The Bright Side of Heteroskedasticity: Measuring Quarterly Economic Growth from Outer Space

Robert C. M. Beyer*, Yingyao Hu† and Jiaxiong Yao†

June 29, 2026

Abstract

Satellite-recorded nighttime lights (NTL) serve as a powerful proxy for economic activity and are increasingly used in a wide array of economic studies. This paper introduces a new framework to estimate the elasticity between NTL and economic activity in the presence of measurement error on both sides. Unlike traditional approaches, our framework achieves identification through the heteroskedasticity of measurement errors in NTL—a feature fundamentally grounded in the satellite data-generating process. Using the latest generation of NTL data, we estimate that the elasticity is around 1.52 at the annual frequency and 2.25 at the quarterly frequency. The estimated elasticity varies only slightly across country groups of different income status and sectoral composition.

Keywords: Nighttime lights, GDP, elasticity, heteroskedasticity.
JEL Classification: C10, E01, R12

*International Monetary Fund, 700 19th St NW, Washington DC 20431 (e-mail: rbeyer@imf.org and jyao@imf.org).

†Johns Hopkins University, Department of Economics, Johns Hopkins University, Wyman Park Building 544E, 3400 N. Charles Street, Baltimore, MD 21218, (email: yhu@jhu.edu).

1 Introduction

Economic research and evidence-based policymaking depend fundamentally on the precise measurement of economic activity, typically quantified through Gross Domestic Product (GDP). Despite its importance, official GDP data suffers from significant gaps in both temporal frequency and spatial granularity. Approximately one-third of nations do not publish quarterly GDP figures, and an even greater number lack subnational data entirely.[1] Where official records exist, they are frequently marred by substantial revisions, prolonged reporting lags, and a lack of geographic detail. These persistent data constraints have prompted the exploration of alternative methods to measure economic activity.

Since the seminal work of Chen and Nordhaus (2011) and Henderson, Storeygard, and Weil (2012), satellite-recorded nighttime light (NTL) data have emerged as a compelling proxy for economic activity. Such data have been increasingly used in a wide array of studies addressing questions related to economic growth(Nordhaus and Chen, 2015; Keola, Andersson, and Hall, 2015; Pinkovskiy and Sala-i Martin, 2016; Morris and Zhang, 2019; Hu and Yao, 2022), regional development (Lessmann and Seidel, 2017; Chanda and Kabiraj, 2020), trade (Henderson, Squires, Storeygard, and Weil, 2018; Hirte, Lessmann, and Seidel, 2020; Chor and Li, 2024), poverty (Jean, Burke, Xie, Alampay Davis, Lobell, and Ermon, 2016; Yeh, Perez, Driscoll, Azzari, Tang, Lobell, Ermon, and Burke, 2020), culture (Gershman, 2020), political economy (Anaxagorou, Efthyvoulou, and Sarantides, 2020; Iacoella, Martorano, Metzger, and Sanfilippo, 2021; Arbatli and Gomtsyan, 2021), and many other areas.

A key component of gaining economic insights from NTL data entails utilizing an estimated elasticity to transform fluctuations in NTL into changes of economic activity. As official GDP measures and NTL both contain unknown measurement errors, identification and estimation of the elasticity are challenging. Consequently, past studies have dealt with the under-identification issue by making assumptions about the signal-to-noise ratios of the data (Henderson, Storeygard, and Weil, 2012), imposing complex structures of measurement errors (Hu and Yao, 2022), or adding an additional signal of GDP (Civelli, Gaduh, and Yousuf, 2024).

In this paper, we provide a new framework for estimating the elasticity between NTL and GDP, achieving identification through heteroskedasticity. Specifically, we show that the measurement errors in NTL are heteroskedastic in the size of a country,

[1] The 2020 IMF's National Accounts Statistics Program survey shows that 63 out of 206 countries surveyed do not produce any quarterly GDP statistics (Silungwe, Baer, and Guerreiro, 2022).

the number of observations per period, and the average intensity of NTL. The heteroskedasticity provides us with additional moment conditions that solve the underidentification problem.

Intuitively, NTL is the sum of nighttime lights across all pixels in a country averaged over a period of time. Each pixel-level light value contains idiosyncratic measurement errors. For countries with a large area, such idiosyncratic errors tend to average out as a result of the law of large numbers. Similarly, within a period of time, as the number of NTL observations increases, the idiosyncratic measurement errors tend to average out. In addition, if the intensity of NTL per pixel is stronger, the noise from the idiosyncratic measurement errors would be smaller compared to the signal that reflects economic activity. As such the signal-to-noise ratio of NTL should be larger for brighter regions. As a result, the measurement errors in NTL are heteroskedastic.

Identification through heteroskedasticity has long been used in monetary economics (Rigobon, 2003; Rigobon and Sack, 2004), where the heteroskedasticity of structural shocks occurs through two regimes and allows for a sample split such that one subsample can act as a probabilistic instrument for the other. In our case, the heteroskedasticity of the measurement errors in NTL allow us to cast each country as a subsample. However, our framework differs from early work in that only one shock—the measurement errors in NTL—is heteroskedastistic.

Our framework based on observable information about measurement errors is applicable to a wider range of settings. It provides a general estimation strategy to gauge the relationship between satellite data and economic activity, because satellite data are usually affected by atmospheric conditions, which would cause heteroskedasticity in their measurement errors. Our estimated elasticity is useful for quantifying changes in economic activity. We show that it can be used in a difference-in-differences style policy analysis as well as in constructing an improved measure of economic activity.

To date, most economic analyses using NTL are still based on the first generation of NTL data from the Defense Meteorological Satellite Program-Operational Linescan System (DMSP-OLS) that covers the period of 1992-2013. Surprisingly few studies employ the latest generation of NTL data from the Visible Infrared Imaging Radiometer Suite (VIIRS), which started in 2012. A recent survey by Gibson, Olivia, and Boe-Gibson (2020) reviews more than 150 economic studies using NTL and finds that they overwhelmingly use DMSP-OLS data. The new NTL data are available at monthly frequency and at high resolution with improved low light imaging, but their limited temporal coverage has sometimes constrained their potential for long-term research applications. Estimates of the elasticity between VIIRS NTL and economic activity over a long period of time and at high frequencies are largely missing in the literature.[2]

In this paper, we apply our new framework to the new NTL data at both annual and quarterly frequencies. First, we utilize the annual simulated VIIRS dataset from Chen, Wang, Zhang, Shen, and Chen (2024) to build a comprehensive panel of 181 countries spanning three decades (1993–2023). The estimated elasticity from this annual-frequency dataset serves as a benchmark for comparison with previous studies that relied on annual DMSP-OLS data. Second, having established this annual-level validation, we transition to a higher temporal resolution by developing a quarterly panel of NTL and GDP growth covering 107 countries from 2013Q2-2025Q2. This quarterly analysis introduces an additional econometric challenge because of seasonality: the number of NTL observations varies substantially in different quarters and GDP growth displays strong seasonal patterns. We address these issues by modifying the moment conditions and focus on year-on-year growth.

For the annual dataset, we estimate that the elasticity between NTL and GDP is 1.52, implying that 1 percent change in GDP is associated with a 1.52 percent change in NTL. The implied signal-to-total variance ratio of GDP growth is 96% for advanced economies, 62% for emerging markets, and 56% for low income countries, confirming the strong signal content of GDP statistics for countries at high levels of economic development and suggesting the usefulness of NTL to improve GDP statistics for countries at low levels of economic development. While the elasticity varies moderately with a country's income status and sectoral composition, the deviation from the global sample average estimate is small.

For the quarterly dataset, we estimate that the elasticity is 2.25. While higher than the annual point estimate, it is not statistically different. The signal to total variance ratio of GDP growth is much lower, indicating the noisiness of GDP data at higher frequencies. As with the annual dataset, the variation of the elasticity across income status and sectoral composition is small.

We conduct extensive robustness checks for both the annual and the quarterly dataset. We show that winsorization of the data is essential for a stable elasticity estimate. Removing the top and bottom 5 percent of the NTL growth distribution yields a robust estimate of the elasticity for both annual and quarterly data.

The rest of the paper is organized as follows. In Section 2, we present our new approach to estimate the elasticity between NTL and GDP and discuss its applications. Section 3 describes the data and important stylized facts. Section 4 presents the

[2] A notable exception is the recent work by Chor and Li (2021), which provides estimates of the Annual inverse elasticity of VIIRS nighttime light intensity in China.

estimation results and Section 5 concludes.

2 Elasticity between Nighttime Lights and Economic Activity

The elasticity is key to quantifying the relationship between NTL and economic activity. However, both NTL and GDP are measured with errors, preventing the estimation of the elasticity without making additional assumptions about measurement errors or introducing new information. In this section, we provide a new framework to identify this elasticity through heteroskedasticity, employing observable information on varying measurement errors in NTL across countries.

2.1 The problem

We assume that NTL growth and GDP growth are both measured with error and have a linear structural relationship.

$$z_{i,t} = z^*_{i,t} + \varepsilon^z_{i,t}, \tag{1}$$

$$y_{i,t} = y^*_{i,t} + \varepsilon^y_{i,t}, \tag{2}$$

$$z^*_{i,t} = \beta y^*_{i,t} + \eta_{i,t}. \tag{3}$$

$z_{i,t}$ and $z^*_{i,t}$ are the measured and true growth rates of NTL of country i in year t. Similarly, $y_{i,t}$ and $y^*_{i,t}$ are measured and true GDP growth rates. Throughout, we define growth rates as the first difference in log levels so that the structural parameter β can be interpreted as the elasticity between NTL and GDP. $\varepsilon^z_{i,t}$ and $\varepsilon^y_{i,t}$ are measurement errors of NTL and GDP growth rates. $\eta_{i,t}$ is the measurement error in the structural relationship. All measurement errors are assumed to be classical with a zero mean and a finite variance.

The assumptions in equations (1)-(3) are standard in the literature. Chen and Nordhaus (2011) make the same assumptions using log levels at the country or grid cell level at the annual frequency. Henderson, Storeygard, and Weil (2012); Hu and Yao (2022) assume the same relationship at the annual frequency combining the error terms $\varepsilon^z_{i,t}$ and $\eta_{i,t}$ by substituting equation (3) for z^* in equation (1):

$$z_{i,t} = \beta y^*_{i,t} + (\varepsilon^z_{i,t} + \eta_{i,t}) \equiv \beta y^*_{i,t} + \mu^z_{i,t}. \tag{4}$$

Under the assumption that measurement errors are classical, the key structural parameter β cannot be identified by the second moments of equations (2) and (4). To see this, we take the second moments of equations (2) and (4):

$$\text{var}(y) = \text{var}(y^*) + \text{var}(\varepsilon^y) \tag{5}$$

$$\begin{aligned}\text{var}(z) &= \beta^2\text{var}(y^*) + \text{var}(\varepsilon^z) + \text{var}(\eta^z) \\ &\equiv \beta^2\text{var}(y^*) + \text{var}(\mu^z)\end{aligned} \tag{6}$$

$$\text{cov}(y, z) = \beta\text{var}(y^*), \tag{7}$$

where we have omitted the subscripts i and t. The system of equations (5)-(7) is not identified because there are three variance-covariance restrictions but four unknowns $(\beta, \text{var}(y^*), \text{var}(\varepsilon^y), \text{var}(\mu^z))$.

To identify β, additional information is needed. One way is to assume some signal-to-noise ratio for GDP, as in Henderson, Storeygard, and Weil (2012). Another way is to impose some relationships that link measurement errors to observables, such as linking nighttime light measurement errors to location and GDP measurement errors to a country's statistical capacity, as in Hu and Yao (2022). A third approach is to add another signal of GDP to overcome the underidentification problem, as in Civelli, Gaduh, and Yousuf (2024).

2.2 Identification through heteroskedasticity

In this paper, we take a different approach and achieve identification through heteroskedasticity of the noise in nighttime lights $\mu^z_{i,t}$.

Notice that $\mu^z_{i,t}$ consists of two components: the measurement error of NTL growth, $\varepsilon^z_{i,t}$, and the measurement error in the relationship between NTL growth and GDP growth, $\eta_{i,t}$. The former is heteroskedastic across countries as a result of spatial and temporal aggregation as well as the average intensity of NTL.

Intuitively, NTL is the sum of nighttime lights across all pixels in a country averaged over a year. Each pixel-level light value has idiosyncratic measurement errors. As a country's size becomes larger, such idiosyncratic errors would average out as a result of the law of large numbers. In the same vein, as the number of daily observations increases, the idiosyncratic measurement errors on each day would average out over time. In addition, if the intensity of NTL per pixel is stronger, the noise from the idiosyncratic measurement errors would be smaller compared to the signal that reflects economic activity. As such the signal to noise ratio of NTL growth should be larger for brighter regions. In Appendix A, we formally show that the variance of the measurement error of NTL growth is approximately proportional to the inverse of the size of a country, the number of daily observations per year, and the squared average intensity of NTL per pixel.

$$\text{var}(\varepsilon^z_{i,t}) \approx \text{var}_i(\varepsilon^z) = \frac{2\sigma^2_z}{\text{Area}_i \cdot N^{obs}_i \cdot (\text{E}z_i)^2}, \tag{8}$$

where σ^2_z is the variance of the pixel-level NTL idiosyncratic error and $\text{E}z_i$ is the average NTL per pixel for country i. The variance of NTL growth is therefore heteroskedastic across countries.

To see how heteroskedasticity helps us achieve identification of the elasticity between NTL and GDP, we replace equation (6) with a country-specific variance

$$\text{var}_i(z) = \beta^2\text{var}(y^*) + \frac{2\sigma^2_z}{\text{Area}_i \cdot N^{obs}_i \cdot (\text{E}z_i)^2} + \text{var}(\eta^z). \tag{9}$$

The system of equations (5), (7), and (9) has five unknowns (β,var(y^*), var(ε^y), σ^2_z, and var(η^z)). Since we have an equation (9) for each country, the system is identified with at least three countries. Note that our identification through heteroskedasticity is different from early work by Rigobon (2003). In their paper, structural shocks are uncorrelated and heteroskedasticity occurs through two regimes in which the variances of all structural shocks differ. By contrast, the implied structural shocks in our case are correlated and only measurement errors of NTL growth are heteroskedastic.[3]

2.3 Dealing with quarterly data

Estimating the elasticity between NTL and GDP at the quarterly frequency poses an additional challenge because of strong seasonality. In principle, we could apply the same framework developed above to the quarterly frequency, replacing annual growth rates with quarter-on-quarter growth rates of NTL and GDP.

However, meteorological and astronomical seasonality implies that the number of NTL observations in each quarter would differ substantially. For example, with the sun setting late in the summer, countries in high latitudes would typically have much fewer observations than in the winter. This means that we cannot simply use N^{obs}_i in equation (9), as it varies substantially across quarters. Moreover, light-generating

[3]To see this, rewriting equations (1) and (4) in terms of the relationship between y and z:

$$z_{i,t} = \beta y_{i,t} + (\mu^z_{i,t} - \beta\varepsilon^y_{i,t})$$

$$y_{i,t} = \frac{1}{\beta}z_{i,t} + (\varepsilon^y_{i,t} - \frac{1}{\beta}\mu^z_{i,t})$$

Cast in such a demand-supply system, the structural shocks to y and z are correlated while the product of the structural parameters is equal to 1 ($\beta * \frac{1}{\beta}$), which Rigobon (2003) does not allow.

activities display strong seasonal patterns, making quarter-on-quarter growth rates noisy and uninformative about the change in economic activity.

To address these issues, we focus on the same quarter year-on-year growth rates of NTL and GDP and make equation (9) quarter-specific:

$$\mathrm{var}_{i,q}(z) = \beta^2 \mathrm{var}(y^*) + \frac{2\sigma_z^2}{\mathrm{Area}_i \cdot N_{i,q}^{obs} \cdot (\mathrm{E}z_{i,q})^2} + \mathrm{var}(\eta^z), \tag{10}$$

where $q = 1, 2, 3, 4$ indicate differnet quarters. z, y and y^* in equations (5), (7), and (10) should be interpreted as year-on-year growth rates.

2.4 Why is the elasticity important?

With NTL gaining prominence as a measure of economic activity, it is important to quantify changes in economic activity through changes in NTL . Below we outline two practical application settings where the elasticity is useful.

2.4.1 Structural vs. predicative relationships

In policy analyses where GDP growth data is missing, the causal effect of a policy change on economic activity is often estimated in panel data using NTL growth as a proxy. In such applications, a classical approach is to use the difference-in-differences regression:

$$z_{i,t} = \psi_0 + \psi_1 \mathrm{Treat}_i + \psi_2 \mathrm{Post}_t + \psi_3 (\mathrm{Treat}_i \times \mathrm{Post}_t) + \epsilon_{i,t}, \tag{11}$$

where Treat is the dummy variable indicating the treatment units (unit affected by the policy change), Post is the dummy variable indicating the treatment time (time of the policy change), and $\epsilon_{i,t}$ is the error term. Because the error term captures the noises in NTL growth, the treatment effect of a policy (ψ_3) should be interpreted as the policy's impact on true NTL growth. As such, to obtain the policy's impact on GDP growth, we should divide the treatment effect by the structural elasticity ($\hat{\psi}_3/\hat{\beta}$).

Other applications are concerned with using NTL growth to predict GDP growth. In those applications, we would conduct the following regression of GDP growth on NTL growth:

$$y_{i,t} = \gamma z_{i,t} + e_{i,t}. \tag{12}$$

where γ is the NTL predictive coefficient and $e_{i,t}$ is the error term. For regions where NTL growth is known but GDP growth is unknown, we could simply multiply the NTL growth with the estimated predictive coefficient $\hat{\gamma}$ to obtain the estimates of

GDP growth for those regions. In such applications, the structural elasticity is not needed because equation 12 already provides a best fit relationship.

To see the relationship between the structural elasticity β and the predictive coefficient γ, notice that γ is essentially $\text{cov}(y,z)/\text{var}(z)$, therefore:

$$\text{plim}\,\gamma = \frac{\beta \text{var}(y^*)}{\beta^2 \text{var}(y^*) + \text{var}(\mu^z)} = \frac{1}{\beta}\left(\frac{\beta^2 \text{var}(y^*)}{\beta^2 \text{var}(y^*) + \text{var}(\mu^z)}\right). \tag{13}$$

γ is the inverse of the elasticity, $1/\beta$, adjusted downward by the signal ratio of NTL growth.

2.4.2 Light-Adjusted GDP Measure

Some applications are interested in an improved measure of GDP growth that combines information in official GDP growth and NTL growth. A light-adjusted GDP measure that is a linear combination of official GDP growth and NTL-predicted GDP growth is often considered:

$$\bar{y} = (1-\lambda)y + \lambda\tilde{y}, \tag{14}$$

where $\tilde{y}$ is the NTL-predicted GDP growth ($\tilde{y}_{i,t} = \gamma z_{i,t}$) and λ is the weight on it.

The optimal weight λ^* is chosen such that the expected mean squared prediction error of true GDP growth is minimized. Appendix B shows that the optimal weight depends on the structural elasticity, the variance of true GDP growth, and the variances of the measurement errors in GDP growth and NTL growth—all of which are identified in Section 2.2:

$$\lambda^* = \frac{\text{var}(\varepsilon^y)}{\text{var}(\varepsilon^y) + (\gamma\beta - 1)^2\text{var}(y^*) + \gamma^2\text{var}(\mu^z)}. \tag{15}$$

Intuitively, the optimal weight gauges the importance of the measurement error in GDP growth in relation to the prediction error of true GDP growth. If the measurement error in official GDP growth is larger, then the light-adjusted GDP growth would optimally a higher weight on NTL-predicted GDP growth and a lower weight on official GDP growth.

The light-adjusted GDP growth proves useful especially in countries where reliable and accurate official GDP figures are lacking.

3 Data and Stylized Facts

3.1 Nighttime Lights

Nighttime light (NTL) data have changed over two different generations of satellites. The first generation of nighttime light data, based on the Defense Meteorological Satellite Program-Operational Linescan System (DMSP-OLS), covers 1992 to 2013. The second generation, based on the Visible and Infrared Imaging Suite (VIIRS) Day Night Band on board onboard the Joint Polar Orbiting Satellite System, covers April 2012 to the present at monthly frequency. Compared with DMSP-OLS, VIIRS nighttime lights have higher resolution and benefit from much improved low light imaging (Elvidge, Baugh, Zhizhin, and Hsu, 2013).[4]

Recently, efforts have been made to synthesize consistent NTL time series from the DMSP-OLS and VIIRS. For example, Chen, Wang, Zhang, Shen, and Chen (2024) generate a continuous and consistent 500-meter global annual simulated VIIRS NTL dataset (SVNL) from 1992 to 2023, using image super-resolution deep learning techniques to transform low-resolution DMSP-NTL data into VIIRS NTL data.

In this paper, we use the SVNL dataset for the annual frequency analysis. It allows us to examine the relationship between high-resolution NTL and economic activity over a long period of time. To obtain country-level NTL, we use the country borders defined by the Database of Global Administrative Areas (GADM, version 4.1) and calculate the sum of NTL pixel values within each country's borders. GADM also contains data on the area of each country, which we use in our estimation.

For the quarterly analysis, we use the VIIRS Nighttime Day/Night Band Composites (Version 1), produced by the Earth Observation Group in the Colorado School of Mines and hosted by the Google Earth Engine. The dataset is monthly from April 2012 to the present. We use the monthly average NTL for a given quarter as the quarterly NTL.

To obtain the number of daily NTL observations per year and per quarter, we use the same VIIRS data, which has a band that contains information on the total number of observations that went into each pixel each month. For each country we take the average number of daily observations across all pixels. While meteorological and astronomical seasonality affects the number of daily NTL observations, the seasonality is usually stable and predictable. As such, for annual data we use the average between 2013-2023 and for the quarterly data we use the same-quarter average over these years.

3.2 GDP Growth, GDP per Capita, and Sectoral Shares

For annual GDP growth rates, we use data from the World Economic Outlook database of the International Monetary Fund (IMF). We transform the growth rates into the first difference in log levels, consistent with the framework in Section 2.1. For quarterly

[4] For more details on different versions of VIIRS nighttime light data, please visit `https://eogdata.mines.edu/products/vnl/`.

GDP rates, we use seasonally adjusted year-on-year growth rates compiled by Haver Analytics.

As we explore the elasticity between NTL and GDP across various subgroups, we distinguish countries based on their income status, level of economic development, and economic structures.

For the income status of countries, we use the IMF's definition of advanced economies (AEs), emerging markets and middle-income countries (EMs), and low-income developing countries (LIDCs) to divide countries into mutually exclusive groups. We also combine the last two into a single group: emerging market and developing economies (EMDEs).

For the level of economic development, we rely on the World Development Indicators (WDI) database from the World Bank and use purchasing power parity-adjusted GDP per capita in constant 2021 international dollars. For economic structures, we use the shares of agriculture, manufacturing, and services in each economy from the same database.

3.3 Summary Statistics

Our analysis focuses on two datasets. The first is a panel dataset of NTL growth and GDP growth at the annual frequency. It covers 181 countries from 1993 to 2023, including 36 AEs, 90 EMs, and 55 LIDCs. We call this the SVNL annual dataset. The second is a a panel dataset at the quarterly frequency, covering 107 countries from 2013Q2-2025Q2, including 34 AEs, 57 EMs, and 16 LIDCs. We call this the VIIRS quarterly dataset. It is worth noting that the smaller number of countries in the VIIRS quarterly dataset is entirely driven by the availability of quarterly GDP data, as NTL has near global and high frequency coverage.

Table 1 presents the summary statistics of the SVNL annual dataset by countries' income status. Across all country groups, the average NTL growth is higher than the average GDP growth. Comparing AEs, EMs, and LIDCs, we see that both NTL growth and GDP growth are higher in countries of lower income status.

NTL growth is typically an order of magnitude more volatile than GDP growth, as is clear from the much wider interquartile range of NTL growth rates. Notice that at the extreme ends, NTL growth is below -61 percent at the bottom 1 percent of the entire sample while above 85 percent at the top 1 percent of the sample. Such outliers disproportionately affect the relationship between NTL growth and GDP growth. Winsorization of data is therefore needed. In subsequent analysis of the SVNL annual dataset, we discard the top and bottom 5 percent of the NTL growth distribution as well as the top and bottom 1 percent of the GDP growth distribution. In the robustness check, we show that such a choice is well grounded.

Table 1: SVNL Annual Dataset: Summary Statistics by Income Classification

	Mean	p1	p10	p25	p50	p75	p90	p99
AEs								
NTL growth	3.2	-67.4	-18.2	-6.6	1.3	11.1	24.8	92.4
GDP growth	2.5	-8.9	-1.1	1.1	2.6	4.2	6.3	10.8
avg. # of obs per year	87.2	47.1	54.0	62.3	73.4	109.3	143.6	178.4
area (thousand sq. km)	1184.7	0.07	3.3	46.3	98.3	458.8	790.9	16864.8
avg. pixel NTL (radiance)	2.3	0.01	0.09	0.3	0.7	1.9	4.6	32.5
EMs								
NTL growth	6.4	-61.1	-13.7	-3.6	5.0	14.6	28.9	84.1
GDP growth	3.3	-15.3	-2.0	1.2	3.6	5.7	8.1	18.6
avg. # of obs per year	109.2	30.9	55.8	77.7	102.2	139.4	171.8	193.9
area (thousand sq. km)	936.8	0.04	0.4	10.4	99.3	722.5	1734.6	29311.1
avg. pixel NTL (radiance)	1.0	0.002	0.02	0.07	0.2	0.4	1.5	20.3
LIDCs								
NTL growth	7.4	-54.2	-16.7	-3.7	6.6	18.1	32.0	78.4
GDP growth	3.7	-17.7	-1.4	2.2	4.5	6.4	8.3	13.9
avg. # of obs per year	101.2	43.1	59.9	71.6	103.9	126.1	144.9	159.3
area (thousand sq. km)	360.8	1.4	21.9	59.4	197.2	507.2	928.7	1900.0
avg. pixel NTL (radiance)	0.04	0.0003	0.002	0.005	0.01	0.04	0.09	0.3
Total								
NTL growth	6.0	-61.1	-15.5	-4.3	4.5	15.1	29.5	85.2
GDP growth	3.3	-14.5	-1.7	1.4	3.6	5.7	7.9	15.9
avg. # of obs per year	102.4	39.7	56.2	71.3	98.3	129.6	159.1	190.1
area (thousand sq. km)	811.7	0.07	1.7	23.2	135.6	535.6	1476.5	16864.8
avg. pixel NTL (radiance)	1.0	0.0008	0.005	0.02	0.1	0.5	1.7	18.8

Notes. This table presents summary statistics of NTL growth and GDP growth between 1993-2023 by country group. The mutually exclusive country groups are advanced economies (AEs), emerging markets and middle-income countries (EMs), and low-income developing countries (LIDCs). Growth rates are defined as 100 times the first different in log levels. The unit of the average pixel NTL is $nWcm^{-2}sr^{-1}$.

Table 1 also provides the distribution of the average number of observations per year, area, and the average intensity of NTLs per pixel (radiance). Since most AEs are located in the northern part of the Earth, their NTLs are affected by late sunsets in summer months and therefore have a lower number of observations on average than EMs and LIDCs. The average AE has about 87 observations per year, whereas EMs and LIDcs have 109 and 101 observations per year, respectively. As NTLs reflect economic activity, the average pixel intensity is the highest in AEs, followed by EMs and then LIDCs, in accordance with their levels of economic development.

The large variations in the number of NTL observations, area, and the average NTL intensity across countries play a central role in the heteroskedasticity of the measurement errors in NTL growth, which allows us to identify the elasticity between NTL growth and GDP growth.

Table 2 presents the summary statistics of the quarterly dataset by quarter. The first column reveals that NTL growth displays strong seasonal patterns. The average NTL growth rate is low in the first and third quarters but high in the second and fourth quarters. The average number of observations also vary substantially over quarters. It falls below 20 days on average in the second quarter but rises to about 30 days in the fourth quarter. As a comparison, year-on-year GDP growth and the average pixel light intensity show much less variation over different quarters.

While the average quarterly year-on-year NTL growth rates are similar to the average annual rates in Table 1, the quarterly growth rates have much fatter tails. For example, the 1st percentile of the total sample shows much more negative growth rates (log difference) of NTL while the 99th percentile shows much higher growth rates. This indicates a potentially greater amount of noise in NTL growth at the quarterly frequency. Similarly, the average number of NTL observations per quarter also has much wider distributions. The 1st percentile of the total sample only has 0.5 days of observation per quarter, which greatly limits the accuracy for NTL measures. In light of this, for subsequent analysis of the VIIRS quarterly dataset, we drop country-quarter observations for which the number of NTL observations is no greater than 5, in addition to discarding the top and bottom 5 percent of the NTL growth distribution and the top and bottom 1 percent of the GDP growth distribution as is done for the SVNL annual dataset.

3.4 Stylized Facts

In this section, we first provide some evidence of the heteroskedasticity of the measurement errors of NTL growth, and then examine the structural and predictive relationships between NTL growth and GDP growth.

3.4.1 Heteroskedasticity of Measurement Errors in NTL Growth

Since the measurement error of NTL growth ε^z in equation (1) is not directly observable, we examine the variance of NTL growth $\mathrm{var}(z)$. If ε^z is heteroskedastic across different countries, by equation (9) we should observe that the variance of NTL growth

Table 2: VIIRS Quarterly Dataset: Summary Statistics by Quarter

	Mean	p1	p10	p25	p50	p75	p90	p99
1st quarter								
NTL growth (yoy)	3.8	-86.2	-24.1	-7.2	3.8	14.7	31.5	107.0
GDP growth (yoy)	2.7	-7.1	-1.1	1.0	2.7	4.6	6.5	13.3
avg. # of obs	27.9	7.4	13.6	24.1	27.4	33.6	38.0	43.6
avg. pixel NTL	0.9	0.02	0.09	0.2	0.3	0.5	1.3	12.9
2nd quarter								
NTL growth (yoy)	9.6	-200.5	-22.4	-4.3	5.6	19.0	45.0	272.6
GDP growth (yoy)	2.8	-22.9	-2.6	1.2	3.2	5.4	8.1	20.8
avg. # of obs	18.3	0.1	2.9	7.3	13.4	29.4	43.1	48.7
avg. pixel NTL	0.8	0.001	0.04	0.09	0.2	0.3	0.8	19.5
3rd quarter								
NTL growth (yoy)	4.9	-150.0	-31.7	-9.8	4.3	18.5	44.6	157.7
GDP growth (yoy)	2.8	-10.9	-1.9	1.2	3.0	5.0	7.1	15.4
avg. # of obs	22.1	0.7	4.9	10.7	17.3	33.7	45.4	53.9
avg. pixel NTL	0.7	0.005	0.04	0.09	0.2	0.3	0.8	18.8
4th quarter								
NTL growth (yoy)	6.5	-76.4	-18.3	-4.7	4.3	14.7	33.3	132.9
GDP growth (yoy)	2.8	-9.0	-1.5	1.1	2.8	4.8	6.7	12.6
avg. # of obs	29.6	6.8	15.8	25.7	29.2	35.6	42.9	48.0
avg. pixel NTL	0.8	0.02	0.09	0.2	0.3	0.5	1.2	11.8
Total								
NTL growth (yoy)	6.2	-120.8	-24.1	-6.5	4.5	16.4	38.2	158.8
GDP growth (yoy)	2.8	-14.6	-1.6	1.1	2.9	4.9	7.0	16.8
avg. # of obs	24.6	0.5	7.1	13.6	26.1	33.9	41.9	51.5
avg. pixel NTL	0.8	0.008	0.07	0.1	0.2	0.4	0.9	18.8

Notes. This table presents summary statistics of NTL growth and GDP growth between 2013Q2-2025Q2 by quarter. Growth rates are year-on-year and defined as 100 times the first different in log levels. The unit of the average pixel NTL is $nWcm^{-2}sr^{-1}$.

var(z) differs across countries for the SVNL annual dataset. In particular, it should be negatively related to country size, observation frequency, and light intensity of a country.

Figure 1 shows that this is indeed the case. In panels (a)-(c), we contrast the standard deviation of NTL growth against the country characteristics mentioned above at the country level. It is clear that for countries that have a larger geographic area, higher number of observations per year, and higher average brightness, the variability of NTL growth is smaller.

Figure 1: Heteroskedasticity of NTL Growth in the SVNL Dataset

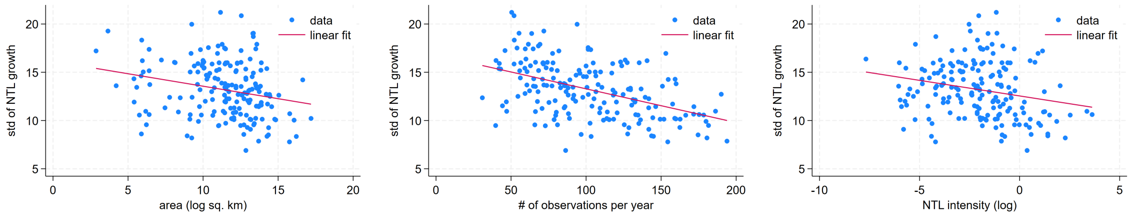

(a) std(z) vs. area (b) std(z) vs. N^{obs} (c) std(z) vs. NTL intensity

Note. This figure contrasts the standard deviation of NTL growth against country characteristics, including the size, the average number of daily observations per year, and the average intensity of NTL in panels (a), (b), and (c), respectively. In each panel, each blue dot represents a country, while the red line is the fit from a linear regression of the standard deviation of NTL growth on the respective characteristic.

Table 3 reports Ordinary Least Squares (OLS) regression estimates examining the relationship between growth volatility and country characteristics across both the SVNL annual and VIIRS quarterly datasets. Column (1) confirms that the standard deviation of NTL growth is negatively and statistically significantly associated with country size, observation frequency, and light intensity. By contrast, column (2) suggests that GDP growth does not exhibit such heteroskedasticity. Columns (3) and (4) show that these patterns largely hold in the VIIRS quarterly sample. While NTL growth volatility remains negatively correlated with country characteristics, there is also some evidence in that quarterly GDP volatility is inversely related to country size and light intensity. The adjusted R^2 indicates that country characteristics account for a sizable share of NTL growth volatility but a negligible share of GDP growth volatility.

3.4.2 Relationships between NTL and GDP growth

We examine the structural and predictive relationships between NTL growth and GDP growth across different country groups based on their income status. To this end, we first regress NTL growth on GDP growth . In the presence of measurement errors, the coefficients we obtain would be a biased estimate of the elasticity. Despite this, they provide initial indication of the lower bound of the elasticity as well as the stability of the relationship between NTL growth and GDP growth. Next we perform the regression in reverse order, which provides a predictive relationship from NTL growth

Table 3: Heteroskedasticity of Measurement Errors in the SVNL and VIIRS Datasets

	SVNL annual		VIIRS quarterly	
	(1)	(2)	(3)	(4)
	std(z)	std(y)	std(z)	std(y)
area (log sq. km)	-0.381***	-0.066	-0.108	-0.231***
	(0.073)	(0.041)	(0.165)	(0.045)
avg. # obs per year	-0.031***	-0.001	-0.155***	0.007
	(0.005)	(0.003)	(0.025)	(0.007)
NTL intensity (log)	-0.472***	-0.060	-1.901***	-0.387***
	(0.095)	(0.054)	(0.300)	(0.082)
constant	19.750***	4.278***	19.650***	5.374***
	(0.887)	(0.503)	(1.783)	(0.489)
Obs	181	181	373	373
Adjusted R^2	0.338	0.000	0.256	0.068

Notes. This table presents regression results of the standard deviation of NTL growth and GDP growth on country characteristics. Standard errors are in parentheses.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

to GDP growth. The results would be useful for predicting GDP growth where official data is not available and constructing an improved measure of GDP growth.

Structural relationship

Table 4 presents the results of regressions of NTL growth on GDP growth for the SVNL annual dataset. Since country specific invariant factors, such as the habit of using lights at night, are already removed when working with growth rates, country fixed effects are not included in these regressions.

Across all columns, the coefficient before GDP growth is statistically significant at the 1 percent level. The coefficients for different country groups are close to each other. When standard errors are taken into account, they are not statistically different. When year fixed effects are included in column (2), the coefficient before GDP growth barely changes compared to that in column (1), suggesting that there is no apparent trend NTL growth beyond GDP growth.

Table 5 presents the results from the same regressions for the VIIRS quarterly

Table 4: SVNL Annual Dataset: Structural Relationship

	NTL growth				
	(1)	(2)	(3)	(4)	(5)
	full sample	full sample	AEs	EMs	LIDCs
GDP growth	0.482***	0.490***	0.457***	0.467***	0.404***
	(0.0478)	(0.0477)	(0.125)	(0.0631)	(0.0914)
constant	3.963***	3.935***	1.987***	4.141***	5.431***
	(0.248)	(0.235)	(0.524)	(0.334)	(0.516)
year fixed effects	-	Yes	-	-	-
Obs	4818	4818	961	2433	1424
Adjusted R^2	0.0204	0.193	0.0127	0.0216	0.0129

Note. This table presents the regression results of NTL growth on GDP growth in the SVNL annual dataset. Mutually exclusive country groups include advanced economies (AEs), emerging markets and middle-income economies (EMs), and low-income developing countries (LIDCs). $^*p < 0.10$, $^{**}p < 0.05$, *** $p < 0.01$.

dataset. Notably, the coefficient before GDP growth is smaller than that in Table 4. In particular, the relationship between NTL growth and GDP growth is negative for LIDCs (column (5)). One possibility for the smaller coefficients is that greater noise in GDP growth measurement errors at the quarterly frequency biases them downward. In Section 4 where we estimate the measurement errors in GDP growth, we show that this is indeed the case.

While Table 4 provides insightful information on the structural relationship between NTL growth and GDP growth, the coefficients cannot be used directly as the elasticity of NTL with respect to GDP, because the measurement errors in both variables would bias the elasticity estimates. In this sense, they are naive estimates of the elasticity. In Section 4, we provide estimates of the elasticity by properly taking measurement errors into account.

Predictive relationship

Table 6 presents the results of regressions of GDP growth on NTL growth for the SVNL annual dataset. As before, country fixed effects are not included while year fixed effects play a limited role changing the coefficient before NTL growth from column (1) to column (2).

Across all columns, the coefficients before NTL growth are statistically significant at the 1 percent level and remains within the 95 percent confidence intervals of each

Table 5: VIIRS Quarterly Dataset: Structural Relationship

	NTL growth				
	(1)	(2)	(3)	(4)	(5)
	full sample	full sample	AEs	EMs	LIDCs
GDP growth	0.194***	0.224***	0.328**	0.208**	-0.394*
	(0.073)	(0.077)	(0.150)	(0.093)	(0.202)
constant	4.914***	4.828***	2.080***	5.726***	9.143***
	(0.340)	(0.335)	(0.581)	(0.441)	(1.168)
year and quarter fixed effects	-	Yes	-	-	-
Obs	3964	3964	1144	2289	531
Adjusted R^2	0.002	0.127	0.003	0.002	0.005

Note. This table presents the regression results of NTL growth on GDP growth in the VIIRS quarterly dataset. Mutually exclusive country groups include advanced economies (AEs), emerging markets and middle-income economies (EMs), and low-income developing countries (LIDCs). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

other. As explained in Section 2.4.1, they can be directly used for predicting GDP growth.

Table 6: SVNL Annual Dataset: Predictive Relationship

	GDP growth				
	(1)	(2)	(3)	(4)	(5)
	full sample	full sample	AEs	EMs	LIDCs
NTL growth	0.0429***	0.0441***	0.0301***	0.0471***	0.0335***
	(0.00425)	(0.00429)	(0.00823)	(0.00637)	(0.00759)
constant	3.118***	3.111***	2.464***	3.023***	3.773***
	(0.0614)	(0.0569)	(0.110)	(0.0905)	(0.117)
year fixed effects	-	Yes	-	-	-
Obs	4818	4818	961	2433	1424
Adjusted R^2	0.0204	0.184	0.0127	0.0216	0.0129

Note. This table presents the regression results of GDP growth on NTL growth in the SVNL annual dataset. Mutually exclusive country groups include advanced economies (AEs), emerging markets and middle-income economies (EMs), and low-income developing countries (LIDCs). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 7 presents the analogous results for the VIIRS quarterly dataset. The coefficient before NTL growth is an order-of-magnitude smaller than that in Table 6 because of large measurement errors in NTL growth at the quarterly frequency. With

Table 7: VIIRS Quarterly Dataset: Predictive Relationship

	GDP growth				
	(1)	(2)	(3)	(4)	(5)
	full sample	full sample	AEs	EMs	LIDCs
NTL growth	0.0092***	0.0095***	0.0128**	0.0105**	-0.0181*
	(0.0035)	(0.0033)	(0.0058)	(0.0047)	(0.0093)
constant	2.814***	2.813***	2.205***	2.812***	4.283***
	(0.061)	(0.055)	(0.095)	(0.084)	(0.187)
year and quarter fixed effects	-	Yes	-	-	-
Obs	3964	3964	1144	2289	531
Adjusted R^2	0.0015	0.221	0.0033	0.0017	0.0052

Note. This table presents the regression results of GDP growth on NTL growth in the VIIRS quarterly dataset. Mutually exclusive country groups include advanced economies (AEs), emerging markets and middle-income economies (EMs), and low-income developing countries (LIDCs). $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

a limited number of observations, the estimated relationship for LIDCs is negative.

4 Estimation and Results

4.1 Baseline

To estimate the elasticity between NTL and GDP, we use the Generalized Method of Moments (GMM). For the baseline, our GMM estimator combines moment conditions as in equations (5), (7) and (9) for the SVNL annual dataset, and replaces equation (9) with equation (10) for the VIIRS dataset. We conduct 400 bootstraps to estimate the standard errors.

Table 8 presents the estimated parameters. The estimated elasticity (β) between NTL and GDP is 1.52 for the SVNL annual dataset with 95 percent confidence interval of $[0.72, 2.31]$. This is much higher than the naive estimate of 0.48 in column (1) of Table 4. The estimated elasticity for the VIIRS quarterly dataset is higher, at 2.25 for the VIIRS quarterly dataset with 95 percent confidence interval of $[1.20, 3.31]$. It is also substantially higher than the naive estimate of 0.194 in column (1) of Table 5.

While the point estimates of the elasticity for the SVNL annual dataset and for the VIIRS quarterly dataset differ, they fall within each other's confidence intervals. In this sense, they are not statistically different. Both estimates are larger than the elasticity estimated using DMSP-OLS data. Henderson, Storeygard, and Weil (2012) estimate an elasticity of 1 and Hu and Yao (2022) estimate an elasticity of 1.3 for the DMSP-OLS data. Taking the confidence intervals into account, however, the difference is statistically small.

Table 8: Baseline Parameter Estimates

	β	σ_{y^*}	σ_{ε^y}	$\log \sigma_z^2$	σ_{η^z}
SVNL annual					
Point Estimate	1.52***	4.19**	3.12*	-9.24***	13.24***
	(0.41)	(1.95)	(1.84)	(3.16)	(0.47)
t-statistic	3.73	2.15	1.69	-2.93	27.89
VIIRS quarterly					
Point Estimate	2.25***	2.90	3.71	-6.01	17.11***
	(0.54)	(1.83)	(3.42)	(16.19)	(1.74)
t-statistic	4.19	1.59	1.09	-0.37	9.84

Note. This table presents estimates of parameters in equations (5), (7) and (9) for the SVNL annual dataset and the VIIRS quarterly dataset. Because σ_z is very large, we estimate and present it in logarithm. Standard errors are in parentheses based on 400 bootstraps. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Along with the elasticity β, our GMM estimator provides estimates of other parameters, including the variance of true GDP growth and its measurement errors. With them, we can calculate the signal to total variance ratio in official GDP growth, defined as the fraction of variance in official GDP growth explained by true economic growth:

$$\text{Sig}_y = \frac{\sigma_{y^*}^2}{\sigma_{y^*}^2 + \sigma_{\varepsilon^y}^2}.$$

Using the point estimates, the signal to total variance ratio equals 64% for the annual data and 38% for the quarterly data. If we allow the measurement errors σ_{ε^y} to differ across different country groups, the implied Sig_y would be 96%, 62%, and 56% for AEs, EMs, and LIDCs, respectively, for annual GDP growth. For quarterly year-on-year GDP growth, Sig_y would be 58%, 39%, and 25% for AEs, EMs, and LIDCs, respectively. Both annual and quarterly results indicate that the accuracy of official GDP growth is highest for AEs, followed by EMs and then LIDCs. However, quarterly GDP growth is much more noisy than annual GDP growth.

As a comparison, Henderson, Storeygard, and Weil (2012) identify the elasticity by assuming that for good data countries Sig_y is 90%, and subsequently calculate that Sig_y is 59% for bad data countries. Our results from the SVNL annual data are close to their assumptions, but our calculation naturally follows the identification of the elasticity through heteroskedasticity of NTL growth without making any assumptions about the measurement errors in GDP growth.

We can also use the estimates of the elasticity and the variances to calculate the optimal weight in a light-adjusted GDP measure, as in equation (15) of Section 2.4. For annual GDP, Using $\gamma = 0.0429$ from column (1) in Table 6, we arrive at an optimal weight of 0.38 for a country with median characteristics in terms of country size, observation frequency, and NTL intensity. For quarterly GDP, the same calculation yields an optimal weight of 0.63. This suggests that NTL data could play a useful role in improving GDP growth estimates.

4.2 Elasticity by country groups

It is sometimes cautioned that NTL is a good proxy of economic activity for middle income countries, but not for low income or advanced economies, because the presence of a large service sector or agriculture sector breaks down the relationship NTL and GDP.

To investigate the heterogeneity of the relationship between NTL growth and GDP growth, we allow the elasticity and the variance of GDP growth measurement errors to vary across country groups, while holding all other parameters constant. The country group classification we consider includes income status and the structure of the economy. Specifically, we categorize countries into AEs and EMDEs, and further partition the sample based on whether countries fall below or above the median values for GDP per capita and the respective sectoral shares of agriculture, industry, and services. Table 9 presents the results.

The elasticity estimate is statistically significant for most columns. Despite differing point estimates, the difference between various country subgroups is modest when standard errors are taken into account. For the SVNL annual dataset, columns (1) and (2) show that for AEs' elasticity is slighly higher than that of EMDEs, but it is imprecisely estimated. Consistent with this, columns (4) and (10) show that for countries whose GDP per capita or service share is above average—which are typical of AEs—the elasticity is not statistically significant. Column (6) shows that even for countries with an above-average agriculture share, NTL growth still reflects GDP growth in a statistically significant way.

For the VIIRS quarterly dataset, the estimated elasticity is generally higher than that from the SVNL annual dataset. Notably at the quarterly frequency, the elasticity for AEs is much higher than that for EMDEs (column (1)). For countries whose agriculture share or service share is above average, the elasticity is not statistically significant.

Table 9: Heterogeneity of Elasticity between NTL and GDP by Country Groups

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Income status		GDP per capita		Agriculture share		Industry share		Services share	
	AEs	EMDEs	Below	Above	Below	Above	Below	Above	Below	Above
SVNL										
β	1.99**	1.57***	1.45***	1.17	1.00***	1.51***	1.64***	1.81***	1.27***	0.70
	(0.91)	(0.37)	(0.37)	(0.73)	(0.38)	(0.35)	(0.39)	(0.35)	(0.34)	(0.62)
t-statistic	2.20	4.22	3.96	1.60	2.66	4.29	4.23	5.11	3.72	1.12
VIIRS										
β	4.55***	2.85***	2.26***	3.23**	2.48**	1.12	5.08***	2.49***	1.71**	-0.54
	(1.62)	(0.63)	(0.75)	(1.27)	(1.17)	(1.55)	(1.13)	(0.71)	(0.76)	(2.03)
t-statistic	2.81	4.53	3.00	2.55	2.12	0.72	4.49	3.50	2.26	-0.26

Note. This table presents the results of the estimates of the elasticity between NTL and GDP by economic structure. The first row represents estimates of the elasticity. Columns (1)-(8) consider grouping countries based on their GDP per capita and sector share relative to the median of the sample. For each variable, below (above) indicates below (above) the median of that variable in the sample. For example, column (1) includes countries whose GDP per capita is below the median. Standard errors are in parentheses based on 400 bootstraps. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Overall, the relationship between NTL and GDP growth remains robust across most economies at both annual and quarterly frequencies. However, this relationship diminishes in high-income nations characterized by a dominant service sector. At the quarterly frequency, it also weakens for economies with a substantial agricultural share.

4.3 Robustness

In this section, we test the model's robustness to data winsorization. As mentioned in Section 3.3, for the baseline sample we have winsorized the data by removing the top and bottom 5 percent of the nighttime light growth distribution as well as the top and bottom 1 percent of the GDP growth distribution. Such winsorization is informed by experiments. It is useful because growth rates at the extreme ends have disproportionate influence on the estimated relationship. Since NTL growth is much more volatile than GDP growth, we chose a higher truncation threshold for nighttime light growth than for GDP growth.

Figures 2 and 3 show the extent to which the removal of the extreme ends of NTL growth and GDP growth distributions affects the estimated relationship for the SVNL annual dataset and the VIIRS quarterly dataset. We define a truncation threshold of

θ as the top and bottom θ% of the distribution. The left panel in each figure controls the threshold of GDP growth and contrasts the elasticity estimate against various NTL growth thresholds, while the right panel in each figure does it in reverse order, controlling the threshold of NTL growth and then contrasting the elasticity estimate against various GDP growth thresholds.

Figure 2: SVNL Dataset: Elasticity Estimates with under Various Winsorization

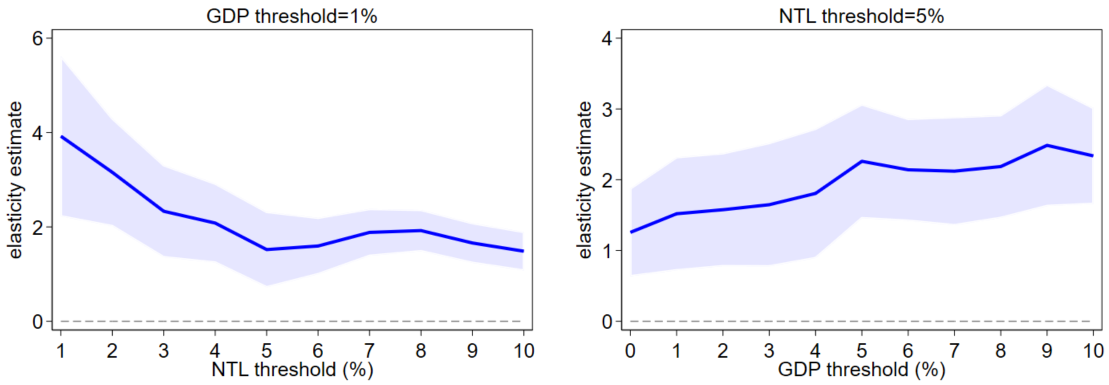

(a) winsorization by NTL (b) winsorization by GDP

Note This figure presents the estimated elasticity between NTL and GDP with various winsorization choices for the SVNL annual dataset. In the left panel, the sample excludes the extreme ends of the NTL growth distribution with progressively higher thresholds, while keeping the threshold for GDP growth fixed. In the right panel, the sample excludes the extreme ends of the GDP growth distribution with progressively higher thresholds, while keeping the threshold for NTL growth fixed. For example, NTL threshold = 5% means that the top and bottom 5 percent of NTL growth distribution is excluded throughout. Blue lines are point estimates and blue bands are 95 percent confidence intervals.

Both figures show that NTL thresholds have a large impact on the elasticity estimate. With no or small winsorization of NTL growth (e.g., NTL threshold = 1%), the elasticity estimate is large and very imprecisely estimated. However, with moderate winsorization of NTL growth, such as removing the top and bottom 5% (i.e., NTL threshold = 5%), the elasticity estimates become much more precise. The left panel of each figure shows that given a GDP winsorization threshold of 1%, the estimated elasticity declines sharply as the NTL threshold increases but stabilizes when the top and bottom 5% of the NTL growth distribution is excluded. The right panel shows that given an NTL winsorization threshold of 5%, the elasticity estimate increases slightly with the threshold of GDP growth. Considering the confidence intervals, however, the estimates do not differ substantially.

In summary, data winsorization impacts the results. The choice of our baseline sample, with a NTL growth threshold of 5 percent and a GDP growth threshold of 1 percent, is supported by the results in Figures 2 and 3. Our estimation results results

Figure 3: VIIRS Dataset: Elasticity Estimates with under Various Winsorization

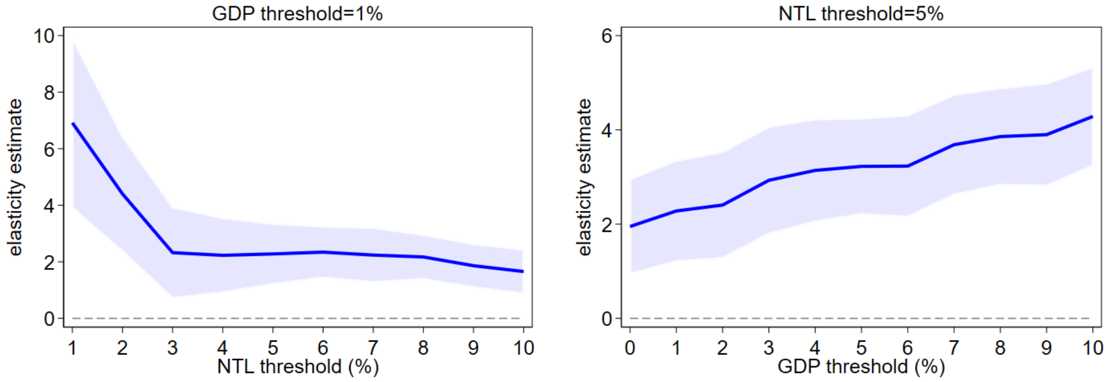

(a) winsorization by NTL (b) winsorization by GDP

Note This figure presents the estimated elasticity between NTL and GDP with various winsorization choices for the VIIRS quarterly dataset. In the left panel, the sample excludes the extreme ends of the NTL growth distribution with progressively higher thresholds, while keeping the threshold for GDP growth fixed. In the right panel, the sample excludes the extreme ends of the GDP growth distribution with progressively higher thresholds, while keeping the threshold for NTL growth fixed. For example, NTL threshold = 2% means that the top and bottom 2 percent of NTL growth distribution is excluded throughout. Blue lines are point estimates and blue bands are 95 percent confidence intervals.

are relatively robust to excluding more observations for NTL growth and GDP growth.

5 Conclusion

The old generation of DMSP-OLS nighttime light data has already contributed greatly to deepening our economic understanding. With the availability of the generation of VIIRS data at an even finer spatial grid and with much improved measurement of low-light areas, nighttime lights are becoming an even more important source of information.

In this paper, we provide a new framework to estimate the elasticity between nighttime lights and GDP, achieving identification through heteroskedasticity of measurement errors in nighttime lights. The flexibility of our framework offers a general estimation strategy to gauge the relationship between satellite data and economic activity.

We have applied the new framework to two datasets: the SVNL annual dataset that synthesized DSMP and VIIRS data and the VIIRS quarterly dataset that is aggregated from the VIIRS monthly series. We estimate the elasticity between nighttime lights and GDP to be 1.52 for the global sample at the annual frequency and 2.25 at the quarterly frequency. The elasticity is relatively stable across country groups with different sectoral compositions, and is robust to data winsorization. Our elasticity estimates can be readily used to transform changes in nighttime lights into changes in economic activity, providing a useful input to applications such as a difference-in-differences style policy analysis.

The implied signal-to-total variance ratio of GDP growth from the SVNL annual dataset is 96% for advanced economies, 62% for emerging markets, and 56% for low income countries. These results suggest substantial measurement errors in the official GDP statistics of emerging markets and low income countries. Furthermore, the VIIRS quarterly dataset yields lower signal-to-total variance ratios across all country classifications compared to the SVNL annual data, underscoring the inherent imprecision of quarterly growth figures. These findings indicate that nighttime lights could serve as a useful tool for refining official GDP statistics, particularly for developing countries and high frequency reporting.

The potential of the VIIRS nighttime light data appears boundless, encompassing questions about the economic costs of natural disasters and conflicts, the local benefits from infrastructure investments, and the spatially diverse effects of economic policies, among many others. We hope that our derived elasticity and subsequent applications will motivate further adoption of VIIRS nighttime light data in economic analyses.

References

Anaxagorou, C., G. Efthyvoulou, and V. Sarantides (2020): "Electoral motives and the subnational allocation of foreign aid in sub-Saharan Africa," *European Economic Review*, 127, 103430.

Arbatli, C. E., and D. Gomtsyan (2021): "Sectarian aid, sanctions and subnational development," *European Economic Review*, 139, 103891.

Chanda, A., and S. Kabiraj (2020): "Shedding Light on Regional Growth and Convergence in India," *World Development*, 133, 104961.

Chen, X., and W. D. Nordhaus (2011): "Using luminosity data as a proxy for economic statistics," *Proceedings of the National Academy of Sciences*, 108(21), 8589–8594.

Chen, X., Z. Wang, F. Zhang, G. Shen, and Q. Chen (2024): "A global annual simulated VIIRS nighttime light dataset from 1992 to 2023," *Scientific Data*, 11(1), 1380.

Chor, D., and B. Li (2021): "Illuminating the Effects of the US-China Tariff War on China's Economy," *National Bureau of Economic Research*.

——— (2024): "Illuminating the effects of the US-China tariff war on China's economy," *Journal of International Economics*, 150, 103926.

Civelli, A., A. Gaduh, and A. S. Yousuf (2024): "Measuring Economic Growth with a Fully Identified Three-Signal Model," *Review of Economics and Statistics*, pp. 1–30.

Elvidge, C. D., K. E. Baugh, M. Zhizhin, and F.-C. Hsu (2013): "Why VIIRS Data are Superior to DMSP for Mapping Nighttime Lights," *Proceedings of the Asia-Pacific Advanced Network*, 35(0), 62.

Gershman, B. (2020): "Witchcraft beliefs as a cultural legacy of the Atlantic slave trade: Evidence from two continents," *European Economic Review*, 122, 103362.

Gibson, J., S. Olivia, and G. Boe-Gibson (2020): "Night Lights in Economics: Sources and Uses," *Journal of Economic Surveys*, 34(5), 955–980.

Henderson, J. V., T. Squires, A. Storeygard, and D. Weil (2018): "The Global Distribution of Economic Activity: Nature, History, and the Role of Trade," *The Quarterly Journal of Economics*, 133(1), 357–406.

Henderson, J. V., A. Storeygard, and D. N. Weil (2012): "Measuring Economic Growth from Outer Space," *American Economic Review*, 102(2), 994–1028.

Hirte, G., C. Lessmann, and A. Seidel (2020): "International trade, geographic heterogeneity and interregional inequality," *European Economic Review*, 127, 103427.

Hu, Y., and J. Yao (2022): "Illuminating economic growth," *Journal of Econometrics*, 228(2), 359–378.

Iacoella, F., B. Martorano, L. Metzger, and M. Sanfilippo (2021): "Chinese official finance and political participation in Africa," *European Economic Review*, 136, 103741.

Jean, N., M. Burke, M. Xie, W. M. Alampay Davis, D. B. Lobell, and S. Ermon (2016): "Combining satellite imagery and machine learning to predict poverty," *Science*, 353(6301), 790–794.

Keola, S., M. Andersson, and O. Hall (2015): "Monitoring Economic Development from Space: Using Nighttime Light and Land Cover Data to Measure Economic Growth," *World Development*, 66, 322–334.

Lessmann, C., and A. Seidel (2017): "Regional inequality, convergence, and its determinants–A view from outer space," *European Economic Review*, 92, 110–132.

Morris, S. D., and J. Zhang (2019): "Validating China's Output Data Using Satellite Observations," *Macroeconomic Dynamics*, 23(8), 3327–3354.

Nordhaus, W., and X. Chen (2015): "A Sharper Image? Estimates of the Precision of Nighttime Lights as a Proxy for Economic Statistics," *Journal of Economic Geography*, 15(1), 217–246.

Pinkovskiy, M., and X. Sala-i Martin (2016): "Newer need not be better: evaluating the Penn World Tables and the World Development Indicators using nighttime lights," Discussion paper, National Bureau of Economic Research.

Rigobon, R. (2003): "Identification through heteroskedasticity," *Review of Economics and Statistics*, 85(4), 777–792.

Rigobon, R., and B. Sack (2004): "The impact of monetary policy on asset prices," *Journal of monetary economics*, 51(8), 1553–1575.

Silungwe, A., A. Baer, and V. Guerreiro (2022): "2020 Global Stocktaking of National Accounts Statistics: Availability for Policy and Surveillance," *IMF Working Papers*, 2022(029).

Yeh, C., A. Perez, A. Driscoll, G. Azzari, Z. Tang, D. Lobell, S. Ermon, and M. Burke (2020): "Using publicly available satellite imagery and deep learning to understand economic well-being in Africa," *Nature communications*, 11(1), 2583.

Appendices

A Heteroskedasticity of NTL Growth

In this section, we show that the variance of measurement errors in NTL growth is heteroskedastic.

A.1 Temporal Averaging at the Pixel Level

Let $\varepsilon^z_{p,d}$ be the idiosyncratic measurement error of nighttime light for pixel p on day d. Assuming they are independent with variance σ^2_z, the annual pixel value measurement error $\varepsilon_{p,t}$ is the mean of N^{obs}_i daily observations:

$$\varepsilon_{p,t} = \frac{1}{N^{obs}_i} \sum_{d=1}^{N^{obs}_i} \varepsilon_{p,d}, \tag{16}$$

which implies that

$$\text{var}(\varepsilon_{p,t}) = \frac{\sigma^2_z}{N^{obs}_i}. \tag{17}$$

A.2 Spatial Aggregation to Country Level

The total measurement error $\varepsilon_{i,t}$ for country i is the sum of all idiosyncratic measurement errors. Let N_i represent the number of pixels in country i. The variance of the total measurement error is:

$$\text{var}(\varepsilon_{i,t}) = \sum_{p=1}^{N_i} \text{var}(\varepsilon_{p,t}) = \frac{N_i \sigma^2_z}{N^{obs}_i}. \tag{18}$$

A.3 Variance of NTL Growth

Let $Z_{i,t}$ and $Z^*_{i,t}$ be the measured and true sum of NLT in country i in year t. Defining NTL growth as the log difference, $z_{i,t} = \log(Z_{i,t}) - \log(Z_{i,t-1})$, we can derive the variance of the measurement error using a first-order Taylor expansion. For small errors $\varepsilon_{i,t}$, we have:

$$\log(Z_{i,t}) = \log(Z^*_{i,t} + \varepsilon_{i,t}) \approx \log(Z^*_{i,t}) + \frac{\varepsilon_{i,t}}{Z^*_{i,t}} \tag{19}$$

The measured growth rate is then:

$$z_{i,t} = z^*_{i,t} + \varepsilon^z_{i,t} \approx z^*_{i,t} + \left(\frac{\varepsilon_{i,t}}{Z^*_{i,t}} - \frac{\varepsilon_{i,t-1}}{Z^*_{i,t-1}} \right) \tag{20}$$

Let $\bar{z}^*_{i,t}$ be the the average NTL intensity per pixel. It follows that $Z^*_{i,t} = \bar{z}^*_{i,t} N_i$. Since $\bar{z}^*_{i,t}$ has much greater variation across countries than over time, we can approximate it by its country average: $\bar{z}^*_{i,t} \approx \bar{z}^*_i \approx \mathrm{E} z_i$.

With the measurement errors assumed to be independent over time and following equation (18), the variance of the NTL growth measurement error $\varepsilon^z_{i,t}$ is:

$$\mathrm{var}(\varepsilon^z_{i,t}) \approx \frac{\mathrm{var}(\varepsilon_{i,t})}{(Z^*_{i,t})^2} + \frac{\mathrm{var}(\varepsilon_{i,t-1})}{(Z^*_{i,t-1})^2} = \frac{\mathrm{var}(\varepsilon_{i,t})}{(\bar{z}^*_{i,t} N_i)^2} + \frac{\mathrm{var}(\varepsilon_{i,t-1})}{(\bar{z}^*_{i,t-1} N_i)^2} \tag{21}$$

$$\approx \frac{2\,\mathrm{var}(\varepsilon_{i,t})}{(\mathrm{E} z_i N_i)^2} = \frac{2\sigma_z^2}{(\mathrm{E} z_i)^2 N_i N_i^{obs}}. \tag{22}$$

With a slight abuse of notation, replacing N_i with Area_i yields the heteroskedastic form:

$$\mathrm{var}(\varepsilon^z_{i,t}) \approx \frac{2\sigma_z^2}{\mathrm{Area}_i \cdot N_i^{obs} \cdot (\mathrm{E} z_i)^2}. \tag{23}$$

B Derivation of Light-Adjusted GDP Measure

Let $\tilde{y}$ be a linear prediction of GDP growth by nighttime light growth, $\tilde{y}_{i,t} = \gamma z_{i,t}$. Some applications consider a linear combination of y and $\tilde{y}$ as a new measure of GDP growth,[5]

$$\bar{y} = (1-\lambda) y + \lambda \tilde{y}, \tag{24}$$

where λ is the weight on nighttime lighted-predicted GDP growth.

The expected mean squared prediction error of the new measure follows:

$$\begin{aligned} E(\bar{y} - y^*)^2 &= E\left[(1-\lambda)(y - y^*) + \lambda(\tilde{y} - y^*)\right]^2 && (25)\\ &= (1-\lambda)^2 \mathrm{var}(\varepsilon^y) + \lambda^2 \mathrm{var}(\tilde{y} - y^*) && (26)\\ &= (1-\lambda)^2 \mathrm{var}(\varepsilon^y) + \lambda^2 \mathrm{var}(\gamma(\beta y^* + \mu^z) - y^*) && (27)\\ &= (1-\lambda)^2 \mathrm{var}(\varepsilon^y) + \lambda^2 \left[(\gamma\beta - 1)^2 \mathrm{var}(y^*) + \gamma^2 \mathrm{var}(\mu^z)\right]. \end{aligned}$$

The optimal linear combination that minimizes the expected mean squared prediction error therefore has the following weight:

$$\lambda^* = \frac{\mathrm{var}(\varepsilon^y)}{\mathrm{var}(\varepsilon^y) + (\gamma\beta - 1)^2 \mathrm{var}(y^*) + \gamma^2 \mathrm{var}(\mu^z)} \tag{28}$$

Empirically, since we have identified β already, the remaining terms can be computed as follows:

[5] For ease of presentation and without loss of clarity, we have omitted the subscripts.

$$\operatorname{var}(y^*) = \operatorname{cov}(y, z)/\beta \tag{29}$$

$$\operatorname{var}(\epsilon^y) = \operatorname{var}(y) - \operatorname{var}(y^*) \tag{30}$$

$$\operatorname{var}(\epsilon^z) = \operatorname{var}(z) - \beta^2 \operatorname{var}(y^*). \tag{31}$$