UNIVERSITY OF WAIKATO

Hamilton
New Zealand

How effective are sanctions on North Korea? Popular DMSP night-lights data may bias evaluations due to blurring and poor low-light detection

Working Paper in Economics 6/22

March 2022

John Gibson
School of Accounting, Finance and Economics
University of Waikato
Private Bag 3105
Hamilton, 3240
New Zealand

Tel: +64 (7) 838 4289
Email: jkgibson@waikato.ac.nz

Bonggeun Kim
Department of Economics
Seoul National University
South Korea

Geua Boe-Gibson
School of Accounting, Finance and Economics
University of Waikato
New Zealand

Abstract

The effect of sanctions on economic activity in targeted countries is increasingly studied with satellite-detected night-lights data because conventional economic activity data for such countries are either unavailable or untrustworthy. Many studies use data from the Defense Meteorological Satellite Program (DMSP), designed for observing clouds for short-term weather forecasts rather than for long-run observation of economic activity on earth. The DMSP data are flawed by blurring, and bottom-coding due to poor low-light detection. These errors may bias evaluation of sanction effectiveness. To show this we use a difference-in-differences analysis of impacts on night-lights of the shutdown of the Kaesong Industrial Zone in North Korea, which South Korea closed in 2016 in response to North Korea's nuclear tests. We estimate impacts of 40–80% declines in luminosity, depending on the choice of comparison region, and these effects are always precisely estimated if data from the accurate Visible Infrared Imaging Radiometer Suite (VIIRS) on the Suomi-NPP satellite are used. Yet with the more widely used DMSP data, apparent impacts are imprecisely estimated and are far smaller. A decomposition suggests much of the attenuation in estimated treatment effects if DMSP data are used comes from false zeroes, which are also likely to matter to evaluations in other poorly lit places.

Keywords

DMSP

mean-reverting error

night lights

sanctions

VIIRS

North Korea

JEL Classification

C80, F51, O11

Acknowledgements

Financial support from Marsden Fund project UOW1901 is acknowledged.

1. Introduction

The imposition of sanctions on Russia following the Ukraine invasion in February 2022 is the latest attempt by the international community to alter behaviour of countries whose actions may violate international norms. Several studies estimate economic impacts of sanctions but face a fundamental difficulty; countries under sanction are unlikely to have trustworthy data available for researchers, especially if such data would show that the sanctions are helping to weaken their economy. In response, a popular recent approach relies on remote sensing data, especially of satellite-detected night-lights. These data cover all countries and should be free from manipulation. Some recent studies of impacts of sanctions imposed on North Korea (Lee, 2018; Son & Cho, 2021; Kim et al, 2021) and Iran (Arbatli & Gomtsyan, 2021; Farzanegan & Fischer, 2021) are good examples of this use of satellite-detected night-lights.

An important measurement issue with undesirable econometric implications does not seem to be considered in this recent literature. Most studies use the Defense Meteorological Satellite Program (DMSP) data, which involves repurposing information originally collected for cloud detection for short-term Air Force weather forecasts. While the seminal study by Henderson et al (2012) showed that DMSP data can also be used to make long-run observations of economic activity on earth, the framework used in that seminal study only required measurement errors in DMSP data to be independent of errors in reported GDP. Independence mattered because an optimally weighted mix of the two types of data was used to proxy for true but unknown economic activity. In contrast, the studies estimating impacts of sanctions use the DMSP data directly as a proxy for economic activity and so the measurement errors in these data can cause an econometric bias that may distort conclusions about the effectiveness of sanctions.

The DMSP data are used on the left-hand side of econometric models in the studies evaluating sanctions, so it requires non-classical measurement errors if there is to be bias in the estimated regression coefficients. The first source of such errors is false zeroes, due to the poor low-light detection capability of DMSP sensors. Given the original purpose of detecting clouds, the sensor needs little amplification when cloud tops are visible in moonlight during half of the lunar cycle, so lights on earth that are not very brightly lit often go undetected. For example, Chen & Nordhaus (2015) study 1°×1° grid cells over Africa, each having 10,000 to 100,000 people and all having anthropomorphic lights according to other data sources, yet in the DMSP data over one-half of these cells are recording as having zero light. Top-coding is a more widely discussed problem in DMSP data (Bluhm & Krause, 2020) but bottom-coding from false zeroes (Abrahams et al, 2018) matters more for poorly-lit areas. Non-classical measurement errors also result from blurring, where night-lights are attributed to places from which light is not emitted. Blurring is an inherent feature of DMSP sensors and data management (Elvidge et al, 2013; Tuttle et al, 2013) and it causes mean-reverting measurement errors (Gibson, 2021).

To provide evidence on the econometric bias caused by these DMSP measurement issues we use a difference-in-differences analysis of the impacts on night-lights of a particular sanction on North Korea: the shutdown of the Kaesong Industrial Zone in February 2016 in response to

North Korea's nuclear tests. This joint venture with South Korea operated about 10 km north of the DMZ, with over 50,000 North Korean workers employed in factories operated by South Korean companies. We estimate 40–80% declines in luminosity, depending on the choice of comparison region, as the impact of this sanction. These effects are always precisely estimated if data from the accurate Visible Infrared Imaging Radiometer Suite (VIIRS) of instruments on the Suomi-NPP satellite are used.[1] Yet if we use the more popular DMSP data for the same years, the same spatial units, and the same sanctions event, the apparent impacts are imprecisely estimated and are far smaller. In other words, using the same data source as in recent studies of sanctions impacts we get far smaller estimated impacts, compared with what we find using a more accurate data source, even though that more accurate data source remains largely unused in the literature that is relying on satellite-detected night-lights to study sanctions impacts.[2]

These results suggest it would be useful to reexamine some of the recently published findings on the impacts of sanctions. If the measurement error bias that we demonstrate here holds more widely then it is possible that effectiveness of sanctions in weakening the economy of targeted countries may have been understated because studies have used satellite-detected night-lights data that are affected by non-classical measurement errors. Importantly, our results should not be affected by sanctions evasion measures because we know the sanction we study was fully implemented, in the sense that all South Korean enterprises shut their activities in Kaesong. In contrast, studies that focus on trade sanctions (e.g. Lee, 2018; Kim et al, 2021) face the difficulty that North Korea has sanctions evasion strategies like falsification of documents and covert ship-to-ship transfers of cargo at sea. Consequently, trade sanctions may appear ineffective either because measurement errors in the night-lights data understate the impacts, or because sanctions actually are partly ineffective due to North Korea's evasions. In contrast, with the Kaesong closure, we have a 'ground-truth' event as a benchmark for providing information about the measurement errors in the data used to evaluate the impact of this event.

We also go beyond showing a bias in the difference-in-differences analysis when the DMSP data are used, by examining the source of the understated impacts. By decomposing the change in results when VIIRS data are used instead of DMSP data it seems that much of the attenuation in estimated treatment effects if DMSP data are used comes from false zeroes. Applied studies using DMSP data—in general, rather than specifically for studying the impacts of sanctions—use various *ad hoc* methods for dealing with apparently unlit observations. A common tactic is to use transformations such as taking the logarithm of $(y + 0.01)$, or some other small constant added to the data (Gibson et al, 2020). Based on the results here, and on more general concerns

[1] A growing empirical literature shows VIIRS data are far more accurate than DMSP data. Early discussion of the theoretical superiority of VIIRS over DMSP for studying night-lights includes Miller et al (2012) and Elvidge et al (2013); this second reference notes that VIIRS data are up to 45-times more accurate than DMSP data. More recently, empirical comparisons with non-nested tests show that models using VIIRS data provide results that are closer to the truth than are models using DMSP data (Gibson, 2021; Zhang & Gibson, 2022).
[2] A *Google Scholar* search on 11 March 2022 for "DMSP" and "sanctions" returned 458 results, while a similar search for "VIIRS" and "sanctions" returned only 138 results. Likewise, using IDEAS/RePEc we found more than twice as many articles using DMSP than using VIIRS, attesting to the comparative popularity of DMSP data in this literature.

about how this sort of *ad hoc* transformation may bias estimated elasticities (Bellemare & Wichman, 2020), finding better solutions for bottom-coded false zeroes would be valuable.

The rest of the paper is set out as follows: Section 2 provides details on the context, Section 3 discusses the data and econometric methods, Section 4 has the results and Section 5 concludes.

2. Context: Sanctions on North Korea

2.1 Previous Evidence on Impacts of the Sanctions

Over the last two decades the international community has responded to North Korea's nuclear testing program by imposing several waves of sanctions. After North Korea's first reported nuclear test in October 2006 the UN Security Council unanimously adopted Resolution 1718 that banned various imports and exports to North Korea, and froze assets and banned travel by participants in the nuclear program. However, China did not participate in these sanctions and China's trade with North Korea appears to have increased to offset some of the loss of trade with other countries. Further UN sanctions were imposed in 2009 following North Korea's second nuclear test and additional sanctions were put in place by South Korea in 2010 after an attack by North Korea on a South Korean navy vessel.

An evaluation of the impact of these sanctions up until 2013 used annual DMSP night-lights data for a micro-grid overlaid on the map of North Korea, where cell sizes for the grid were about 2.6 km^2 (Lee, 2018). According to the results of this study, while there appeared to be little overall effect of sanctions on luminosity, some redistribution of economic activity to the regions of North Korea that were more linked to trade with China was detected. One concern with results of this study is that the cell size for the grid is far smaller than the ground footprint of DMSP pixels, which is 25 km^2 at nadir and may be up to 60 km^2 away from the nadir (Elvidge et al, 2013).[3] Thus, the luminosity values attributed to particular cells are likely to include lights from nearby cells. This blurring of the observations will create a form of mean-reverting error that biases regression coefficients towards zero, in proportion to the strength of the mean-reversion in the left-hand side variable (Gibson, 2021).

A more aggregated approach with night-lights data is used by Son & Cho (2021), based on the growth in night-lights for 25 major cities in North Korea. A range of new sanctions imposed in 2016, which notably included China as one of the countries limiting trade with North Korea, appears to have decreased the growth in night-lights, particularly in cities where coal mining is important. Notably, coal was included in the list of banned North Korean exports under two UN resolutions (#2270 in March 2016 and # 2371 in August 2017) and prior to then coal had been North Korea's largest export, contributing about one-third of total export revenues (Kim

[3] The ground footprint expands away from the nadir because of the angle of viewing the earth. Although DMSP has a 3000 km wide swath, only data from the central 1500 km of the swath are used to form annual composites because further out towards the extremes the ground footprint grows to be four-times as large as at the nadir (Falchi and Cinzano, 1998). Thus, at swath extremities the ground footprint covers approximately 100 km^2 and all light from that area is wrongly attributed to the pixel at the centre of the footprint (Abrahams et al, 2018).

et al, 2021). The study by Son & Cho (2021) used a time-series of night-lights data that spliced records from the two data sources: DMSP and VIIRS. However, there was no allowance for the quite different measurement error properties of the two data sources (with a mean-reverting error in the DMSP data there is no simple adjustment factor to 'line up' the two data sources because the error will vary with the true, but unknown, level of luminosity). Moreover, splicing together these two data sources to give a longer time-series (the DMSP time-series previously ended in 2013) is no longer needed because the DMSP time-series was recently extended through to 2019 (Ghosh et al, 2021).

Another study of the impacts of various layers of sanctions on North Korea is Kim et al (2021), who use night-lights data aggregated to the second sub-national level (this corresponds to counties for rural areas and districts within urban areas). This study uses VIIRS data rather than DMSP data, because the focus is on sanctions imposed in 2016 and 2017, near the middle of the 2014-19 time-series that is used (VIIRS became operation from 2012). The sanctions on exports and intermediate inputs were found to cause declines in luminosity, while product-level market price data suggest that North Korea faced significant price increases for imports of sanctioned products. This study also attempted to distinguish between the effects of different waves of sanctions. The layering of various sanctions over time, and of the countries adhering to them, with China eventually joining with other countries in imposing sanctions, points to the difficulty of the empirical task in trying to estimate impacts of the sanctions. Sanctions may not seem to have much overall impact, as Lee (2018) found, either due to the luminosity data used having particular types of measurement errors or alternatively because North Korea successfully evades sanctions that are not applied in an immediate and complete manner. So it requires a different type of sanctions event to use as a benchmark for learning about the nature of the measurement errors in the luminosity data used to evaluate impacts of the sanctions.

2.2 North Korea's Administrative Geography

A basic overview of the administrative geography may help readers to interpret the difference-in-differences results of the sanctions event studied below. At the first sub-national level, North Korea has nine provinces and four special municipalities that are centrally-controlled rather than under a province. Each of these centrally-controlled municipalities is called a "city" but this is an administrative term rather than a functional description, with low-density rural areas also within these municipalities.[4] One of these centrally-controlled municipalities is "Kaesong City", which is on the border with South Korea and is about 50 km northwest of Seoul. At the second sub-national level, North Korea is divided into 186 areas, which are named as districts within the four centrally controlled municipalities and usually are named as counties in the more rural areas.

The Kaesong Industrial Zone is located in Panmun District of Kaesong City (Figure 1). The only other district in Kaesong City is Kaepyong District. So one set of difference-in-differences

[4] This is equivalent to the situation in China, where Chongqing is one of four centrally-controlled municipalities and this "city" covers as area as large as Austria, and only a small part of it is densely built-up urban area.

results involves comparing luminosity trends in Panmum and Kaepyong districts, where the treatment (closure of the industrial zone) occurred sufficiently early in 2016 (on February 10) that we can use annual data from 2016 onwards as the post-treatment period, and values from 2015 and earlier as the pre-treatment period.[5] Although Kaepyong District is physically closest to the treated unit there are two other candidates for control group membership that are more distant, where their candidature is based on their having special economic zone status: Rason and Sinuiju. These two cities have economic activity that is linked closely to China, and they are more market oriented than other parts of North Korea. Examples of these links include China's domestic shipments of coal from mines in northeast China going via Rason to Shanghai and other coastal destinations in China, and the use of Sinuiju as a site for entrepôt trade (both legal and illegal) across the Yalu river into China.

Figure 1: Administrative Areas of North Korea

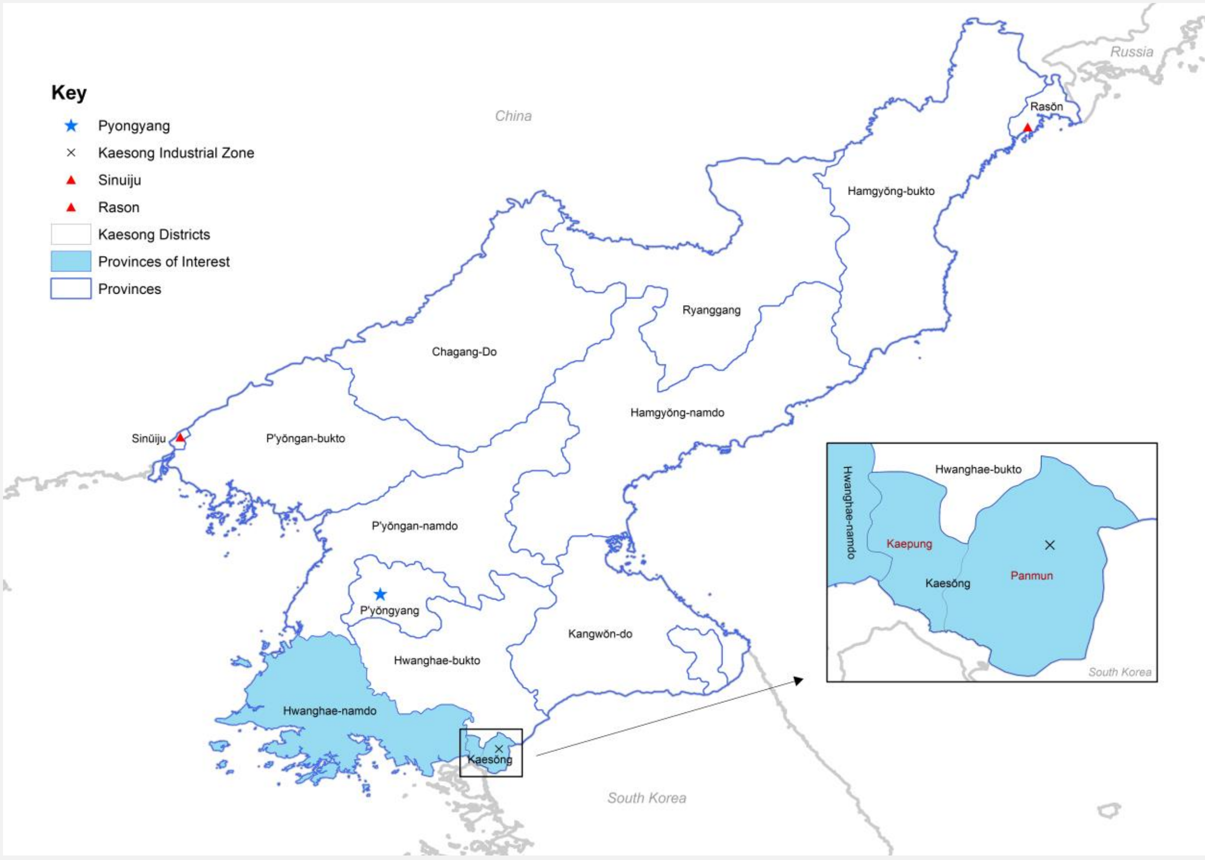

There is no guarantee that the impacts of the closure of the Kaesong Industrial Zone will be confined to Panmun District, so we also present difference-in-differences results at a more spatially aggregated level, using Kaesong City as the treated unit. The reason for this broader treatment unit is that employment at Kaesong Industrial Zone was equivalent to one-sixth of the entire population of Kaesong City so impacts beyond the Panmun District are quite likely. The neighbouring province, Hwanghae-namdo (South Hwanghae) is a natural comparison unit

[5] While monthly data are available, they are not cleaned and processed to the same extent as annual composites and there are often missing observations due to the impacts of stray light in some seasons (Gibson, 2021).

and we also use all parts of North Korea outside of Kaesong City as a potential comparison group. In contrast to using Sinuiju and Rason as comparison groups, which are spatial units that are somewhat urbanized and illuminated, using an entire province like Hwanghae-namdo or even more so for using other parts of North Korea as comparison groups brings the very low levels of lighting in these areas into consideration, to see how the problem of false zeroes affects results from the two night-lights data sources.

3. Data and Econometric Methods

We use two sources of night-lights data: DMSP stable lights annual composites and VIIRS Day-Night Band (DNB) version 2 (V.2 VNL) masked average radiance annual composites.[6] Key references with details on how these annual composites are formed are Baugh et al (2010) and Ghosh et al (2021) for DMSP and Elvidge et al (2017, 2021) for VIIRS.

The DMSP data are 6-bit digital numbers (DN) ranging from 0 to 63 (with lower values for less brightly lit areas), that are reported for each 30 arc-second output pixel (which is roughly 0.9×0.7 km at the latitude of North Korea). Importantly, the output pixel size is far smaller than the spatial resolution of what the sensor can detect, where this coarse resolution of the sensor is due to effects of geolocation errors, to the aggregation of pixels to save on limited data storage and to the expansion of the footprint from viewing the earth at an angle away from the nadir (Elvidge et al, 2013; Tuttle et al, 2013; Abrahams et al, 2018). In the annual composites, 'stable' simply means that ephemeral lights, from sources such as fires and gas flaring, are removed before the annual composite is built up from nightly images. An alternative meaning of 'stable', in terms of temporal consistency, does not apply as there is no in-built calibration of the DMSP data to adjust for the constant changes in sensor amplification (that are made without any record kept, where the amplification changes are so that cloud-tops can be viewed with similar brightness across the light and dark part of the lunar cycle). Consequently, it is best to think of the DN value as a relative measure of brightness, because the same DN value in different years could correspond to different radiance values (Doll, 2008). While there are radiance-calibrated DMSP data (Elvidge et al, 1999), these are only available in certain years, and not since 2010, and so are not available for assessing the impact of the sanction of closing the Kaesong Industrial Zone in 2016.

[6] The DMSP data are available from: https://eogdata.mines.edu/products/dmsp/#download and the VIIRS data from https://eogdata.mines.edu/products/vnl/#annual_v2. The DMSP data are from satellite F18 for 2012-13 and from F15 from 2014 onwards. There is a difference in observation time for these two satellites with F18 having a late evening (ca. 10pm) observation time and F15 an early morning (ca. 3.30am) observation time (while the observation time of VIIRS is in the middle, at ca. 1.30am). We used a test for known break-points due to Ditzen et al (2021), to see if the change in the time of observing earth for the two DMSP satellites affected the difference-in-differences results and it did not. We also compared the proportion of unlit observations at the second sub-national level for North Korea according to F18 (0.48) and to F15 (0.50), with no significant difference due to the change in observation time apparent. In terms of VIIRS, comparisons between the various V.2 VNL data products and DMSP data show that the V.2 VNL masked average radiance product that we use here provides the best proxy (amongst the various VIIRS data products) for economic activity, using a panel of United States county-level and state-level GDP as a benchmark (Gibson & Boe-Gibson, 2021).

The V.2 VNL annual composites are produced from monthly cloud-free radiance averages, with initial filtering by remote sensing specialists (Elvidge et al, 2021) to remove extraneous features, such as fires and aurora, before a further set of outlier removal procedures are used to isolate lit grid cells from background. The data are in units of nano Watts per square centimeter per steradian (nW/cm^2/sr) reported on a 15 arc-second output grid (equivalent to 0.47×0.38km at the latitude of North Korea). Unlike DMSP data, there are no blurring issues with V.2 VNL data, with the sensor compensating for the change in ground footprint size as the earth is viewed at an angle and with no data storage constraints causing pixels to be aggregated. Moreover, the VIIRS DNB sensor can detect lights on earth in a far wider range of lighting conditions, that covers seven orders of magnitude from minimum luminosity to maximum luminosity while the DMSP sensor covers only two orders of magnitude (Gibson et al, 2020).

The inherent blurring in the DMSP images is illustrated in Figure 2, which compares the VIIRS annual composite for 2015 with the DMSP annual composite for the same year. The choice of year is deliberate, to show the situation immediately prior to the Kaesong Industrial Zone being shut down from early February 2016. In the VIIRS image the Kaesong Industrial Zone is shown as a distinct location of lights with unlit space between there and the border with South Korea (the overall far greater luminosity level in South Korea is also clearly apparent). In the VIIRS image there are 75 illuminated pixels in the Kaesong Industrial Zone, which covers an area of about 14 km^2 (equivalent to the combined area of 20 soccer fields).

Figure 2: Illustrating Blurring in the DMSP Images: Kaesong Industrial Zone in 2015

VIIRS DNB 2015 Annual Composite | DMSP Stable Lights 2015 Annual Composite

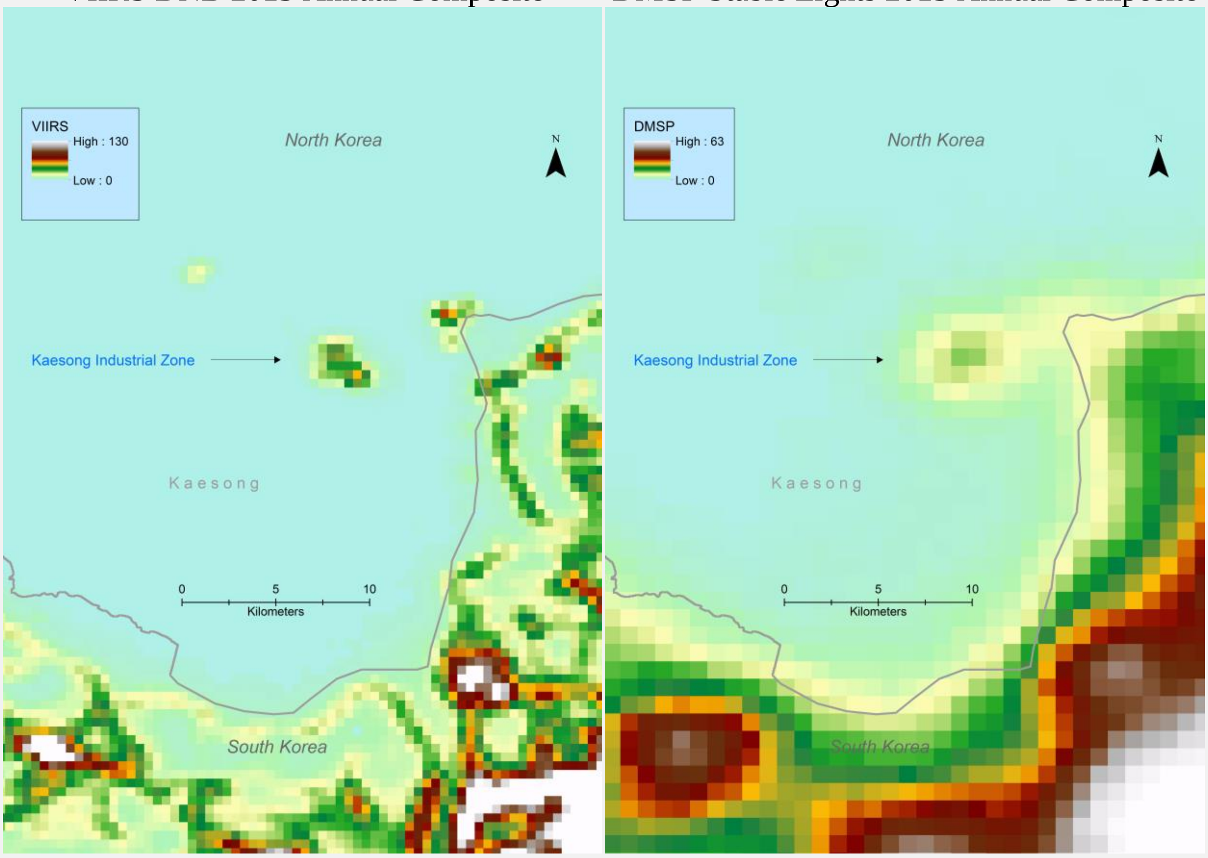

In the DMSP image the illuminated area from the Kaesong Industrial Zone appears far larger, at almost 90 km^2 (so over 6-times as large as what VIIRS shows).[7] Moreover, lit area seems to extend all the way to the border with South Korea and so luminosity is attributed to unlit places (such as between the border and the Kaesong Industrial Zone), which is a common failing of the DMSP data. The blurring shown in Figure 2 suggests that results using DMSP data for micro-grids, such as the grid used by Lee (2018), may have some biases because the blurring causes a mean-reverting error which has been shown in other settings to cause econometric bias in results estimated with DMSP data (Gibson, 2021; Gibson & Boe-Gibson, 2021).

It is unclear from Figure 2 whether blurring of DMSP data is sufficient to cause bias when the spatial units are counties and districts rather than a micro-grid. In other words, even though the Kaesong Industrial Zone appears much bigger in the DMSP image in Figure 2, with unlit areas being attributed some of the light from elsewhere (which causes an upwards reversion towards the mean) the overstatement of illuminated area does not appear to be so large that it spills over the District boundary. Therefore to study this question in another way we use a decomposition of the variance into between-area and within-area components, where blurring is expected to reduce the share of within-area variation. Specifically, an ANOVA using data from 2012 to 2019 for the second-level sub-national units decomposes variation in log luminosity into the part due to differences between first sub-national level units (provinces and centrally controlled municipalities) and the part due to differences within those units (that is, between counties or districts within the same province). With DMSP the partial R^2 values for the between and within components are 0.213 and 0.159, so there appears to be less variation within provinces than between provinces. Yet with the VIIRS data the situation is reversed, with the partial R^2 values for the between and within components being 0.117 and 0.133. In other words, the blurring in DMSP data that is illustrated in Figure 2 at a fine scale also operates at a larger spatial scale in making districts within a province or centrally controlled municipality appear more alike. This reversion towards the mean caused by blurring would be expected to attenuate econometric estimates of difference-in-differences results when a treatment like the shutdown of the Kaesong Industrial Zone is heavily concentrated in one district because the differences between that district and the other parts of the same province will be blurred.

The other relevant flaw in the DMSP data, of false zeroes, is harder to show on a map. Instead, to illustrate this flaw we aggregated the night-lights data for 2015 to the district/county level (the second sub-national level) and compared distributions of the DMSP and VIIRS data. Of the 186 spatial units, 56% of them (n=105) appear to be entirely unlit in the DMSP annual composite for 2015. However, the VIIRS composite for the same year shows that over 60% of these 105 apparently unlit second-level units actually did have positive levels of luminosity. In other words, false zeroes made up more than one-half of the apparently (according to DMSP data) unlit counties and districts in 2015, which is somewhat higher than the rate found in a different setting (and using larger spatial units) by Chen & Nordhaus (2015). Across all years,

[7] Examples from elsewhere also show that DMSP overstates lit area by 6-fold or more (Gibson et al, 2020) due to the inherent blurring of the images.

the fraction of false zeroes in areas that DMSP indicated as unlit ranged from 66% to 78%, and averaged 70%. These false zeroes provide another form of non-classical measurement error (in addition to the mean-reverting error due to the blurring of the DMSP images) which could bias regression coefficients even when night-lights data are the left-hand side variable.

The basis of the econometric models is shown graphically in Figure 3, which has a comparison of mean radiance over time in Panmun District (where the Kaesong Industrial Zone was based) and in the neighbouring Kaepyong District. The average radiance level in Panmum was much higher than in Kaepyong, especially prior to the closure of the Industrial Zone, so the chart is based on movements relative to the 2012 baseline in each district. In the econometric model, this same effect of adjusting for pre-treatment differences in the levels of the outcome variable is dealt with using dummy variables for the treated unit(s).[8]

Figure 3: Changes in Average Radiance (V.2 VNL) in Selected Areas of North Korea

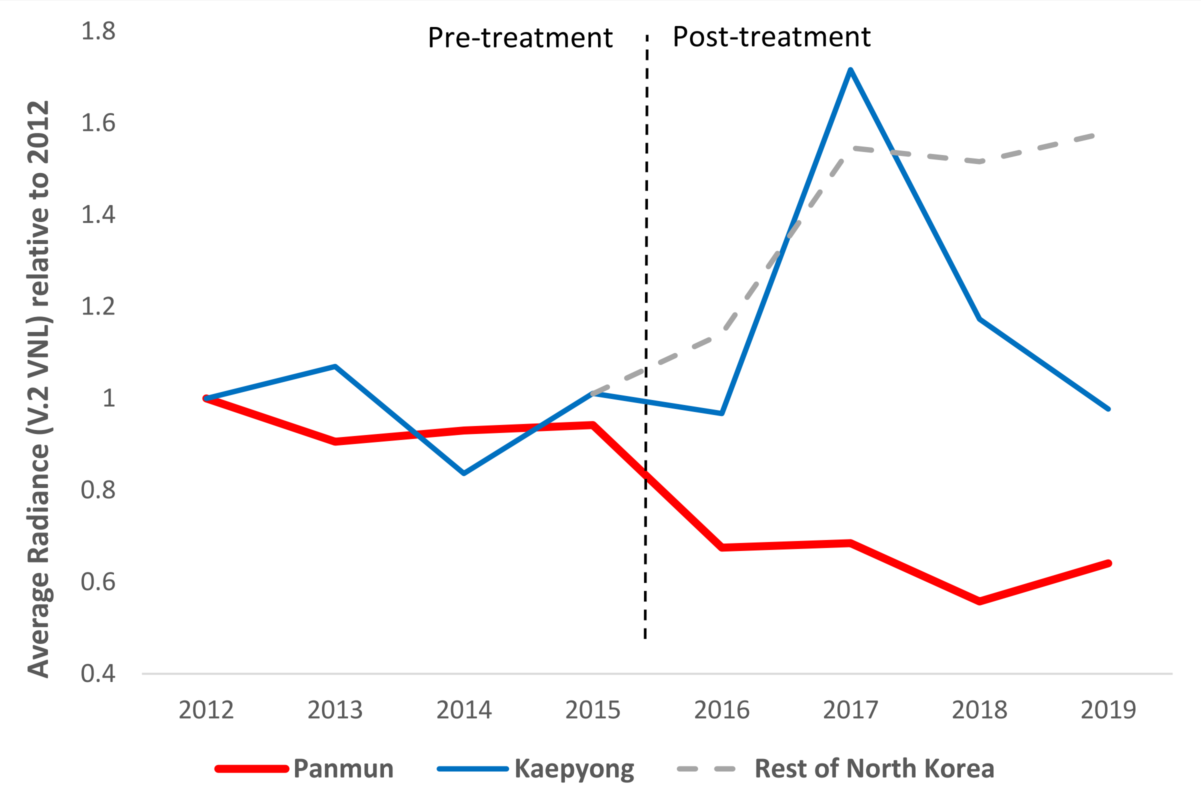

In the years before the closure of the industrial zone, average radiance in Panmun was 94% of the baseline radiance, with a range of 91–100%. The average radiance dropped almost 30% in 2016, after closure of the industrial zone, and thereafter averaged just 64% of baseline. This single difference, of a drop of about 36% based on the before and after comparison, is only a valid estimate of the impact of the sanction if everything elsewhere in North Korea stayed the same. In fact, there was considerable growth in luminosity in other parts of North Korea, even

[8] A test for parallel trends in the pre-treatment 2012-15 period does not reject (p<0.49) the null of no difference in the trends for Panmun and Kaepyong.

with the various layers of sanctions imposed. In the neighbouring Kaepyong District, mean radiance had averaged 98% of baseline during 2012-15, and from 2016 onwards it was 21% above baseline with particularly strong growth in mean radiance from 2016 to 2017 (Figure 3). The combination of the fall in average radiance in Panmun after closure of the industrial zone, and the rise in average radiance emitted from Kaepyong, gives a difference-in-difference estimate of the impact of the sanction as a 57% fall in luminosity relative to the baseline level of luminosity. Thus, rather than a decline of about one-third, as shown with a single difference approach, the actual impact of the closure of the industrial zone was more like a three-fifths decline in luminosity, reflecting the greatly reduced economic activity post-closure.

The increased luminosity in the Figure 3 control group region, Kaepyong, contributes to the difference-in-differences estimate of the impact of closing the industrial zone. This increase is not a fluke, as other parts of North Korea also had increased luminosity from 2016 onwards, despite various trade sanctions in place. If we use the movements in average radiance in the rest of North Korea as the counterfactual, which is shown in Figure 3 by the dashed line, the difference-in-difference estimate of the impact of the closure on Panmun is even larger than when using Kaepyong as the control group, with a fall in luminosity equivalent to about 70% of the baseline level.

Although Figure 3 does not show it, so as to avoid clutter, movements in average DN values from DMSP data imply substantially smaller impacts of the closing of the Kaesong Industrial Zone. Rather than a difference-in-difference estimate of a 57% fall in luminosity due to the sanction, which is what the VIIRS data show (when using Kaepyong as the control group), the DMSP data suggest just a 16% fall in luminosity due to closure of the industrial zone.

In order to explore these differences more formally we use the following econometric model:

$$ln(luminosity)_{it} = \beta_0 + \beta_1 D_i + \beta_2 T_t + \beta_3 (D_i \times T_t) + \varepsilon_{it} \qquad (1)$$

where the "luminosity" left-hand side variable is either the V.2 VNL masked average radiance or the DMSP digital number. The spatial dummy variable $D_i = 1$ for the treated unit and 0 for other areas. The treated unit is usually chosen to be Panmun District but we also use more spatially aggregated data and in that case set Kaesong City as the treated unit. The temporal dummy variable T_t has a value of 1 for the post-sanction period (2016 onwards) and 0 for the pre-sanction period. The key coefficient estimate of interest is $\hat{\beta}_3$, which subtracts the change in the sample mean of the control group before and after the sanction from the change in the mean of the treatment group before and after the sanction, or in other words provides an estimate of the following double-difference: $(\bar{y}_{T=1} - \bar{y}_{T=0})_{D=1} - (\bar{y}_{T=1} - \bar{y}_{T=0})_{D=0}$.

The difference-in-differences results from the graphical approach in Figure 3 correspond to the type of values derived from the $\hat{\beta}_3$ estimates (with a logarithmic dependent variable and dummy variable regressors we use the Van Garderen & Shah (2002) estimator to calculate exact percentages and their standard errors). Our interest is in seeing how these results differ when the blurred and bottom-coded DMSP data are used instead of the more accurate VIIRS data.

4. Econometric Results

When we use VIIRS data to estimate equation (1) for getting difference-in-differences impacts of closing the Kaesong Industrial Zone, we estimate 40–80% declines in luminosity, depending on choice of comparison region (Table 1). The largest impacts are if using Hwanghae-namdo as the comparison region (a 77% decline in luminosity in Panmum District relative to the control group) and the smallest from using Kaepyong District (a 39% impact on Panmum). The average impact, across the five different control group regions used in Table 1, is a 61% decline in the luminosity of Panmun District after closure of the Kaesong Industrial zone compared to what would have been expected from temporal changes in other parts of North Korea.

Table 1: Using VIIRS Data to Estimate Difference-in-Differences Impacts on Luminosity of Closing the Kaesong Industrial Zone: Panmun District as the Treated Area with Various Control Groups

	------------------------ Control Group ------------------------				
	Kaepyong District	Rason (Special City)	Sinuiju City	Hwanghae-namdo	Rest of North Korea
Constant, (β_0)	–5.678***	–3.801***	0.428*	–6.044***	-5.579***
	(0.052)	(0.319)	(0.211)	(0.150)	(0.168)
Treated group (D), (β_1)	4.568***	2.691***	-1.538***	4.934***	4.469***
	(0.056)	(0.320)	(0.212)	(0.151)	(0.169)
Sanction period (T), (β_2)	0.092	0.777**	0.353**	1.042***	0.660***
	(0.151)	(0.350)	(0.218)	(0.158)	(0.208)
Treated × Sanction, (β_3)	–0.487**	–1.173***	–0.749***	–1.438***	-1.056***
	(0.156)	(0.356)	(0.222)	(0.163)	(0.213)
D-in-D percentage impact	-39.3%	-71.0%	-53.9%	-76.6%	-66.0%
(Standard error of impact)	(9.4)	(10.0)	(10.1)	(3.8)	(7.2)

Note: Each column represents a separate estimate of equation (1), using VIIRS V.2 VNL masked average radiance annual composites from 2012-19, with samples varying according to the control group used in columns (1) to (5). Robust standard errors in ( ), with the ***, **, * representing statistical significance at the 1%, 5% and 10% levels. The percentage impact values and their standard errors are derived from the estimates for β_3 in the table, using the approximate unbiased variance estimator of van Garderen and Shah (2002).

The difference-in-differences estimates of the impacts, based on $\hat{\beta}_3$, are surrounded by small standard errors (on average, the robust standard errors are just one-quarter of the absolute value of the coefficient, with a range from one-tenth to one-third). Thus, the impact of this particular sanction is quite precisely estimated, even when there are not many annual observations before and after the event that we study.

We also get similar results if we study the impacts of closing the Kaesong Industrial Zone with a more spatially aggregated dataset. Specifically, the results in Table 2 are based on using the entire Kaesong City as the treated unit. This aggregates Panmun District, where the industrial zone was located, and Kaepyong District which is the other sub-unit of Kaesong City, to allow for spillovers. Such spillovers are possible, given that the total employment in the industrial zone was equivalent to one-sixth of the entire population – rather than just of the working age population, which is not available to us – of Kaesong City. With the industrial zone workforce being such a large fraction of Kaesong City population, some workers were likely in Kaepyong District, especially as transit stations in that District were only 15 km from the industrial zone.

Table 2: Using VIIRS Data to Estimate Difference-in-Differences Impacts on Luminosity of Closing the Kaesong Industrial Zone: Kaesong City as the Treated Area with Various Control Groups

	------------------- Control Group			-------------------
	Rason (Special City)	Sinuiju City	Hwanghae-namdo	Rest of North Korea
Constant, (β_0)	–3.801**	0.428*	–6.498***	–5.579***
	(0.319)	(0.211)	(0.123)	(0.168)
Treated group (D), (β_1)	0.407	-3.822***	5.389***	2.185***
	(0.320)	(0.213)	(0.125)	(0.170)
Sanction period (T), (β_2)	0.777**	0.353	0.412	0.660***
	(0.354)	(0.218)	(0.245)	(0.208)
Treated × Sanction, (β_3)	–0.929**	–0.505**	–0.808***	-0.812***
	(0.363)	(0.233)	(0.249)	(0.224)
D-in-D percentage impact	-63.0%	-41.3%	-56.8%	-56.7%
(Standard error of impact)	(13.0)	(13.5)	(10.6)	(9.6)

Notes: See Table 1.

The difference-in-differences estimates show that across the four different control group regions used in Table 2, there is an average 55% decline (ranging from -41% to -63%) in the luminosity of Kaesong City following the closure of the industrial zone, compared to what would have been expected from movements in luminosity in other parts of North Korea. This estimated treatment impact is similar to the average of the estimates in Table 1, so using the more spatially aggregated data does not greatly change the conclusions about the effectiveness of this particular sanction. However, using the more spatially aggregated treated unit does result in a slight reduction in precision compared to the results in Table 1, with the standard errors for $\hat{\beta}_3$ averaging about one-third the size of the coefficient. Nevertheless, the difference-in-differences effects are still statistically significant at the 5% level or 1% level.

4.1 Results Using DMSP Data

In contrast to the precisely estimated impacts with the VIIRS data, which were reported in Table 1 and Table 2, if the DMSP data are used for the same years and the same treatment and control groups, the apparent impacts on economic activity of closing the Kaesong Industrial Zone are far smaller and are imprecisely estimated. The results in Table 3 show that the estimated impacts of this sanction, when using the DMSP data, range from +23% to -68%. Across the five estimates in Table 3, the DMSP data suggest that, on average, there was a 32% decline in luminosity in Panmun District relative to what would be expected from the changes over time in other parts of North Korea. Notably, this average is just one-half of the average impact in Table 1 that was estimated using VIIRS data with the same specification for the same spatial units and with the same time period. Moreover, only one of the five estimates of $\hat{\beta}_3$ in Table 3 is statistically significant (and only at the 10% level) when using the DMSP data, whereas four of the estimates of $\hat{\beta}_3$ in Table 1 were statistically significant at the 1% level and one was significant at the 5% level, when using the VIIRS data. In other words, if an evaluation relied on DMSP data, conclusions about the impact of this sanction on North Korea are likely to be very distorted because of the underestimation of the difference-in-differences impacts.

Table 3: Using DMSP Data to Estimate Difference-in-Differences Impacts on Luminosity of Closing the Kaesong Industrial Zone: Panmun District as the Treated Area with Various Control Groups

	------------------------ Control Group ------------------------				
	Kaepyong District	Rason (Special City)	Sinuiju City	Hwanghae-namdo	Rest of North Korea
Constant, (β_0)	0.415	–1.102***	3.129***	–1.635***	-1.486***
	(0.523)	(0.330)	(0.138)	(0.414)	(0.203)
Treated group (D), (β_1)	1.637**	3.155***	-1.076***	3.688***	3.539***
	(0.601)	(0.443)	(0.326)	(0.509)	(0.359)
Sanction period (T), (β_2)	-1.284**	0.159	-0.293*	-0.506	-0.579**
	(0.579)	(0.353)	(0.142)	(0.455)	(0.206)
Treated × Sanction, (β_3)	0.427	–1.017*	–0.563	–0.351	-0.277
	(0.658)	(0.472)	(0.345)	(0.553)	(0.375)
D-in-D percentage impact	23.4%	-67.6%	-46.3%	-39.6%	-29.3%
(Standard error of impact)	(73.2)	(14.5)	(18.0)	(31.0)	(25.6)

Note: Each column represents a separate estimate of equation (1), using DMSP stable lights annual composites from 2012-19, with samples varying according to the control group used in columns (1) to (5). Other notes, see Table 1.

A consideration of the measurement error properties of DMSP data may help explain why the impact estimates in Table 3 are only about one-half as large as the VIIRS-based estimates were in Table 1. Consider a regression model: $y = \alpha + \beta x + u$, where y is an outcome variable, the independent variable is x, the response coefficient is β, and there is a pure random error, u. The outcome variable has an observed value y^* that is related to the true value by:

$$y^* = \theta + \lambda y + v \tag{2}$$

The textbook case of classical measurement error makes the assumptions that $\theta = 0, \lambda = 1$ and $E(v) = cov(y, v) = cov(x, v) = cov(u, v) = 0$, so that just white noise is added to the true value. However, existing evidence on measurement errors in DMSP data is inconsistent with these textbook assumptions. Specifically, these studies treat VIIRS data as the true value, y and the DMSP data as the mis-measured variable, y^* in order to empirically estimate equation (2).[9] With pixel-level data for urban areas that are predominantly in developing countries, Alimi et al (2022) estimate $\hat{\lambda} = 0.3$. A higher estimate, of $\hat{\lambda} = 0.7$, comes from regional data for Europe at the NUTS2 level (Gibson, 2021); this level of aggregation uses provinces in some countries and groupings of counties in others (with a median population of 1.5 million per spatial unit). When we use the second sub-national level data for North Korea to estimate equation (2) we get $\hat{\lambda} = 0.55$ (± 0.04). Thus, our estimate is mid-way between these two prior estimates, which is consistent with our use of a level of spatial aggregation that is mid-way between the spatial aggregation level used in these two prior studies.[10]

With the textbook case of classical measurement error, the $\hat{\beta}$ is not biased if the mis-measured variable is on the left-hand side of a regression. However, when the errors are mean-reverting,

[9] Support for the assumption that the VIIRS data can be used as the true value comes from Vuong (1989) tests that show that models using VIIRS data provide results that are closer to the truth than are models using DMSP data (Gibson, 2021; Zhang & Gibson, 2022).
[10] We estimated equation (2) using standardized variables, because of the different units and scales of DMSP and VIIRS data.

$0 < \lambda < 1$, which is the situation for DMSP data, the estimator of the response coefficient is biased if the error-ridden dependent variable is used. Specifically, the estimator is:

$$\beta_{y^*x} = \frac{cov(y^*,x)}{var(x)} = \frac{cov(\lambda\alpha+\lambda\beta x+\lambda u-v,x)}{var(x)} = \lambda\beta \qquad (3)$$

In other words, with mean-reverting errors in the DMSP data, the difference-in-differences estimator of equation (1), for the impact of closing the Kaesong Industrial Zone, is expected to be attenuated in proportion to the estimated mean-reversion parameter $\hat{\lambda}$. The comparison of the average impacts calculated from the results in Table 1 using VIIRS data with the average impacts in Table 3 using DMPS data are consistent with what equation (3) shows. Specifically, using DMSP data as the outcome measure for the difference-in-differences estimates will give apparent impacts that are only one-half as large as the actual impacts.

4.2 Effects of Different Treatments for False Zeros

The mean-reverting errors in the DMSP data that are causing attenuation of the impact estimates from the difference-in-differences regressions have at least two sources. One is the blurring of the DMSP images, which was shown visually in Figure 2. Another factor is the false zeros, where dimly lit areas are wrongly assessed by DMSP data as being totally unlit. The comparisons reported in Section 3 suggest that over one-half of the occurrences of zeros in the DMSP data for the second sub-national level of North Korea are false, because the VIIRS data show positive levels of luminosity in these same areas in the same year. This issue of the false zeros should matter especially to the difference-in-differences results that are using either Hwanghae-namdo or the rest of North Korea as the control group because compared to the other control groups (such as Rason and Sinuiju) a greater proportion of the observations in these two groups come from dimly-lit areas rather than from built-up urban areas. The sensitivity of results when these two control groups are used is shown by the greater attenuation in estimated impacts (based on the estimates of $\hat{\beta}_3$) when using the DMSP data reported in Table 3 when the regressions are using either Hwanghae-namdo or the rest of North Korea as the control group compared to the attenuation in the results that use more urbanized control groups such as Rason or Sinuiju.

The results for the difference-in-differences estimator, $\hat{\beta}_3$ of using various approaches to deal with zeros in the DMSP data are shown in Table 4. The first two rows of the table reproduce the coefficients reported in Table 1 using VIIRS data, and the coefficients reported in Table 3 using DMSP data. We are interested in whether any of the approaches for dealing with zeros in DMSP data move the coefficients from their Table 3 values to being closer to the Table 1 values, which we view as a movement closer to the truth given the greater accuracy of the VIIRS data. We consider four approaches to zeros in the DMSP data: (a) the false zeros are replaced by the corresponding VIIRS values (which, by definition, are non-zero if the DMSP data are a false zero), (b) the false zeros are given an *ad hoc* transformation, of adding 0.01 so that logarithms can be taken, (c) all zeros in the DMSP data—whether false or not—are given a value of 0.01, and (d) the arcsinh (or inverse-hyperbolic sine) transformation is used, consistently on all observations whether zero or not. The choices we study reflect some of the common approaches in the applied literature. For example, in his North Korea sanctions study Lee (2018) added 0.01 to all DMSP values before they were logarithmically transformed, so as to get around the problem of the logarithm of zero being undefined.[11]

Table 4: Effect on Estimates of the Key Difference-in-Differences Parameter, (β_3) of Different Approaches to Dealing With Zeros in DMSP Data

	Control Group	
	Hwanghae-namdo	Rest of North Korea
Baseline results		
1. Using VIIRS data for all observations (copied from Table 1)	–1.438***	-1.056***
	(0.414)	(0.203)
2. Using DMSP data for all observations (copied from Table 3)	–0.351	-0.277
	(0.509)	(0.359)
Treatment of zeros in DMSP data		
3. Replacing false zeros in DMSP data with VIIRS values	-1.274***	-0.922**
	(0.364)	(0.344)
4. Replacing false zeros in DMSP data with 0.01 prior to taking logs	-0.465	-0.570
	(0.445)	(0.371)
5. Replacing all zeros in DMSP data with 0.01 prior to taking logs	-0.631	-0.666*
	(0.401)	(0.365)
6. Inverse-hyperbolic sine transformation applied to all DMSP data	-0.646**	-0.631**
	(0.265)	(0.266)

Note: Each cell represents results of a separate estimate of equation (1), with the $\beta_0, \beta_1,$ and β_2 coefficient estimates not reported. The treated unit is Panmun District, and the control groups and noted in the column headings. Other notes, see Table 1.

The problem of false zeros in the DMSP data is a large contributor to the mean-reverting errors that attenuate difference-in-differences estimates of the impacts of closing the Kaesong Industrial Zone. Over 35% of the DMSP observations are for areas (and years) that seem to be completely unlit when in fact in those same times and places the VIIRS data show that they are illuminated. If these false zeros are replaced by the more accurate VIIRS value, the estimated mean-reversion parameter from equation (2), $\hat{\lambda}$ rises from 0.55 (±0.04) to 0.80 (±0.02), so with these corrected data there should be less attenuation in the difference-in-differences estimator of the sanction impacts (according to equation (3)). The results in row (3) of Table 4 bear this out, in showing that much of the attenuation of the difference-in-difference estimator is due to the false zeros, with about 85% of the gap between the row (1) and row (2) values—that is, the smaller apparent impact of closing the industrial zone when using DMSP data rather than VIIRS data—removed by replacing the false zeros with the VIIRS data. Of course this treatment is not possible if a study relies totally on the DMSP data, which is typically the case in the literature estimating impacts of sanctions.

[11] This type of *ad hoc* transformation is common in the economics literature using DMSP data. In the review by Gibson et al (2020), almost 40% of studies with DMSP night lights data on the left-hand side of regressions added a small constant to the data before creating the logarithms.

If we follow the common practice of adding a small constant, such as 0.01, to the data prior to creating the logarithm of the DMSP values, about one-quarter of the gap between the row (1) and row (2) values is closed, if this *ad hoc* transformation is just done for the false zeros (with these results in row (4)). Between one-quarter to one-half of the gap is closed if all of the DMSP zeros (including those times and places that truly are unlit, according to VIIRS) have the small constant added to them (row (5)). With the exception of the result when the rest of North Korea is used as the control group and 0.01 is added to all zeros, the resulting estimates of the closure impacts reported in Table 4 remain statistically insignificant and in all cases they are far smaller than the impacts according to the VIIRS data so these treatments for zeros in the DMSP data are not very successful at producing results that approximate the benchmark results in Table 1. Moreover, the addition of a small constant is a very arbitrary transformation, because if we were to use even smaller constants, such as 0.0001 or 0.000001, the estimate of $\hat{\beta}_3$ goes from the -0.67 value reported in row (5) to -0.85 (±0.47) and -1.13 (±0.69). This sensitivity to what is an arbitrary choice of an *ad hoc* transformation undermines the confidence that can be placed in published results that have used such an approach.

A less *ad hoc* approach is to apply the inverse-hyperbolic sine transformation to all of the DMSP values, whether they are zero or not. For non-zero observations this is equivalent to using logarithms, as has been shown in previous studies using DMSP data (Gibson et al, 2017). The results in row (6) of Table 4 show that this approach provides more precise estimates, that are statistically significant at the 5% level, but still only closes about one-third of the gap between the row (1) and row (2) values. Thus, based on the exercises reported in Table 4, finding better solutions for the issue of bottom-coded false zeros when DMSP data are used for evaluations in poorly illuminated places would be valuable.

5. Conclusions

In this paper we have used two different data sources to study a particular sanction on North Korea, the closure of the Kaesong Industrial Zone in early 2016. The widely used DMSP night-time lights data provide estimates of the impact of this sanction which are far smaller, and are generally imprecise, than what is revealed by the more accurate VIIRS data. The understated impact of the sanction when DMSP data are used is due to two main issues: (a) mean-reverting errors, which are likely due to blurring caused by poor spatial resolution of the DMSP sensor and by limited on-board data storage that result in pixel aggregation, and (b) false zeros, which are due to the poor low-light detection capability, recalling that the original objective of DMSP was to observe clouds and little amplification of the sensor is needed for this purpose during the brightly lit part of the lunar cycle. While previous studies have noted that DMSP data have some measurement errors, the implications of these errors for popular econometric models such as the difference-in-differences estimator have not been highlighted. To the extent that the patterns that we find hold more generally, conclusions about the impact of sanctions that have relied upon DMSP data may need to be revisited because it is likely that the impacts have been understated due to the mean-reverting measurement errors in DMSP data.

References

Abrahams, Alexei, Christopher Oram and Nancy Lozano-Gracia (2018) Deblurring DMSP night-time lights: A new method using Gaussian filters and frequencies of illumination. *Remote Sensing of Environment* 210(1): 242-258.

Arbatli, Cemal Eren, and David Gomtsyan (2021) Sectarian aid, sanctions and subnational development. *European Economic Review* 139(1): 103891.

Baugh, Kimberly, Christopher Elvidge, Tilottama Ghosh, and Daniel Ziskin (2010) Development of a 2009 stable lights product using DMSP-OLS data. *Proceedings of the Asia-Pacific Advanced Network* 30(1): 114.

Bellemare, Marc, and Casey Wichman (2020) Elasticities and the inverse hyperbolic sine transformation. *Oxford Bulletin of Economics and Statistics* 82(1): 50-61.

Bluhm, Richard, and Melanie Krause (2020) Top lights - bright cities and their contribution to economic development. *SoDA Laboratories Working Paper* No. 2020-08.

Bound, John, Charles Brown, and Nancy Mathiowetz (2001) Measurement error in survey data, in James Heckman and Edward Leamer (editors) *Handbook of Econometrics*, vol. 5, pp. 3705-3843. Elsevier, Amsterdam.

Chen, Xi, and William Nordhaus (2015) A test of the new VIIRS lights data set: Population and economic output in Africa. *Remote Sensing* 7(4): 4937-4947.

Ditzen, Jan, Yiannis Karavias, and Joakim Westerlund (2021) Xtbreak: Estimating and testing for structural breaks in Stata. https://janditzen.github.io/xtbreak/

Doll, Christopher (2008) *CIESIN Thematic Guide to Night-time Light Remote Sensing and its Applications*. Center for International Earth Science Information Network, Columbia University, New York.

Elvidge, Christopher, Kimberly Baugh, John Dietz, Theodore Bland, Paul Sutton, and Herbert Kroehl (1999) Radiance calibration of DMSP-OLS low-light imaging data of human settlements. *Remote Sensing of Environment* 68(1): 77-88.

Elvidge, Christopher, Kimberly Baugh, Mikhail Zhizhin, and Feng-Chi Hsu (2013) Why VIIRS data are superior to DMSP for mapping night time lights. *Proceedings of the Asia-Pacific Advanced Network* 35(1): 62-69.

Elvidge, Christopher, Kimberly Baugh, Mikhail Zhizhin, Feng Chi Hsu, and Tilottama Ghosh (2017) VIIRS night-time lights. *International Journal of Remote Sensing* 38(21): 5860-5879.

Elvidge, Christopher, Mikhail Zhizhin, Tilottama Ghosh, Feng-Chi Hsu, and Jay Taneja (2021) Annual time series of global VIIRS nighttime lights derived from monthly averages: 2012 to 2019. *Remote Sensing* 13(5): 922

Falchi, Fabio, and Pierantonio Cinzano (1998) Maps of artificial sky brightness and upward emission in Italy from DMSP measurements. *arXiv preprint astro-ph/9811234*.

Farzanegan, Mohammad Reza, and Sven Fischer (2021) Lifting of international sanctions and the shadow economy in Iran-A view from outer space. *Remote Sensing* 13(14): 4620.

Ghosh, Tilottama, Kimberly Baugh, Christopher Elvidge, Mikhail Zhizhin, Alexey Poyda, and Feng-Chi Hsu (2021) Extending the DMSP nighttime lights time series beyond 2013. *Remote Sensing* 13(24): 5004.

Gibson, John (2021) Better night lights data, for longer. *Oxford Bulletin of Economics and Statistics* 83(3): 770-791.

Gibson, John, and Geua Boe-Gibson (2021) Nighttime lights and county-level economic activity in the United States: 2001 to 2019. *Remote Sensing* 13(14): 2741.

Gibson, John, Susan Olivia and Geua Boe-Gibson (2020) Night lights in economics: Sources and uses. *Journal of Economic Surveys* 34(5): 955-980.

Gibson, John, Susan Olivia, Geua Boe-Gibson and Chao Li (2021) Which night lights data should we use in economics, and where? *Journal of Development Economics* 149(1): 102602.

Gibson, John, Gaurav Datt, Rinku Murgai, and Martin Ravallion (2017) For India's rural poor, growing towns matter more than growing cities. *World Development* 98(1): 413-429.

Heger, Martin Philipp, and Eric Neumayer (2019) The impact of the Indian Ocean tsunami on Aceh's long-term economic growth. *Journal of Development Economics* 141(1): 102365.

Henderson, Vernon, Adam Storeygard, and David Weil (2012) Measuring economic growth from outer space. *American Economic Review* 102(2): 994-1028.

Kim, Bonggeun, and Gary Solon (2005) Implications of mean-reverting measurement error for longitudinal studies of wages and employment. *Review of Economics and Statistics* 87(1): 193-196.

Kim, Jihee, Kyoochul Kim, Sangyoon Park, and Chang Sun (2021) The economic costs of trade sanctions: Evidence from North Korea. *Mimeo* http://jiheekim.net/files/paper_NorthKorea.pdf

Lee, Yong Suk (2018) International isolation and regional inequality: Evidence from sanctions on North Korea *Journal of Urban Economics* 103(1): 34-51.

Miller, Steven, Stephen Mills, Christopher Elvidge, Daniel Lindsey, Thomas Lee and Jeffrey Hawkins (2012) Suomi satellite brings to light a unique frontier of nighttime environmental sensing capabilities. *Proceedings of the National Academy of Sciences* 109(39): 15706-15711.

Nordhaus, William and Xi Chen (2015) A sharper image? Estimates of the precision of night time lights as a proxy for economic statistics. *Journal of Economic Geography* 15(1): 217-246.

Son, Sung Hyun and Joonmo Cho (2021) Are economic sanctions against North Korea effective? Assessing nighttime light in 25 major cities. *Pacific Affairs* 94(3): 464-492.

Tuttle, Benjamin, Sharolyn Anderson, Paul Sutton, Christopher Elvidge, and Kim Baugh (2013) It used to be dark here. *Photogrammetric Engineering & Remote Sensing* *79*(3): 287-297.

van Garderen, Kees and Chandra Shah (2002) Exact interpretation of dummy variables in semilogarithmic equations. *The Econometrics Journal* 5(1): 149-159.

Vuong, Quang (1989) Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica* 57(2): 307-333.

Zhang, Xiaoxuan and John Gibson (2022) Using multi-source nighttime lights data to proxy for county-level economic activity in China from 2012 to 2019. *Remote Sensing* 14(5): 1282.