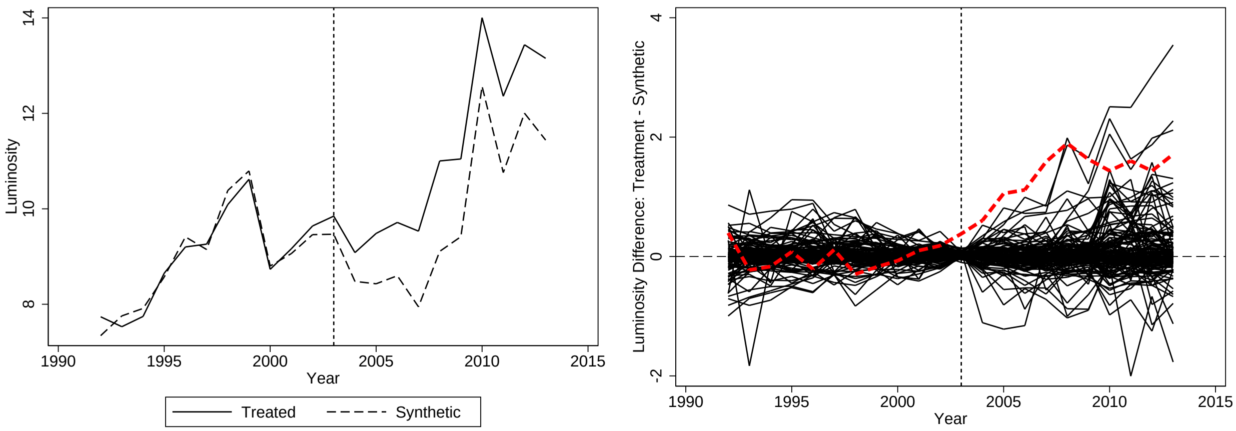

(a) Trends in Luminosity: Treated vs. Synthetic (b) Luminosity Gaps: Cup Venue vs. Placebos

Note: The vertical dashed line indicates the end of the pre-treatment period (2003). Panel (a) displays the average World Cup venue and its synthetic counterpart. Panel (b) plots luminosity gaps (treatment minus synthetic) for the average World Cup venue and placebo units: the black solid lines and the red dashed line represent the placebos and the treated unit, respectively.

Figure C.1: Estimation Results: The Average World Cup Venue w/o Johannesburg, Durban, and Cape Town

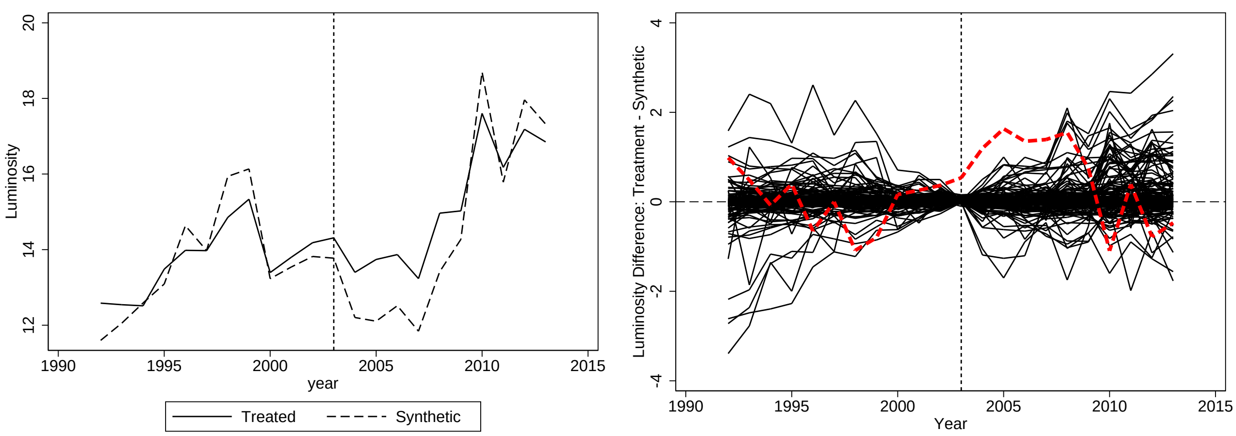

(a) Trends in Luminosity: Treated vs. Synthetic (b) Luminosity Gaps: Cup Venue vs. Placebos

Note: The vertical dashed line indicates the end of the pre-treatment period (2003). Panel (a) displays the average World Cup venue and its synthetic counterpart. Panel (b) plots luminosity gaps (treatment minus synthetic) for the average World Cup venue and placebo units: the black solid lines and the red dashed line represent the placebos and the treated unit, respectively.

Figure C.2: Estimation Results: The Average World Cup Venue including more Treated Municipalities

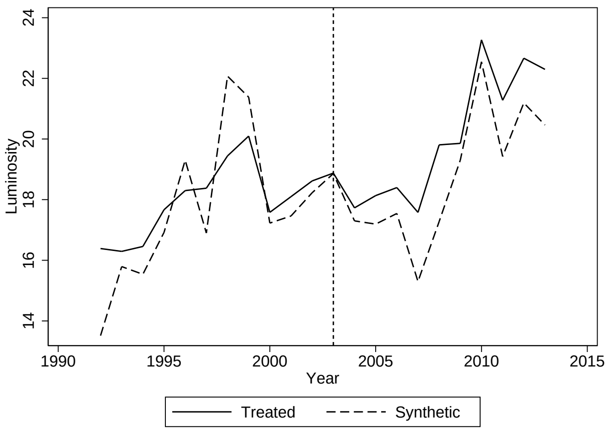

(a) Trends in Luminosity: Treated vs. Synthetic

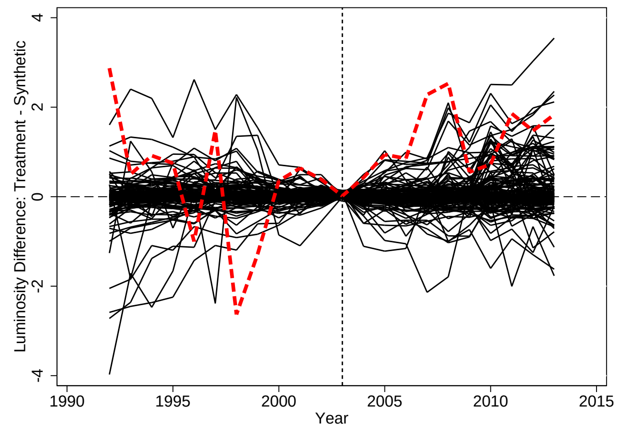

(b) Luminosity Gaps: Cup Venue vs. Placebos

Note: The vertical dashed line indicates the end of the pre-treatment period (2003). Panel (a) displays the average World Cup venue and its synthetic counterpart. Panel (b) plots luminosity gaps (treatment minus synthetic) for the average World Cup venue and placebo units: the black solid lines and the red dashed line represent the placebos and the treated unit, respectively.

Figure C.3: Estimation Results: The Average World Cup Venue without Control Unit 'uMhlathuze Local Municipality'

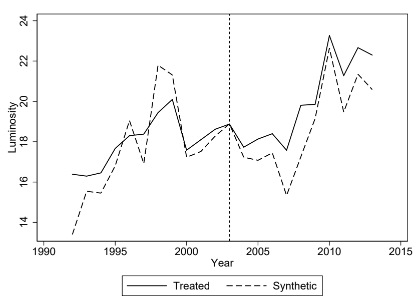

(a) Time-invariant Predictors

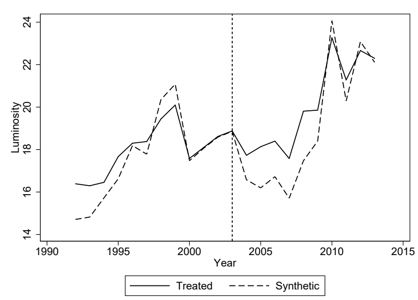

(b) Distance to Coast

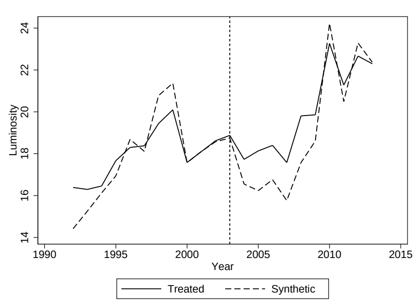

(c) Population Density

Note: Each subfigure shows the comparison of trends in luminosity of the treated unit (average World Cup venue) and the synthetic unit. The vertical dashed line indicates the end of the pre-treatment period (2003). Panel (a) corresponds to data with time-invariant predictors. Panel (b) corresponds to data with distance to coast as an additional predictor. Panel (c) corresponds to data with population density as an additional predictor.

Figure C.4: Estimation Results for the Average World Cup Venue: Three Different Alternative Predictor Sets

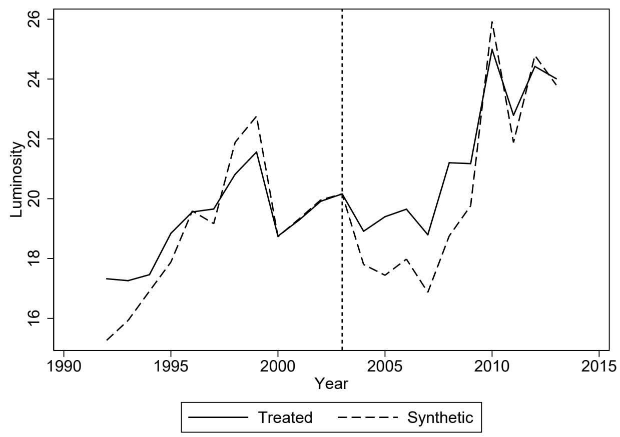

Note: The figure shows the comparison of trends in luminosity of the treated unit (average World Cup venue) and the synthetic unit. The vertical dashed line indicates the end of the pre-treatment period (2003).

Figure C.5: Estimation Results for the Average World Cup Venue: Data Set without Zero-Luminosity Pixels

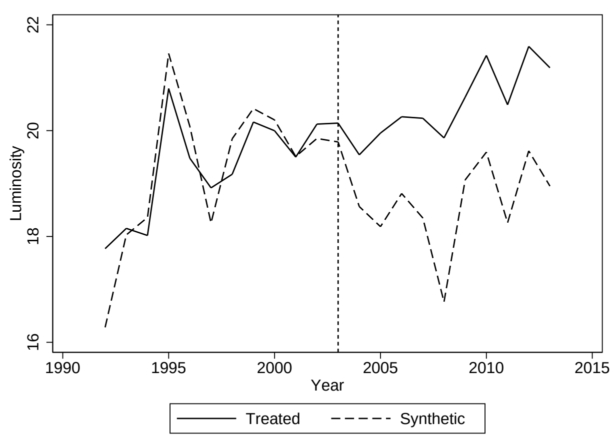

(a) Trends in Luminosity: Treated vs. Synthetic

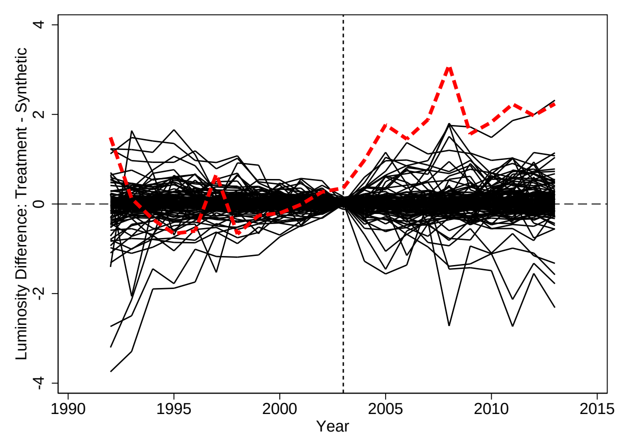

(b) Luminosity Gaps: World Cup Venue vs. Placebos

Note: The vertical dashed line indicates the end of the pre-treatment period (2003). Panel (a) displays the average World Cup venue and its synthetic counterpart. Panel (b) plots luminosity gaps (treatment minus synthetic) for the average World Cup venue and placebo units: the black solid lines and the red dashed line represent the placebos and the treated unit, respectively.

Figure C.6: Estimation Results for the Average World Cup Venue: Inter-calibrated Luminosity Values

C.2 Analysis on the Ward Level

Alternative treatment cutoff

In most cases of individual projects (that in our analysis on ward level are considered as individual treatments) we have the information about the project begin. Using this knowledge, we challenge our default treatment cutoff, which refers to the year 2004, to check how results might change if we use more precise information on the timing of the projects. As an example, we re-run our SCM analysis for Polokwane International Airport, employing the year 2008 as the beginning of the treatment period. Corresponding results are shown in Figure C.8, emphasizing that the generally defined treatment period from our initial SCM analysis has not led to misguided conclusions. Using the new, more precise cutoff, the airport effect is at least as pronounced as before. At its maximum, the luminosity gap relatively to the actual pre-treatment luminosity level amounts to 50%, notably stronger and more stable than for the case of the Peter Mokaba Stadium.

Another example we present here is the case of the Durban King Shaka International Airport. We shift the treatment cutoff to 2007 which was the actual beginning of the airport construction. From the evidence presented in Figure C.7 it can be inferred that the results look virtually identical to the baseline case which assumes the treatment to start in 2004. However, the figure also suggests that using 2007 as treatment year enables us to identify the actual onset of the increasing luminosity gap between the actual and synthetic airport ward much more precisely than before. For all examples shown in Section 4.2 of the paper, moving the cutoff towards later periods rather increases corresponding effects, and never undermines our previous findings. Thus, the results reported throughout the main analysis are always conservative lower bounds. Contrarily, moving the (placebo) cutoff to an early pre-treatment period, e.g., 1997, delivers zero-results as expected.

Analysis excluding lights produced by the treatment itself

Potential argument that can be raised in the context of our analysis on ward level is that the results might be driven by a World Cup project itself. It can be argued that, e.g., airports produce a lot of luminosity by themselves, e.g., through illuminated runways. To check if the investment projects generate effects for the remaining part of the ward, we exclude the luminosity pixels in the area affected by a project from our data set and re-run SCM. We consider two cases: the Durban King Shaka International Airport and the Polokwane airport. In each of these cases, we overlay the borders of the airport ward with a map of satellite pictures from google maps. We then draw a polygon of the exact airport area and excluded this area from the calculation of average luminosity of the airport ward. The results corresponding to the Durban King Shaka International Airport and the Polokwane airport are depicted in Figure C.9 and C.10, respectively. One can easily see that they remain virtually unchanged as compared with the initial results of Section 4.2 of the paper. Therefore, our findings are not solely driven by the airport itself, but (also) strongly by an increase in economic activity around it, e.g., by firms and people that settle nearby the airport due to its amenities.

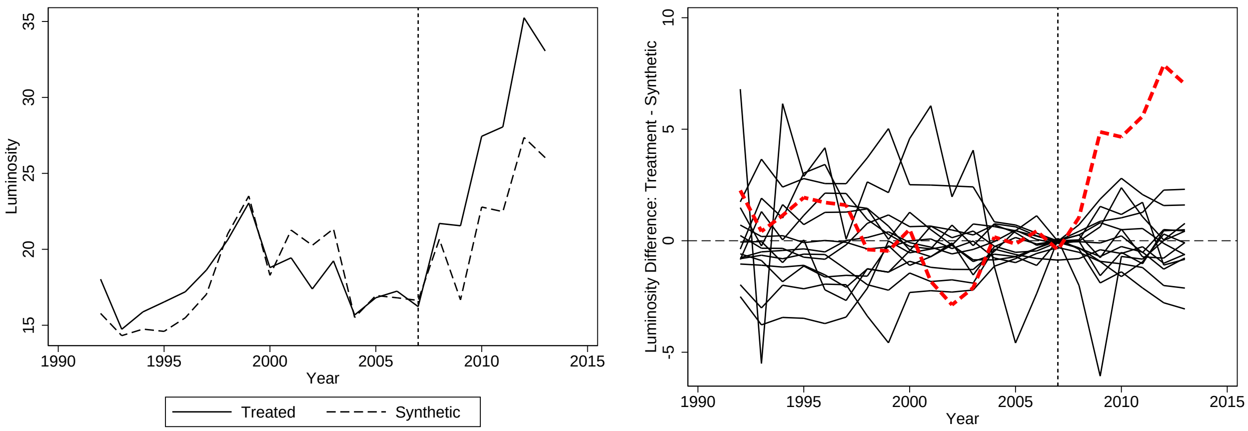

(a) Trends in Luminosity: Treated vs. Synthetic (b) Luminosity Gaps: Airport Ward vs. Placebos

Note: The vertical dashed line indicates the end of the pre-treatment period (2007). Panel (a) displays the airport ward and its synthetic counterpart. Panel (b) plots luminosity gaps (treatment minus synthetic) for the airport ward and placebo units: the black solid lines and the red dashed line represent the placebos and the treated unit, respectively.

Figure C.8: Estimation Results: Polokwane International Airport—Treatment Beginning in 2008

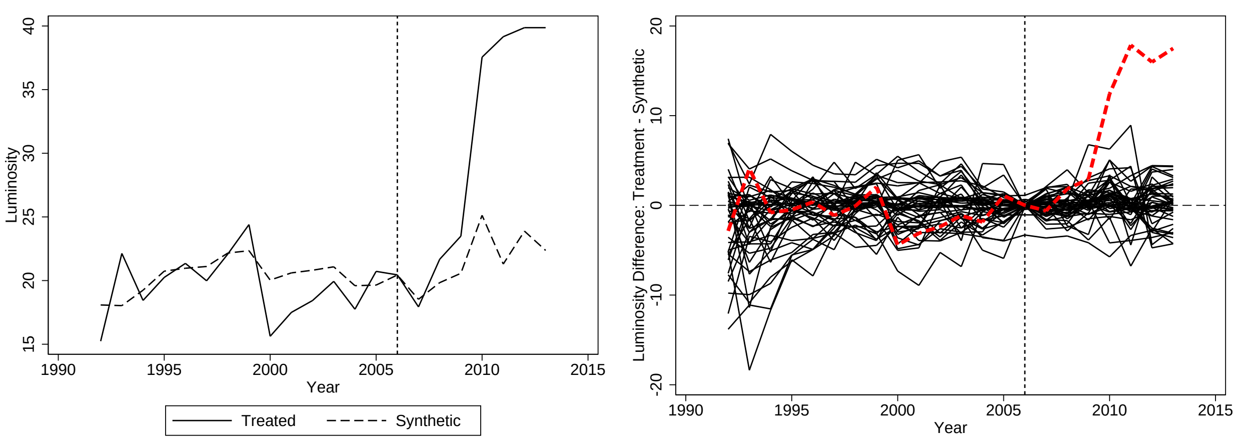

(a) Trends in Luminosity: Treated vs. Synthetic (b) Luminosity Gaps: Airport Ward vs. Placebos

Note: The vertical dashed line indicates the end of the pre-treatment period (2006). Panel (a) displays the airport ward and its synthetic counterpart. Panel (b) plots luminosity gaps (treatment minus synthetic) for the airport ward and placebo units: the black solid lines and the red dashed line represent the placebos and the treated unit, respectively.

Figure C.7: Estimation Results: King Shaka International Airport Durban—Treatment Beginning in 2007

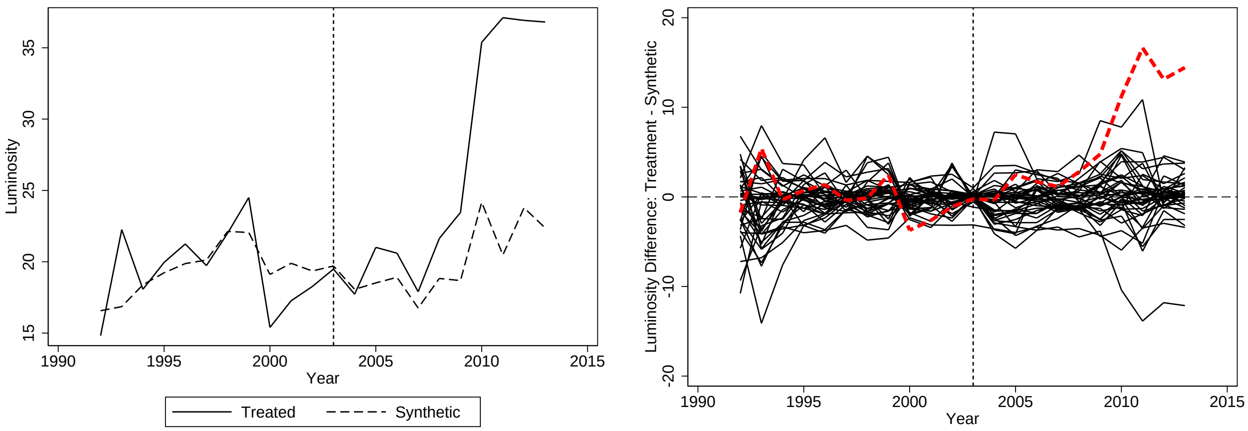

(a) Trends in Luminosity: Treated vs. Synthetic (b) Luminosity Gaps: Stadium Ward vs. Placebos

Note: The vertical dashed line indicates the end of the pre-treatment period (2003). Panel (a) displays the airport ward and its synthetic counterpart. Panel (b) plots luminosity gaps (treatment minus synthetic) for the airport ward and placebo units: the black solid lines and the red dashed line represent the placebos and the treated unit, respectively.

Figure C.9: Estimation Results for King Shaka International Airport Durban Without Airport Lights

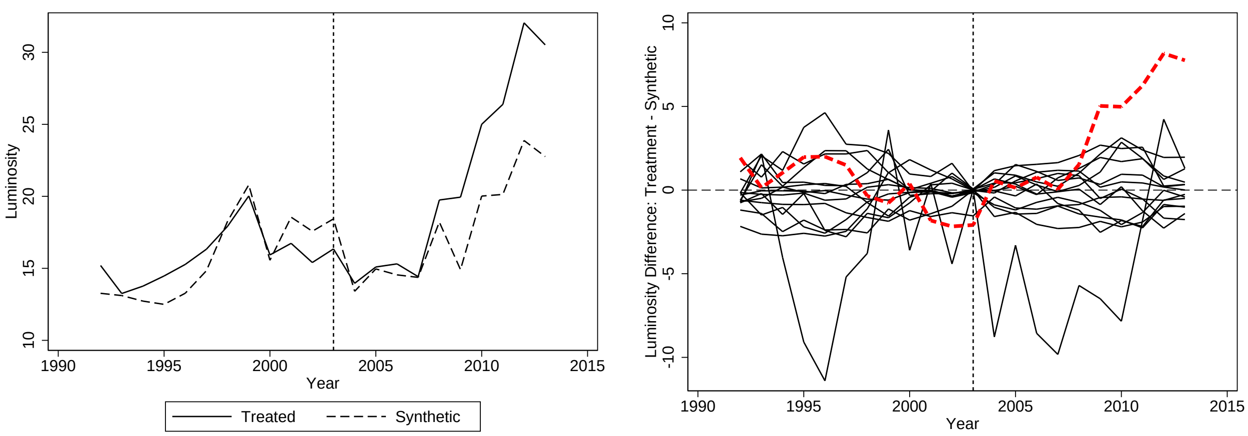

(a) Trends in Luminosity: Treated vs. Synthetic (b) Luminosity Gaps: Stadium Ward vs. Placebos

Note: The vertical dashed line indicates the end of the pre-treatment period (2003). Panel (a) displays the airport ward and its synthetic counterpart. Panel (b) plots luminosity gaps (treatment minus synthetic) for the airport ward and placebo units: the black solid lines and the red dashed line represent the placebos and the treated unit, respectively.

Figure C.10: Estimation Results for Polokwane International Airport Without Airport Lights

D Difference-in-Differences Analysis

Difference-in-Differences (DiD) approaches, seemingly suitable in a setting as ours, have been criticized for offering too much leeway in choosing the respective control unit, giving room for manipulation and questioning the external validity of corresponding results (e.g., Bertrand et al., 2004; Hansen, 2007a,b). Moreover, the common trends assumption needed for identification might be too restrictive in our context, which is why we apply a causality test in the spirit of Granger (1969). This is done by regressing the outcome of interest on a treatment dummy, period-fixed effects, as well as several successive leads (anticipatory effects), i.e., interacted placebo indicators ranging from one to (in this case) five years prior to our treatment. More precisely, luminosity is regressed on a World Cup venue indicator, year-fixed effects, and interacted year-treatment fixed effects. The latter comprise the coefficients of interest described above: they should be close to zero and statistically insignificant if the common trends assumption is about to hold. Under the common trends assumption, time-specific effects should be fully captured by the set of period dummies in the pre-intervention phase, i.e., the placebo indicators should not matter. However, Table D.1 shows that every single placebo effect is statistically significant at the one-percent level, strongly indicating that the identifying assumption of DiD is violated.

Abstracting from the violation of the common trends assumption shown, we still estimate a DiD specification for comparing our SCM estimates on municipality level. We arrive at a virtually identical treatment effect of 1.74 (p-value = 0.000)—see Column (1) of Table D.2. Therefore, our result does not depend on the particular estimation method used, which is what one would expect comparing DiD with SCM.

Table D.1: Common Trends: Plausibility Check

Dep. Var. is Luminosity					
Year	1999	2000	2001	2002	2003
Placebo Indicator	2.917***	0.870***	1.339***	1.760***	2.037***
	(0.590)	(0.299)	(0.360)	(0.377)	(0.408)

Note: Luminosity is regressed on a treatment dummy, year-fixed effects, and interacted year-treatment fixed effects. Merely the latter are reported, i.e., *Placebo Indicator*s. Level of observation is municipality ($N = 5,148$). Standard errors in parentheses are clustered on municipality level. Coefficient is statistically different from zero at the ***1 % level.

Table D.2: Main Specification (Municipality Level): Difference-in-Differences Results

Dep. Var. is Luminosity	(1)
Treatment Period	2004-2013
DiD Estimate	1.737***
	(0.407)
Year Dummies	Yes
Municipality Dummies	Yes
Observations	5,148
Within R^2	0.324
Number of Municipalities	234

Note: Standard errors in parentheses are clustered on municipality level. Coefficient is statistically different from zero at the ***1 % level. The regression additionally includes a dummy variable indicating the treatment period (i.e., variables equal to one for the period after 2004). The DiD estimate shows the coefficient of the interaction term of a dummy variable indicating treated municipalities and the treatment period dummy.

References

Abadie, A., Diamond, A., and Hainmueller, J. (2015). Comparative Politics and the Synthetic Control Method. *American Journal of Political Science*, 59(2):495–510.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should We trust Differences-in-Differences Estimates? *The Quarterly Journal of Economics*, 119(1):249–275.

Chen, X. and Nordhaus, W. D. (2011). Using Luminosity Data as a Proxy For Economic Statistics. *Proceedings of the National Academy of Sciences of the United States of America*, 108(21):8589–8594.

Elliot, R. J., Strobl, E., and Sun, P. (2015). The Local Impact of Typhoons on Economic Activity in China: A View from Outer Space. *Journal of Urban Economics*, 88:50–66.

Elvidge, C. D., Hsu, F.-C., Baugh, K. E., and Ghosh, T. (2014). National Trends in Satellite-Observed Lighting 1992–2012. In Weng, Q., editor, *Global Urban Monitoring and Assessment through Earth Observation*, pages 595–622. CRC Press, Boca Ranton, FL.

Elvidge, C. D., Ziskin, D., Baugh, K. E., Tuttle, B. T., Ghosh, T., Pack, D. W., Erwin, E. H., and Zhizhin, M. (2009). A Fifteen Year Record of Global Natural Gas Flaring Derived from Satellite Data. *Energies*, 2:595–622.

Gonzalez-Navarro, M. and Turner, M. A. (2016). Subways and Urban Growth: Evidence from Earth. SERC Discussion Paper 195.

Granger, C. W. J. (1969). Investigating Causal Relations by Econometric Models and Cross-Spectral Methods. *Econometrica*, 37(3):424–438.

Hansen, C. B. (2007a). Asymptotic Properties of a Robust Variance Matrix Estimator for Panel Data when T is large. *Journal of Econometrics*, 141(2):597–620.

Hansen, C. B. (2007b). Generalized Least Squares Inference in Panel and Multilevel Models with Serial Correlation and Fixed Effects. *Journal of Econometrics*, 140(2):670–694.

McGregor, T. and Wills, S. (2017). Surfing a Wave of Economic Growth. *OxCarre Working Paper No. 170*.

Zabel, F., Putzenlechner, B., and Mauser, W. (2014). Global Agricultural Land Resources. A High Resolution Suitability Evaluation and its Perspectives until 2100 under Climate Change Conditions. *PLOS One*, 9(9):1–12.