Journal of Econometrics 228 (2022) 359–378

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Illuminating economic growth☆

Yingyao Hu [a,*], Jiaxiong Yao [b]

[a] Johns Hopkins University, Department of Economics, Wyman Park Building 544E, 3400 N. Charles Street, Baltimore, MD 21218, United States of America

[b] International Monetary Fund, 700 19th St NW, Washington, DC 20431, United States of America

ARTICLE INFO

Article history:
Received 27 January 2020
Received in revised form 19 May 2021
Accepted 19 May 2021
Available online 7 July 2021

JEL classification:
C10
C14
E01
R12

Keywords:
Measurement error
Nighttime lights
GDP growth

ABSTRACT

This paper seeks to illuminate national accounts GDP growth using satellite-recorded nighttime lights in a measurement error model framework. Using recently developed results in conjunction with reasonable assumptions about the exogeneity of the lights data generating process, we identify and estimate the relationship between nighttime light growth and GDP growth, as well as the nonparametric distribution of errors in both measures. We obtain three key results: (i) the elasticity of nighttime lights to GDP is about 1.3; (ii) national accounts GDP growth measures are less precise for low and middle income countries, and nighttime lights can play a big role in improving such measures; and (iii) our new measure of GDP growth, based on the optimal combination of nighttime lights and national accounts data under our identification assumptions, implies that China and India had considerably lower growth rates than official data suggested between 1993 and 2013. We expect our statistical framework and methodology to have a broad impact on measuring GDP using additional information.

1. Introduction

Economic growth is one of the most important goals of macroeconomic policy-making, but measuring it is not easy. National accounts GDP growth, as the official measure of economic growth, can be quite uncertain. Inability of statistical agencies to produce quality statistics, mismeasurement of the economy, and the existence of informal economy, among others, can all subject GDP growth measures to substantial uncertainty. This problem becomes more acute for low and middle income countries where the data collection and compilation process is less sophisticated. Understanding the uncertainty of these measures and constructing more accurate measures are therefore of great importance to assess economic performance, facilitate cross-country comparisons, and inform policy decisions.

This paper attempts to use satellite-recorded nighttime lights, or night lights, to illuminate national accounts GDP growth.[1] Mostly generated by human activity, night lights are visible from outer space and recorded by satellites. Their global coverage and strong correlation with real economic activity make them attractive as an alternative measure of

☆ The authors are grateful to Robert Barbera, Olivier Blanchard, Xiaodong Zhu, and seminar participants at the International Monetary Fund, Michigan State University, University of Colorado Boulder, Tsinghua University, SEA 2018 annual meeting, Beijing University of Aeronautics and Astronautics, and the 6th International Conference on NSE in Peking University. We also thank the editor, the associate editor, and the anonymous referees for their valuable suggestions. All errors are our own. The views expressed here are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

* Corresponding author.
E-mail addresses: yhu@jhu.edu (Y. Hu), jyao@imf.org (J. Yao).

[1] We use national accounts GDP growth and official GDP growth interchangeably in this paper.

https://doi.org/10.1016/j.jeconom.2021.05.007
0304-4076/

GDP. However, despite their economic relevance and increasing use in the literature, the extent to which night lights can be a useful proxy for GDP is unclear.

Both GDP and night lights contain measurement errors that vary across countries and over time, and the relationship between night light growth and true GDP growth is unknown. For countries with low levels of electrification – most of which are in Africa – night lights reflective of economic activity may be eclipsed by background noise. However, if official measures of GDP growth in those countries are poor, night lights can still serve as a useful proxy for economic activity in spite of the noise. Similarly, for countries with reasonably good GDP growth statistics, the extent to which night lights add additional information on economic activity is unclear. To determine the usefulness of night lights, we need a framework that uncovers the distribution of measurement errors in both national accounts GDP growth and night light growth as well as their functional relationship simultaneously.

In this paper, we provide such a statistical framework based on nonclassical and nonlinear measurement error models. In this framework, the error in national accounts GDP growth may depend on a country's statistical capacity – which is a country's ability to collect, analyze, and disseminate data – and the error in night light growth may depend on geographic locations. The relationship between night light growth and true GDP growth can be nonlinear and vary with geographic locations. Based on recently developed results for measurement error models and using variation across geographic locations and different levels of statistical capacity, we establish identification of the distribution function of night light growth conditional on true GDP growth under fairly weak and reasonable statistical assumptions.

Given our nonparametric identification results, the elasticity of night lights to true GDP is naturally obtained. The elasticity would be useful in imputing changes of economic activity in regions where official data are not available or reliable. With the estimated distributions of measurement errors, we assess the relative uncertainty of night light growth compared to national accounts GDP growth, which informs us about the usefulness of night lights in improving GDP growth measures. We then construct a new GDP growth measure by optimally combining national accounts GDP growth and night light-predicted GDP growth at the observation level. For each observation of official GDP growth, we provide an estimate of the optimal weight on night lights.

To our best knowledge, this is the first paper to estimate the distribution of measurement errors in national accounts GDP growth and night light growth directly from data. The error distributions are crucial for both understanding the uncertainty in national accounts GDP growth and constructing more precise measures. Intuitively, we can use night light growth to infer the accuracy of national accounts GDP growth. This is because night lights reflect real economic activity, and therefore, are correlated with true GDP. Meanwhile, night lights are independent of measurement errors in official GDP. After identifying the functional relationship between night light growth and true GDP growth, we could infer the distribution of true GDP growth from the observed joint distribution of night light growth and national accounts GDP growth. Comparing the distribution of true GDP growth with the observed distribution of national accounts GDP growth, we could then pin down the signal-to-noise ratio of the latter. Therefore and more generally, with additional information such as night lights, we can provide better measures of true GDP growth.

There are three main findings in this paper. First, the relationship between night light growth and GDP growth can be captured by a linear production function and year fixed effects. In our baseline specification, we estimate that one percentage point increase in true GDP growth increases night light growth by about 1.3 percentage points. In other words, the elasticity of night lights to GDP is about 1.3.

Second, we find, perhaps not surprisingly, that measurement errors in national accounts GDP growth are larger for countries at the low end of the income spectrum, and night lights are useful for improving GDP growth measures in low and middle income countries. While the distribution of measurement errors of GDP growth for high income countries is concentrated at zero, that for low and middle income countries exhibits fatter tails. In other words, there is greater uncertainty in low and middle income countries' national accounts GDP growth measures and the measurement errors therein can be substantial at times. Our new measure of GDP growth, which is an optimal linear combination of national accounts GDP growth and light-predicted GDP growth, has a non-trivial weight on the latter. The optimal weight ranges from 0.14 to 0.88 with a mean of 0.65.

Finally, comparing our new measure with national accounts GDP growth, we find that for China and India, official GDP growth rates were substantially higher than true GDP growth. We estimate, for example, that China's true GDP growth could be 3.6 percentage points below official GDP growth on average between 1993 and 2013. For conflict-torn African countries, we find that true GDP growth tends to be smoother than official GDP growth. We speculate that this is likely because of the existence of the informal sector that transitions into the formal sector in good times and out of it in adverse times.

The rest of the paper is organized as follows. Section 2 briefly reviews the related literature. Sections 3 and 4 describe the data and statistical framework, respectively. We present our results in Section 5. Section 6 concludes. Additional information on data is in Appendices A and B. Robustness checks, including simulations and different estimates, are discussed in Appendix C. Mathematical proofs of identification and asymptotic properties are provided in online Appendix D.

2. Related literature

This paper is closely related to several strands of literature.

First, a wealth of literature has shown that night lights are correlated with economic activity at various geographic aggregation levels. For example, Elvidge et al. (1997) focus on city levels, Ghosh et al. (2010) on subnational levels, and Chen and Nordhaus (2011) on national levels. More recently, Gibson et al. (2021) compare night lights with official GDP at the second subnational level. However, the literature also finds that night lights are not a suitable proxy for GDP outside of cities. In this paper, we focus on the econometric relationship between night light growth and GDP growth at the country level.

Second, we contribute to the growing literature on understanding economic growth through the lens of satellite-recorded night lights. Since the seminal work of Chen and Nordhaus (2011), Henderson et al. (2012), night lights have been increasingly used as a proxy for economic activity. Pinkovskiy and Sala-i-Martin (2016) assess the relative quality of GDP per capita and survey means by comparing them to night lights. From a statistical perspective, while those papers construct new measures of GDP – whether in growth, per unit area, or per capita terms – through the combination of night lights with official GDP measures using a constant weight, we show that the information content on GDP from night lights differs for each observation. For each country at each point in time, our optimal linear measure uses a different weight on light-predicted GDP growth, and this can only be achieved when we uncover the entire distribution of measurement errors in both night light growth and official GDP growth.

Third, this paper is related to the measurement error literature on identification and estimation of measurement error models. Our statistical framework is based on recently developed results for nonclassical measurement error models. Since Hu and Schennach (2008), we have been able to generally identify and estimate nonlinear models with nonclassical measurement errors in a continuous variable. When there are only two continuous measurements for a continuous latent variable as in the current paper, nonparametric identification requires additional data information or extra restrictions. Carroll et al. (2010) use a secondary survey sample to achieve nonparametric identification, which can be interpreted as identification with two continuous measurements and two discrete instruments. Our method relies on these two papers with national accounts GDP growth and night light growth as two continuous measurements and statistical capacity and geographic location as two discrete instruments.[2]

Fourth, we contribute to the literature on improving the measurement of the economy from a measurement-error perspective. Aruoba et al. (2016) improve historical United States' GDP growth at relatively high frequency and find the persistence of aggregate output dynamics to be stronger than previously thought. Feng and Hu (2013) show that the official US unemployment rate substantially underestimates the true level of unemployment. This paper aims to improve GDP growth estimates in a measurement error model setting for low and middle income countries.

Finally, we make a contribution to the burgeoning literature on bringing satellite data to economic analysis. Donaldson and Storeygard (2016) provide a comprehensive review of applications of satellite data in economics. While many applications focus on converting satellite images to physical quantities relevant for economics, such as night lights, greenness, or temperature, we focus on examining the relationship between such quantities and economic variables of interest from an econometric perspective. Our method can be applied broadly to a wide range of remote sensing data, as they inevitably contain measurement errors and their relationship with economic variables of interest may not be simple and linear.

3. Data and stylized facts

3.1. Nighttime lights

The Colorado School of Mines hosts average-radiance cloud-free composites of nighttime lights data, made available by Elvidge et al. (1997, 2017). At the time of writing, there are two types of nighttime lights data available that cover different time periods: (1) annual composites between 1992 and 2013 produced by the Operational Linescan Sensor (OLSS) onboard the Defense Meteorological Satellite Program (DMSP), and (2) monthly composites from April 2012 onward produced by the Visible and Infrared Imaging Suite Day Night Band on board of the Joint Polar-orbiting Satellite System satellites. For the DMSP/OLSS[3] annual composites, observations affected by sunlight, moonlight, glare, aurora, and other temporal lights have been removed. However, monthly composites have not been filtered to screen out those lights. For the purpose of this paper, we focus on DMSP/OLSS data.

To get an intuitive sense of night lights data, Fig. 1 presents some examples of satellite images of night lights for Africa and mainland China. Notably, Africa's night lights remained subdued throughout the period of 1992–2013 despite strong economic growth and improvement in electrification. By contrast, China's rapid economic growth during the same period was accompanied by visible changes in luminosity.

Night light data are a type of geospatial data with geographic information associated with each data point. Every pixel of a night light satellite image represents a data point for a certain location at a particular point in time. The value of the pixel indicates the intensity of lights at night, which was originally measured in radiance and then transformed into a discrete digital number between 0 and 63.

2 See Hu (2021) for a survey of recent developments in this literature.

3 DMSP/OLSS is commonly abbreviated as DMSP/OLS. Here, to avoid confusion with Ordinary Least Squares (OLS), we use the abbreviation DMSP/OLSS.

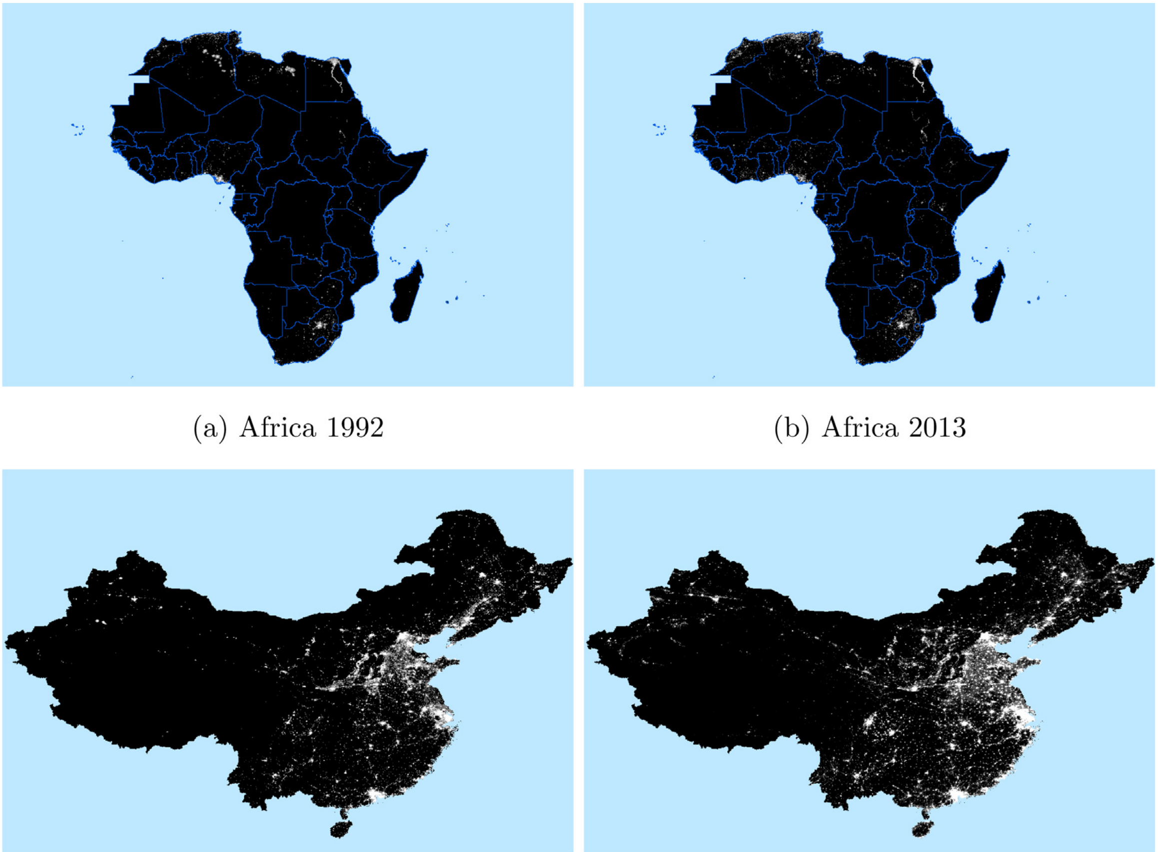

(a) Africa 1992 (b) Africa 2013

(c) China (mainland) 1992 (d) China (mainland) 2013

Fig. 1. Examples of satellite images of night lights.

One of the limitations of DMSP/OLSS data is that they are top-coded, with the maximum pixel value corresponding to the saturation radiance of the satellite sensors. Despite such top coding, variation in night lights still contains ample information on economic activity at the country level.[4] Appendix A provides more details on DMSP/OLSS data and on statistical distributions of night light data.

The Database of Global Administrative Areas (known as GADM) provides administrative shape files for all countries in the world. For each country, we clip night lights at each point in time to their country borders and sum the numeric pixel values of all pixels within their borders. We further calculate night light growth, a proxy for GDP growth that is used extensively in this paper, as the first difference of the logarithm of the sum of night lights between adjacent years.

3.2. GDP, population, statistical capacity, and cloud cover

We obtain GDP growth from the World Development Indicators Database Archive (WDI-DA)[5] of the World Bank, and winsorize the data by discarding the top and bottom 0.5% of GDP growth.[6] For low and middle income countries, the World Bank also provides statistical capacity ratings, which measure a nation's ability to collect, analyze, and disseminate high-quality data about its population and economy. While the ratings change over time, the change is small for most countries. For this reason we group low and middle countries into two categories: those below the median of statistical capacity ratings (low statistical capacity) and those above it (high statistical capacity) . For completeness, we add an

[4] As shown in Appendix A, while top coding might be a concern for high income city states, such as Singapore, where almost all night lights reach saturation radiance, saturated pixels only account for a very small fraction of the pixels that have positive values for most countries. For example, they account for less than 2.4% in the United States, 1.1% in China, and 0% in Sierra Leone.

[5] We use the October 2019 version.

[6] The top 0.5% of GDP growth is typically in high double-digits and mainly consists of country-year observations where there was discovery of natural resources, such as Equatorial Guinea in 1996 and 1997, or recovery from war, such as Iraq in 2004 and Libya in 2012. The bottom of 0.5% of GDP growth is typically in negative double-digits and mainly consists of country-year observations where there was conflict or war.

Table 1
Summary statistics (DMSP/OLSS).

Statistical capacity	Cloud cover	Night lights per 1000 people	GDP growth	Night light growth	# of countries	# of obs
Low	Low	24	0.043	0.055	35	681
Low	High	17	0.042	0.056	26	528
High	Low	35	0.045	0.050	35	726
High	High	51	0.043	0.048	31	636
(High income)	Low	122	0.034	0.030	12	241
(High income)	High	215	0.023	0.021	23	483
Total	–	66	0.040	0.046	162	3295

additional category of high income countries. Such discretization has the additional benefit of reducing the measurement errors in statistical capacity itself.

Measurement errors in night lights are correlated with locations, as atmospheric conditions in different parts of the world affect satellite sensors' ability to capture night lights. We consider a country's annual average cloud cover[7] to account for location-specific measurement errors in night lights, because the more cloudy days there are in a year, the less number of effective daily observations there is to produce night lights at the annual frequency.[8] We further discretize cloud cover of countries into binary values based on whether it is below or above the median. In other words, we group countries into two categories of low and high cloud cover.

3.3. Summary statistics

In total we have an unbalanced panel of 162 countries and 3295 observations based on DMSP/OLSS night light data. The majority of countries have data spanning from 1993 to 2013. Tables 1 presents the summary statistics. As can be seen from the table, richer countries tend to be brighter at night while poorer countries tend to have high GDP growth and night light growth.

4. Statistical framework

In this section, we present a statistical framework to analyze night light data in relation to GDP. Under reasonable assumptions, the relationship between night lights and GDP and the distributions of their respective measurement errors are identified. The elasticity of night lights with respect to true GDP is naturally obtained. We then construct a more accurate measure of GDP growth by optimally combining national accounts data and prediction by night lights.

4.1. Baseline setup

Let $Y^*_{i,t}$ denote true GDP growth for country i in year t. It is measured as $Y_{i,t}$ with an error. Let $Z_{i,t}$ denote night light growth. It is related to true GDP growth but also contains a measurement error. Let S_i and L_i stand for the statistical capacity and the location-specific characteristics of country i, respectively. Both S_i and L_i are binned into groups and thus discrete. Our results are based on a panel data $\left\{Z_{i,t}, Y_{i,t}, S_i, L_i\right\}_{i=1,\dots,N;t=1,\dots,T}$ with a large N and a finite T. To identify the model, we require an i.i.d sample across countries.[9] In other words, the model is identified for each time period t. For the simplicity of estimation, we assume the distributions of two measurement errors are time-invariant, but keep the data dependence across time flexible.

Our baseline specification assumes that official GDP growth contains an additive measurement error that is independent of true GDP growth. The measurement error's distribution varies with statistical capacity S_i as follows:

$$Y_{i,t} = Y^*_{i,t} + \epsilon^Y_{i,t}(S_i). \tag{1}$$

For a subsample of countries with $S_i = s$, the marginal distribution of measurement error $\epsilon^Y_{i,t}(s)$ is denoted as $f_{\epsilon^Y|S}(\cdot|s)$. We assume the error distribution is the same for all t in the estimation, which is not necessary for identification. The past literature has assumed measurement errors to be either in output levels (Chen and Nordhaus, 2011) or in growth rates (Henderson et al., 2012). If the true measurement error is in levels, the structure of measurement error in growth rates might be complex, including possibly serial correlation. On the other hand, if GDP levels are consistently

7 Cloud cover is computed as the area-weighted average between 1901 and 2018 for a country, using high-resolution gridded data sets from the Climate Research Unit, University of East Anglia (CRU TS v.4.03).

8 Since DMSP/OLSS satellites are polar orbiting—their orbits are perpendicular to the direction of Earth's rotation, we also explore the possibility of using a country's latitude to account for measurement errors in night lights in Appendix B.

9 This requirement rules out cross-country dependence, which is a challenging issue to deal with in our nonparametric analysis. Nevertheless, we partially address this issue in our reduced-form analysis in Section 4.3 by examining standard errors clustered by year.

mis-measured over time, such as in the case when an economy has a sizable informal sector, assuming measurement errors in growth rates is more appropriate.[10]

We assume night light growth is related to the true latent GDP growth through an unknown production function $m(\cdot)$ and an additive error term whose distribution may differ by geographic location L_i,

$$Z_{i,t} = m(Y^*_{i,t}) + \epsilon^Z_{i,t}(L_i). \tag{2}$$

For a subsample of countries with $L_i = l$, the marginal distribution of measurement error $\epsilon^Z_{i,t}(l)$ is denoted as $f_{\epsilon^Z|L}(\cdot|l)$. For simplicity, we assume this error distribution is the same for all t in the estimation, which again is not necessary for identification. The additive nature of measurement errors is not necessary for identification, but is assumed for the simplicity of estimation.

A heuristic example

A special case of Eqs. (1) and (2) is when the night light production function is linear and the measurement error distributions are the same for all countries:

$$\begin{aligned} Y_{i,t} &= Y^*_{i,t} + \epsilon^Y_{i,t}, \\ Z_{i,t} &= \beta Y^*_{i,t} + \epsilon^Z_{i,t}. \end{aligned} \tag{3}$$

Under the assumption that measurement errors are classical, β cannot be identified from second moments. To see this, we take variances and covariances of the above equations,

$$\begin{aligned} \text{var}(Y_{i,t}) &= \text{var}(Y^*_{i,t}) + \text{var}(\epsilon^Y_{i,t}), \\ \text{var}(Z_{i,t}) &= \beta^2 \text{var}(Y^*_{i,t}) + \text{var}(\epsilon^Z_{i,t}), \\ \text{cov}(Y_{i,t}, Z_{i,t}) &= \beta \text{var}(Y^*_{i,t}). \end{aligned}$$

Because there are only three equations but four unknowns ($\text{var}(\epsilon^Y_{i,t})$, $\text{var}(\epsilon^Z_{i,t})$, $\text{var}(Y^*_{i,t})$, and β), the system cannot be identified. The literature has dealt with identification in this case by either assuming additional relationship between $\text{var}(\epsilon^Y_{i,t})$ and $\text{var}(\epsilon^Z_{i,t})$ (Henderson et al., 2012) or adding auxiliary data (Pinkovskiy and Sala-i-Martin, 2016).

In our baseline specification (Eqs. (1) and (2)), we relax the assumptions in the linear special case (3) substantially to allow the error terms to vary with a country's statistical capacity and location, which is more realistic.

4.2. Nonparametric identification

We obtain nonparametric identification in a more general setting – where the error terms need not be additive – than our baseline specification. Our identification relies on an i.i.d sample across countries. In other words, we can identify the model using only one period of the data. As such we omit the time script for ease of presentation. To make things clear below, we follow the convention that variables are in uppercase and their realizations in lowercase.

Let $f(\cdot|\cdot)$ be a generic conditional probability density function. Throughout, we use the notation $f_{X|Z}(x|z)$ to stand for a conditional probability density of random variables X and Z in upper case with their corresponding realizations x and z in lower case.

We make the following assumptions.

Assumption 1. Night light growth Z satisfies

$$f_{Z|Y,Y^*,S,L}(z|y, y^*, s, l) = f_{Z|Y^*,L}(z|y^*, l). \tag{4}$$

This assumption implies that night light growth is related to a country's true GDP growth and geographic location, but has nothing to do with how GDP growth is measured or the country's statistical capacity.

Assumption 2. Official GDP growth Y satisfies

$$f_{Y|Y^*,S,L}(y|y^*, s, l) = f_{Y|Y^*,S}(y|y^*, s). \tag{5}$$

This assumption implies that statistical capacity captures how accurate GDP growth is measured regardless of the location of the country.[11]

[10] Relatedly, the literature on survey data has discussed the properties of measurement error in panel data and the consequences of different transformations of the data. See, for example, Griliches and Hausman (1986) and Bound et al. (2001). The literature on household earnings and consumption typically assumes measurement error in levels. However, there is also evidence that measurement error in earnings is not classical and that first-differenced data are more reliable in the presence of positive serial correlation and mean reversion (Bound and Krueger, 1991).

[11] Statistical capacity is likely to be mis-measured. If the measurement error in statistical capacity is independent of other variables, exclusion restrictions on statistical capacity may still hold. But if a conditional independence assumption is conditional on statistical capacity, the measurement error in statistical capacity may fail the conditional independence condition. Since statistical capacity for a country does not change much over our sample period, discretization could alleviate the concerns of mismeasurement. This is the reason that we discretize statistical capacity, have to assume the discrete version is accurate, and use it directly in the assumption. The same argument also applies to the location-specific variable.

In addition to Assumptions 1 and 2, we need some technical assumptions for identification. Let (Z_i, Y_i, S_i, L_i) be a random draw from distribution $f_{Z,Y,S,L}(z, y, s, l)$ for country i in a given year, with the supports $z \in \mathcal{Z} \subseteq \mathbb{R}$, $y \in \mathcal{Y} \subseteq \mathbb{R}$, $y^* \in \mathcal{Y}^* \subseteq \mathbb{R}$, $s \in \mathcal{S} = \{s^1, s^2, \ldots, s^J\}$ with $J \geq 2$, and $l \in L = \{l^1, l^2, \ldots, l^K\}$ with $K \geq 2$, respectively. Note that we have implicitly assumed away cross-sectional dependence. While cross-sectional dependence may well be present in the data, Section 4.3 provides some evidence that it is not a major source of concern at the country level.

We provide sufficient conditions for the identification of latent distributions: $f_{Z|Y^*,L}(z|y^*, l)$, $f_{Y|Y^*,S}(y|y^*, s)$, and $f_{Y^*|S,L}(y^*|s, l)$ from the observed distribution $f_{Z,Y,S,L}(z, y, s, l)$.

We assume

Assumption 3. $f_{Z,Y,Y^*,S,L}(z, y, y^*, s, l)$ is positive, bounded on its support $\mathcal{Z} \times \mathcal{Y} \times \mathcal{Y}^* \times \mathcal{S} \times \mathcal{L}$, and is continuous in $(z, y, y^*) \in \mathcal{Z} \times \mathcal{Y} \times \mathcal{Y}^*$.

Assumption 4. (i) For each given $s \in S$, $\int f_{Y|Y^*,S}(y|y^*, s)h(y^*)dy^* = 0$ for all $y \in \mathcal{Y}$ for a bounded function h implies that $h \equiv 0$ almost everywhere over $\mathcal{Y}^*$;
(ii) For each given (s, l), $\int f_{Z,Y|S,L}(z, y|s, l)h(y)dy = 0$, i.e., $E[h(Y)|Z = z, S = s, L = l] = 0$, for all $z \in \mathcal{Z}$ for a bounded function h implies that $h \equiv 0$ almost everywhere over $\mathcal{Y}$.

Assumption 4(i) is the *bounded completeness* of the conditional density $f_{Y^*|Y,S}(\cdot|\cdot, s)$ (See, for example, Mattner, 1993). As shown in Canay et al. (2013) and Freyberger (2017), these completeness assumptions are not testable without any additional restrictions. Assumption 4 is a high level condition mainly because we are achieving a nonparametric identification result. Assumption 4(i) requires that the distribution of official GDP growth should vary enough for different values of true GDP growth, for each given level of statistical capacity, and Assumption 4(ii) requires that the distribution of night light growth should vary enough for different values of GDP growth, for each given level of statistical capacity and cloud cover.

Define

$$\kappa^{abjk}\left(y^*\right) \equiv \frac{f_{Y^*|S,L}\left(y^*|s^a, l^j\right) f_{Y^*|S,L}\left(y^*|s^b, l^k\right)}{f_{Y^*|S,L}\left(y^*|s^b, l^j\right) f_{Y^*|S,L}\left(y^*|s^a, l^k\right)} \quad \text{for } y^* \in \mathcal{Y}^*. \tag{6}$$

Assumption 5. For any $y_1^* \neq y_2^*$ and $s^a \in \mathcal{S}$, there exist $s^b \in \mathcal{S}$ and $l^j, l^k \in \mathcal{L}$, such that $\kappa^{abjk}\left(y_1^*\right) \neq \kappa^{abjk}\left(y_2^*\right)$ and $\sup_{y^* \in \mathcal{Y}^*} \kappa^{abjk}\left(y^*\right) < \infty$.

This assumption requires that the distribution of true GDP growth varies with countries' statistical capacity and geographic location, which is quite reasonable. In the data, we observe that higher income countries have more effective statistical institutions. In the meantime, it is well known that GDP growth is highly correlated with countries' geographic locations.

Since y^* is not observed, we need a normalization assumption as follows:

Assumption 6. One of the followings holds for all $Y^* \in \mathcal{Y}^*$: for some $s \in \mathcal{S}$, (i) (mean) $E[Y|Y^* = y^*, S = s] = y^*$; or (ii) (mode) $\arg\max_y f_{Y|Y^*,S}\left(y|y^*, s\right) = y^*$; or (iii) (median) $\inf\{v : \int_{-\infty}^{v} f_{Y|Y^*,S}\left(y|y^*, s\right) dy \geq 0.5\} = y^*$.

Assumption 6 says that the reported GDP growth from some country with statistical capacity $S = s$ is targeted for the true y^*. Specifically, either the mean, mode or median of the distribution of Y given $Y^* = y^*$ and $S = s$ is equal to y^*. This condition is not required for other countries with different levels of statistical capacity.

With these assumptions, we have the following nonparametric identification result:

Theorem 1. *Suppose Assumptions 1–6 hold. Then, the distribution function $f_{Z,Y,S,L}(z, y, s, l)$ uniquely determines the joint distribution function $f_{Z,Y,Y^*,S,L}(z, y, y^*, s, l)$ satisfying*

$$f_{Z,Y,Y^*,S,L}(z, y, y^*, s, l) = f_{Z|Y^*,L}(z|y^*, l)f_{Y|Y^*,S}(y|y^*, s)f_{Y^*,S,L}(y^*, s, l). \tag{7}$$

A brief proof (*See Online Appendix D.1 for Details*)**.** For each value (s, l), Assumptions 1 and 2 imply that

$$f_{Y,Z|S,L}(y, z|s^a, l^j) = \int f_{Y|Y^*,S}\left(y|y^*, s^a\right) f_{Z|Y^*,L}(z|y^*, l^j)f_{Y^*|S,L}(y^*|s^a, l^j)dy^*, \tag{8}$$

By Eq. (8) and the definition of the operators in the Appendix, we have the operator equivalence

$$M_{Y,Z|s^a,l^j} = M_{Y|Y^*,s^a} D_{Y^*|s^a,l^j} M_{Z|Y^*,l^j}. \tag{9}$$

Assumption 4 implies that all the operators involved in Eq. (9) are invertible. Hence

$$M_{Y,Z|s^a,l^j} M^{-1}_{Y,Z|s^b,l^j} = M_{Y|Y^*,s^a} D_{Y^*|s^a,l^j} D^{-1}_{Y^*|s^b,l^j} M^{-1}_{Y|Y^*,s^b}. \tag{10}$$

This equation holds for all l^j and l^k so that we may then eliminate $M_{Y|Y^*,s^b}$ to have

$$M_{Y,Y}^{abjk} \equiv (M_{Y,Z|s^a,l^j} M_{Y,Z|s^b,l^j}^{-1})(M_{Y,Z|s^a,l^k} M_{Y,Z|s^b,l^k}^{-1})^{-1} = M_{Y|Y^*,s^a} D_{Y^*}^{abjk} M_{Y|Y^*,s^a}^{-1}, \tag{11}$$

where $D_{Y^*}^{abjk} : \mathcal{L}^2(\mathcal{Y}^*) \to \mathcal{L}^2(\mathcal{Y}^*)$ is still a diagonal operator

$$D_{Y^*}^{abjk} \equiv D_{Y^*|s^a,l^j} D_{Y^*|s^b,l^j}^{-1} (D_{Y^*|s^a,l^k} D_{Y^*|s^b,l^k}^{-1})^{-1}. \tag{12}$$

In fact, this diagonal operator can be defined as $(D_{Y^*}^{abjk} h)(y^*) \equiv \kappa^{abjk}(y^*)h(y^*)$ with κ^{abjk} defined in Eq. (6). Eq. (11) implies a diagonalization of an observed operator $M_{Y,Y}^{abjk}$, where an eigenvalue of $M_{Y,Y}^{abjk}$ equals $\kappa^{abjk}(y^*)$ for a value of y^* with corresponding eigenfunction $f_{Y|Y^*,S}(\cdot|y^*, s^a)$. Notice that each eigenfunction is a conditional density, and therefore, is automatically normalized.

Eq. (11) implies that the operator $M_{Y,Y}^{abjk}$ has the same spectrum as the diagonal operator $D_{Y^*}^{abjk}$. Since an operator is bounded by the largest element of its spectrum, Assumption 5 guarantees that the operator $M_{Y,Y}^{abjk}$ is bounded with distinctive eigenvalues. Following theorem XV.4.3.5 in Dunford and Schwartz (1971), we have that the diagonal decomposition of $M_{Y,Y}^{abjk}$ is unique up to the index of eigenvalues and eigenfunctions.

In order to fully identify each eigenfunction, we need to identify the exact value of y^* in each eigenfunction $f_{Y|Y^*,S}(\cdot|y^*, s^a)$. Here we use the ordering assumption in Hu and Schennach (2008), i.e., Assumption 6, to pin down the exact value of y^* for each eigenfunction $f_{Y|Y^*,S}(\cdot|y^*, s^a)$. Such an identification procedure can be applied to each subpopulation with a different value of s. Thus, we have fully identified the conditional density $f_{Y|Y^*,S}$.

Given $f_{Y|Y^*,S}$, other densities containing y^* can also be identified due to the injectivity of operator $M_{Y|Y^*,s}$ as follows: for given (z, s, l)

$$f_{Z,Y^*,S,L}(z, \cdot, s, l) = M_{Y|Y^*,s}^{-1} f_{Z,Y,S,L}(z, \cdot, s, l) \tag{13}$$

In summary, we have shown that the density $f_{Z,Y,S,L}$ uniquely determines the joint density $f_{Z,Y,Y^*,S,L}$ satisfying $f_{Z,Y,Y^*,S,L} = f_{Y|Y^*,S} f_{Z,Y^*,S,L}$. ■

The general results in Theorem 1 directly apply to the baseline setup in Eqs. (1) and (2) under assumptions as follows:

Assumption 7. In Eq. (1), $\epsilon_{i,t}^Y(s)$ for each $s \in \mathcal{S}$ is independent of $Y_{i,t}^*$,

and

Assumption 8. In Eq. (2), $\epsilon_{i,t}^Z(l)$ for each $l \in \mathcal{L}$ is independent of $Y_{i,t}^*$ and $\epsilon_{i,t}^Y(s)$ for all $s \in \mathcal{S}$.

We summarize the identification result under the baseline setup as follows:

Corollary 1. *Under Assumptions 3–8, the function $m(.)$ and the error distributions $f_{\epsilon^Y|S}(\cdot|s)$ and $f_{\epsilon^Z|L}(\cdot|l)$ for all s and l are nonparametrically identified in the baseline setup in Eqs. (1) and (2).*

This corollary is a direct application of Theorem 1. Assumptions 7 and 8 are for the simplicity of estimation and are stronger than necessary. The general identification results in Theorem 1 hold with a slightly more general specification, such as[12]

$$Y_{i,t} = Y_{i,t}^* + \sigma(Y_{i,t}^*)\epsilon^Y(S_i). \tag{14}$$

under the same assumptions as in Corollary 1, where $\sigma(.)$ is a nonparametric function.

In summary, we have presented a set of sufficient conditions under which all the distributions containing the latent true GDP growth can be uniquely determined by the observed joint distribution of GDP growth and night light growth from countries with different statistical capacity and at different locations:

Such a nonparametric identification result implies that consistent estimation is possible for parametric, semiparametric, or nonparametric specifications. In order to focus on the relationship between night light growth and the latent true GDP growth and also to take the sample size into account, we adopt the baseline specification to simplify the measurement error structure.

4.3. Specifications

The specification of the production function $m(\cdot)$ and the choice of the location variable are informed by the data and the assumptions required for identification.

To examine whether there is nonlinearity in $m(\cdot)$, we conduct ordinary least squares regression of night light growth on GDP growth, its second order term, as well as the cross-section average growth. Table 2 presents the regression results

[12] We thank an anonymous referee for pointing this out.

Table 2
Relationship between night light growth and GDP growth: Ordinary least squares estimation.

	Night light growth				
	(1)	(2)	(3)	(4)	(5)
GDP growth	0.531***	0.475***	0.595***	0.515***	0.662***
se, no clustering	(0.0869)	(0.0974)	(0.0681)	(0.109)	(0.0854)
se, clustered by year	(0.129)	(0.138)	(0.0761)	(0.166)	(0.0983)
average GDP growth		−0.842			
se, no clustering		(0.329)			
se, clustered by year		(1.520)			
GDP growth squared				0.227	−0.880
se, no clustering				(0.885)	(0.674)
se, clustered by year				(0.833)	(0.698)
Year fixed effects	–	–	Yes	–	Yes
Obs	3295	3295	3295	3295	3295
Adjusted R^2	0.0109	0.0126	0.439	0.0106	0.439

Notes: This table presents the results of ordinary least squares regressions of night light growth on GDP growth, average GDP growth by country, GDP growth squared, and year fixed effects. Standard errors (se) are calculated analytically. Standard errors without clustering are presented in the first row below the point estimates of each variable and standard errors clustered by year in the second row. Stars indicate the lower significance level implied by standard errors with and without clustering, i.e., we present the one with fewer stars. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

with and without clustered standard errors. They show that a linear production function with year fixed effects is sufficient to capture $m(\cdot)$. The coefficient on GDP growth is statistically significant and stable throughout, but the coefficient on the second order term of GDP growth is not. Year fixed effects increase the adjusted R^2 substantially, most likely because of differing atmospheric conditions in different years and the decay of satellite sensors over time that contribute to temporal variation in night light growth across all countries. Country fixed effects are not included because with first-differencing of night lights, time-invariant factors, such as the habit of using lights at night, are already taken into account. In fact, further including country fixed effects adds little to the explanatory power of columns (3) and (5). Clustering by year increases standard errors as expected but does not change the significance level of the point estimates of GDP growth, indicating that cross-sectional dependence at the country level, while present in the data, is not a major concern in our analysis.

Assumption 1 requires that night light growth is related to a country's geographic location, while Assumption 2 requires that the accuracy of GDP growth conditional on statistical capacity has nothing to do with location. To choose a location variable that satisfies these assumptions, we use a country's average cloud cover. Appendix B provides some evidence that cloud cover is a suitable choice.

Given the general nonparametric identification, we provide a sieve maximum likelihood (ML) estimator as suggested in Carroll et al. (2010).[13] We develop our estimator based on a panel data $\left\{Z_{i,t}, Y_{i,t}, S_i, L_i\right\}_{i=1,\ldots,N;t=1,\ldots,T}$ with a large N and a finite T, which is i.i.d across country i. As alluded to earlier, our estimator works with a cross-sectional data with $T = 1$. For the simplicity of estimation, we assume that distribution $f_{Z_{i,t},Y_{i,t},Y^*_{i,t}|S_i,L_i}(z, y, y^*|s, l)$ is time-invariant, i.e., the same for different t. The specification of the production function $m(\cdot)$ is informed by the data and assumed to be a linear function with year fixed effects. The error terms $\epsilon^Y_{i,t}(S)$ and $\epsilon^Z_{i,t}(L)$ are allowed to have a general density function. Therefore, in this empirical study, we adopt a parametric specification of function $m(\cdot;\theta)$ and leave other elements nonparametrically specified in the baseline specification in Eqs. (1) and (2). The asymptotic properties of the sieve ML estimator are presented in Online Appendix D with proofs.[14]

4.4. Constructing a new measure of GDP growth

With the conditional distribution of night light growth and national accounts GDP growth rates at hand, we construct a new and more accurate measure of GDP growth rates that optimally combines the information in these two measures.

Henderson et al. (2012) and Pinkovskiy and Sala-i-Martin (2016) construct new measures of GDP growth and GDP per capita through a linear combination of night lights with official or survey-based GDP measures. However, their combination uses constant weights. The new measure we construct here, which we call the optimal linear measure, will have a time-varying optimal weight on light-predicted GDP growth, because the entire distribution of measurement errors in both night light growth and national accounts GDP growth is uncovered. The information in night lights on GDP growth will differ for each observation.

[13] For a general description of sieve estimators, we refer to Chen (2007).

[14] The smoothness conditions are imposed through the sieve approximation by choosing a sieve functional space satisfying those smoothness conditions. In the meantime, we want to point out that not all the assumptions can be imposed in the estimation, which is not a unique issue in this paper.

The optimal linear measure is based on a linear combination of national accounts GDP growth and prediction by night lights:

$$\hat{Y}^*_{i,t} = \lambda_{i,t}\tilde{Y}_{i,t} + (1-\lambda_{i,t})Y_{i,t}, \tag{15}$$

where $\tilde{Y}_{i,t}$ is night light-predicted GDP growth, $Y_{i,t}$ is national accounts GDP growth, and $\lambda_{i,t}$ is the weight.

Obtaining prediction by night lights is necessary for converting night light growth into measures of GDP growth. The weight $\lambda_{i,t}$ determines the extent to which the new GDP growth measure depends on the prediction by night lights.

Crucially, to obtain the optimal weight $\lambda^*_{i,t}$, we need information on measurement errors in both $\tilde{Y}_{i,t}$ and $Y_{i,t}$, which in turn relies on the distributions of measurement errors in night light growth and national accounts GDP growth identified previously.

Night light-predicted GDP growth

Following the same specification in Section 4.3, we assume that night light-predicted GDP growth is a linear function of night light growth with year fixed effects,

$$\tilde{Y}_{i,t} = \delta Z_{i,t} + \gamma_t D^Y_t, \tag{16}$$

where we assume year dummies D^Y are deterministic functions of $(Z_{i,t}, S_i, L_i)$. Under this assumption, we have $\tilde{Y}_{i,t} = \tilde{Y}_{i,t}(Z_{i,t}, S_i, L_i)$. In practice, we regress national accounts GDP growth on night light growth plus year fixed effects and estimate $\tilde{Y}$ as the prediction of the regression.

Optimal linear combination

To obtain the optimal linear combination, we minimize the conditional mean squared error of our new measure, i.e.,

$$\lambda_{i,t} = \arg\min_{\lambda} E_{Y^*,\epsilon^Y}\left[(\hat{Y}^*_{i,t} - Y^*_{i,t})^2 | Z_{i,t}, S_i, L_i\right], \tag{17}$$

where $Y^*_{i,t}$ is true GDP growth. The conditional expectation is taken with respect to both the distribution of true GDP growth Y^* and the distribution of GDP growth measurement errors.

By Eqs. (1) and (2), this conditional mean squared error can be decomposed into two parts:

$$\begin{aligned} & E\left[(\hat{Y}^*_{i,t} - Y^*_{i,t})^2 | Z_{i,t}, S_i, L_i\right] \\ &= \lambda^2 E\left[\left(\tilde{Y}_{i,t} - Y^*_{i,t}\right)^2 | Z_{i,t}, S_i, L_i\right] + (1-\lambda)^2 E\left[\left(\epsilon^Y_{i,t}\right)^2 | S_i\right], \end{aligned} \tag{18}$$

where the equality holds because $\left(\tilde{Y}(Z_{i,t}, S_i, L_i) - Y^*_{i,t}\right)$ is conditionally independent of $\epsilon^Y_{i,t}$ as a result of Assumption 1. The first term in Eq. (18) captures the uncertainty in night light-predicted GDP growth, whereas the second term captures the uncertainty in national accounts GDP growth. Following Assumption 1, the first term can be calculated as

$$\begin{aligned} & E\left[\left(\tilde{Y}_{i,t} - Y^*_{i,t}\right)^2 | Z_{i,t}, S_i, L_i\right] \\ &= \frac{\int \left(\tilde{Y}(Z_{i,t}, S_i, L_i) - y^*\right)^2 f_{\epsilon^Z|L}(Z_{i,t} - m(y^*)|L_i) f_{Y^*|S,L}(y^*|S_i, L_i)\, dy^*}{\int f_{\epsilon^Z|L}(Z_{i,t} - m(y^*)|L_i) f_{Y^*|S,L}(y^*|S_i, L_i)\, dy^*}. \end{aligned}$$

With both $f_{\epsilon^Z|L}(Z_{i,t} - m(y^*)|L_i)$ and $f_{Y^*|S,L}(y^*|S_i, L_i)$ identified and estimated, this term can be calculated straightforwardly through numerical integration.

The optimal weight then depends on the relative uncertainty in these two measures of GDP growth and equals,

$$\lambda_{i,t} = \frac{E\left[\left(\epsilon^Y_{i,t}\right)^2 | S_i\right]}{E\left[\left(\tilde{Y}_{i,t} - Y^*_{i,t}\right)^2 | Z_{i,t}, S_i, L_i\right] + E\left[\left(\epsilon^Y_{i,t}\right)^2 | S_i\right]}. \tag{19}$$

$\lambda_{i,t}$ is always in $[0, 1]$, which makes the optimal linear measure bounded by national accounts GDP growth and light-predicted GDP growth. In addition, $\lambda_{i,t}$ is country and time dependent.

5. Results

5.1. Night light production function

Table 3 displays the estimated night light production function for DMSP/OLSS data. The estimates are statistically significant at the 0.01 level. As the specification is linear, the elasticity of night lights with respect to GDP is conveniently θ_0. It is larger than the point estimate in column (2) of Table 2, suggesting that measurement errors play an important role in biasing the ordinary least squares estimate downward.

Table 3
Estimated light production function: Sieve maximum likelihood estimation.

Parameter	θ_0	θ_1
Point estimate	1.317	0.133
Bootstrap standard error	(0.171)	(0.004)

Notes: Standard errors are based on 400 sample bootstraps.

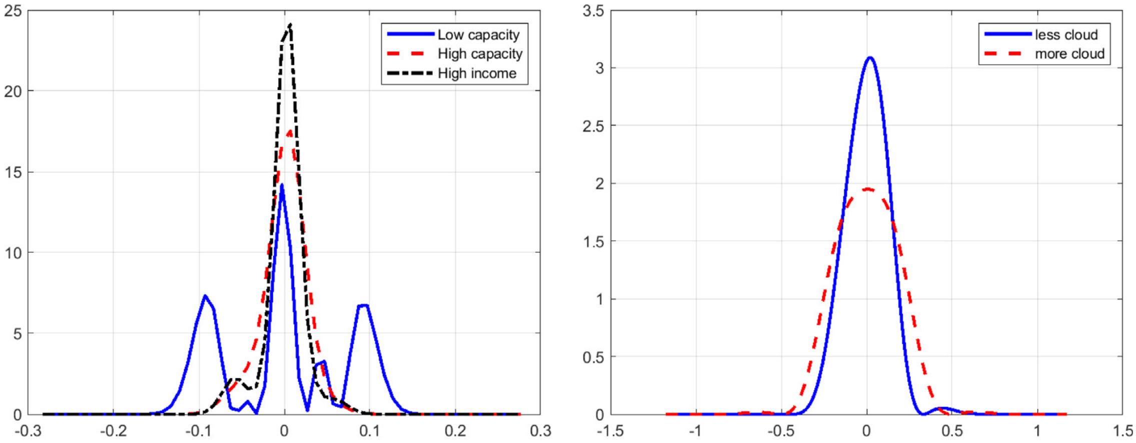

(a) GDP growth

(b) Night light growth

Fig. 2. Distribution of measurement errors. Notes: Of low and middle income countries, high capacity and low capacity refer to those above and below the median of statistical capacity, respectively. The spikes in the estimated probability density functions result from the choice of base functions (Hermite orthogonal polynomials) used to approximate the distributions of measurement errors. As the sample size increases, these spikes will typically be smoothed out..

5.2. Uncertainty in national accounts GDP growth

Fig. 2 presents the estimated probability density function of measurement errors in national accounts GDP growth and night light growth. The probability density functions are each approximated by Hermite orthogonal polynomials.

Panel (a) shows that the distribution of measurement errors of high income countries' GDP growth is concentrated at zero, indicating relatively high precision in official figures. By contrast, low and middle income countries' distribution of measurement errors has fat tails, suggesting that measurement errors are generally of larger size. Among low and middle income countries, those with high statistical capacity tend to have smaller measurement errors than those with low statistical capacity.

Panel (b) displays the distribution of measurement errors in night light growth. It has a much wider domain compared to the distribution of measurement errors in GDP growth, suggesting that night light growth data contain sizable noise. Observations that are less affected by cloud cover have a tighter distribution of measurement errors than those more affected by cloud cover. This is intuitive as cloud cover tends to block night lights, with more cloud cover leading to more measurement errors of night lights.

5.3. New measure of GDP growth

The new measure of GDP growth presented in Eq. (15), namely, the optimal linear measure, depends on night light-predicted GDP growth, national accounts GDP growth, as well as the optimal weight to combine them. There are two factors that determine the extent to which the optimal linear measure differs from official GDP growth. One is the discrepancy between night light-predicted GDP growth and official GDP growth and the other is the optimal weight.

In general, night light-predicted GDP growth is close to national accounts GDP growth, as implied by the relatively high correlation between night light growth and GDP growth when year fixed effects are taken into account (Table 2). Table 4 compares the mean squared error of regression (16) between countries with different income status. The mean squared error of light-predicted GDP growth for low and middle income countries is larger than that of high income countries, indicating greater discrepancy between light-predicted GDP growth and official figures for low and middle income countries.

Table 4
Root mean squared error by income status.

World Bank country classification	Low income	Lower middle income	Upper middle income	High income
Root mean squared error	0.044	0.041	0.044	0.037
# observations	1126	957	488	724

Notes: This table compares the MSE of the regression $Y_{i,t} = \delta Z_{i,t} + \gamma_t D_t^y + \eta_{i,t}$ for countries of different income groups.

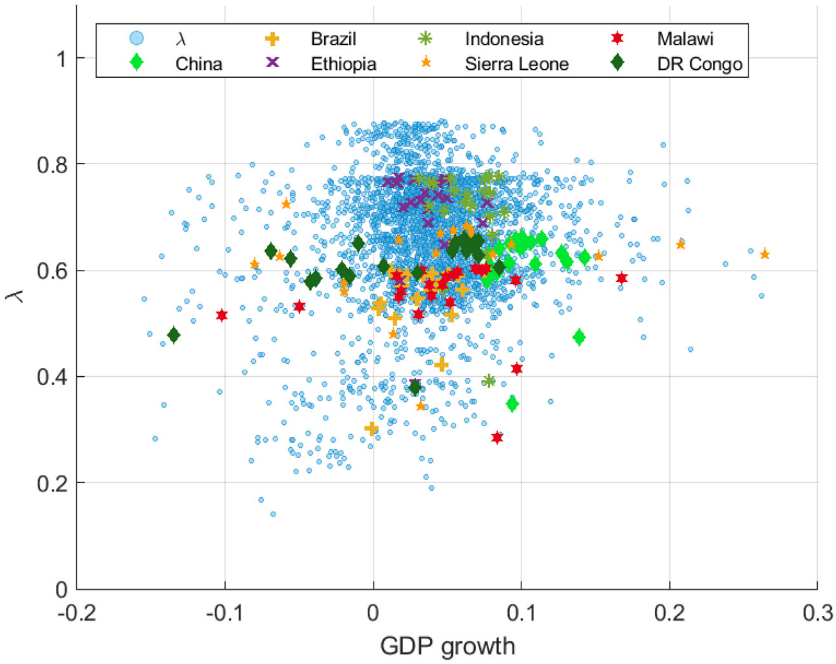

Fig. 3. Optimal weights and GDP growth: DMSP/OLSS 1992–2013.

Fig. 3 contrasts the optimal weights against GDP growth where we highlight a few low and middle income countries. The optimal weights are calculated according to Eq. (19) and they range from 0.14 to 0.88 with a mean of 0.65. Such non-negligible weights suggest that night lights can play a useful role in improving GDP growth measures for the majority of middle and low income countries.

5.4. Comparison of official and new measures

We examine two types of countries of particular interest. The first type is large emerging market economies for which official GDP growth is occasionally cast into doubt. The second type is low income countries that are often torn by war and conflicts and have weak statistical capacity.

Fig. 4 compares the new measure of GDP growth with official GDP growth for some of the BRICS countries: China, India, Brazil, and South Africa.[15] Panel (a) indicates that based on the new measure, China's GDP growth was much lower than suggested by national accounts data over the past decades.[16] Between 1993 and 2013, the difference between their geometric average is 3.6 percentage points every year. Panel (b) shows a similar result for India, though the difference is much smaller compared to China, averaging 1.7 percentage points every year. Panels (c) and (d) show that for Brazil and South Africa, the new measure is slightly higher than official GDP growth.[17]

Table 5 breaks down the sample period into three segments, separated by the 2001 and 2008 recessions in the United States, and presents the contrast between official and new GDP growth. For the BRICS countries, the difference between official and new GDP growth is systematic: for China and India, official GDP growth is consistently below the new measure; for Brazil and South African, it is consistently above the new measure.

15 We exclude Russia because it is a major oil producer, where night lights can be affected by gas flares substantially.

16 A possible explanation is that a nontrivial portion of China's investment was on unproductive assets, which were recorded in national accounts, but were not reflected in night lights.

17 It is quite challenging to analyze the asymptotic properties of the estimated $\lambda_{i,t}$, which corresponds to each observation. As such we leave the standard errors of these estimates for future research.

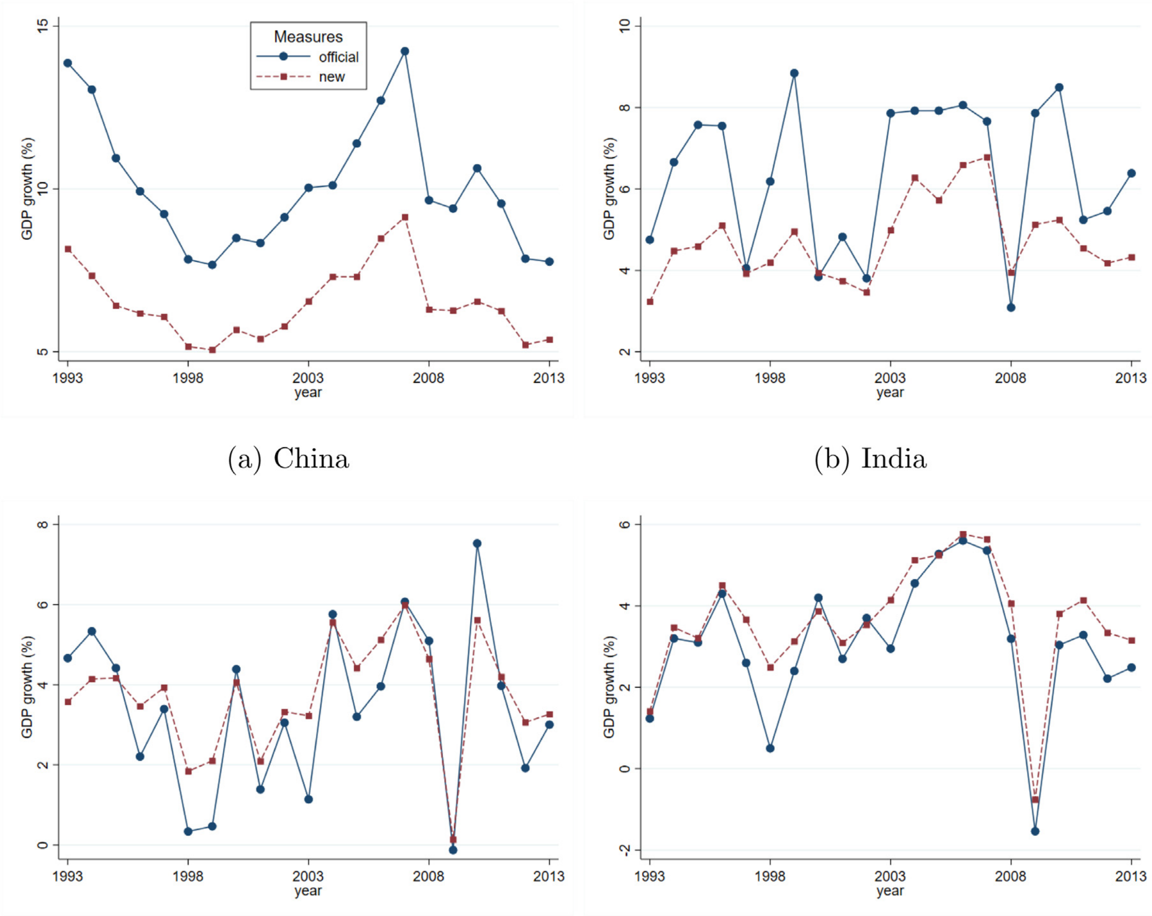

(a) China (b) India

(c) Brazil (d) South Africa

Fig. 4. GDP growth for selected emerging markets: 1993–2013.

Table 5
Average GDP growth for selected countries by official and new measures.

	1992–2001		2001–2007		2008–2013	
	Official	New	Official	New	Official	New
China	9.9	6.2	11.3	7.4	9.1	6.0
India	6.0	4.2	7.2	5.6	6.1	4.6
Brazil	2.9	3.3	3.9	4.6	3.5	3.5
South Africa	2.7	3.2	4.6	4.9	2.1	2.9
Dem. Rep. of Congo	−4.3	0.2	5.5	5.2	6.4	4.7
Ethiopia	5.9	4.4	7.7	6.2	10.4	5.7
Kenya	2.4	3.1	4.6	4.8	4.7	4.2
Sierra Leone	−1.5	1.5	9.6	6.9	9.2	6.0

Notes: Average growth rates for each period are geometric average and in percent.

That the new measure suggests lower GDP growth for China and India means that night lights did not grow as fast as their official GDP growth would have implied. In other words, light-predicted GDP growth is lower than official GDP growth, and because the new measure is bounded by the two growth rates, it is also lower than official GDP growth.

Fig. 5 examines the new measure against official GDP growth for a few low income countries that experienced war or conflict between 1993 and 2013, including the Democratic Republic of the Congo, Ethiopia, Kenya, and Sierra Leone. The

Fig. 5. GDP growth for selected low income countries: 1993–2013.

new measure tends to be smoother than official GDP growth. A potential factor behind this is the interaction between the formal and informal sectors of these economies. During economic downturns, formal economic activity disappears and informal economic activity springs up. While not recorded by official data, informal economic activity is captured by night lights. As a result, the new measure suggests higher growth than official data. Conversely, the flow of economic activity from the informal to the formal sector during economic upturns could make official GDP growth higher than the new measure.

6. Conclusion

In this paper, we first provide a statistical framework to describe the relationship among night light growth, national accounts GDP growth, and true GDP growth. We make use of the variation of observed night light growth and national accounts GDP growth across different levels of statistical capacity and geographic locations, and provide sufficient conditions under which the joint distribution of observables and the latent true GDP growth is uniquely determined by the distribution of observables. We estimate the linear production function of night lights and obtain an elasticity estimate of night lights to GDP, which is about 1.3. We find that national accounts GDP growth is less precise for low and middle income countries, and night lights can play a big role in improving such measures. Based on the estimated distributions, we construct a new measure of GDP growth. For China and India, our new measure is consistently below official GDP growth. For low income countries, our new measure is generally smoother than official GDP growth, likely due to the existence of the informal sector. We expect our statistical framework and methodology to have a broad impact on measuring GDP using additional information.

Appendix A. More details on nighttime lights data

A.1. DMSP/OLSS

The U.S. Air Force Defense Meteorological Satellite Program (DMSP) Operational Linescan System (OLSS) has been collecting global low light imaging data since the 1970s. DMSP satellites overpass at local time in the 7 p.m. to 9 p.m. range,[18] and nighttime lights are a class of derived products of the low light imaging data in spectral bands where electric lights emissions are observed. Elvidge et al. (1997) produced cloud-free composites from the raw data and the Colorado School of Mines currently hosts a digital archive of them that span from 1992 to 2013.[19] The cloud-free composites of nighttime lights were produced based on a set of quality criteria that remove observations affected by sunlight, moonlight, glare, aurora, and the edges of the DMSP/OLSS swaths.[20] For some years, there were two satellites collecting data and two composites were produced. In those cases, we use the average of the two composites. Each pixel of the DMSP/OLSS nighttime lights images is a 30 arc-second grid (slightly less than 1 square kilometer). It is associated with a digital number from 0 to 63 that is increasing with brightness.

A.2. Distributions

DMSP/OLSS nighttime lights are top-coded as a result of sensor saturation, often raising concerns about their validity as a proxy for economic activity. We examine this issue using data in 2013, which is one of the brightest years in the DMSP/OLSS sample–the world economy is growing and the world generally becomes brighter over time.

Fig. 6 presents the distributions of nighttime lights in 2013 for selective countries representative of distinct country types. For each country, we count the number of pixels for each discrete value $(1, 2, \ldots, 63)$ and calculate its ratio to the total number of pixels with positive values in that country. For mainland China, the maximum value (63) accounts for only 1.07% of all lit pixels. For the United States, it is 2.37%. Singapore has 81.2% of lit pixels reach saturation while Sierra Leone has none. This example highlights that top coding in DMSP/OLSS nighttime lights makes it not suitable for studying high-income cities or city states, which tend to have quality statistics anyway, but may be much less of a problem if the focus is on low and middle income countries or even advanced countries with large rural areas.

A.3. Gas flaring in the data

While nighttime lights primarily reflect economic activity for a majority of countries, it is recognized in the literature (for example, Henderson et al., 2012) that the flaring of natural gas might make nighttime lights incommensurate with the level of economic development. To examine the extent to which gas flares affect our results, we use gas flare shapefiles (polygons) provided by the National Oceanic and Atmospheric Administration[21] and calculate the fraction of nighttime lights within gas flare shapefiles in a country.

As an example, Fig. 7 shows a map of Nigeria where the white area is the shapefile that contains gas flares and the green area is the rest of the country. We obtain the fraction of nighttime lights in gas flare shapefiles (denoted by τ) by taking the sum of the nighttime lights within the white area first and then divide the sum by the total sum of lights in Nigeria.

Table 6 ranks countries by the fraction of nighttime lights in gas flare shapefiles in descending order. Among the highest are Equatorial Guinea, Gabon, and Nigeria. Most of the countries with high ranks are African and Middle East oil producers. While China and the United States produce oil, the vast majority of nighttime lights were produced in areas outside of the gas flare shapefiles. In particular, nighttime lights within gas flare shapefiles account for about 1 percent of total nighttime lights in China between 1992–2013 and even less for the United States. When dropping countries with $\tau \geq 0.02$ (34 countries) and reestimating our model, we find no statistically significant difference in estimated coefficients of the nighttime light production function compared to the full sample.

Appendix B. Cloud cover as the location variable

The nonparametric identification in Section 4.2 requires a location variable that satisfies Assumptions 1 and 2. In other words, it requires that night light growth is related to the location variable but the accuracy of GDP growth conditional on statistical capacity unrelated to the location variable.

[18] See Elvidge et al. (2009) for an overview.

[19] Night lights data can be downloaded here: https://eogdata.mines.edu/dmsp/downloadV4composites.html.

[20] A detailed description of the selection criteria can be found here: https://eogdata.mines.edu/dmsp/gcv4_readme.txt.

[21] https://ngdc.noaa.gov/eog/interest/gas_flares_countries_shapefiles.html

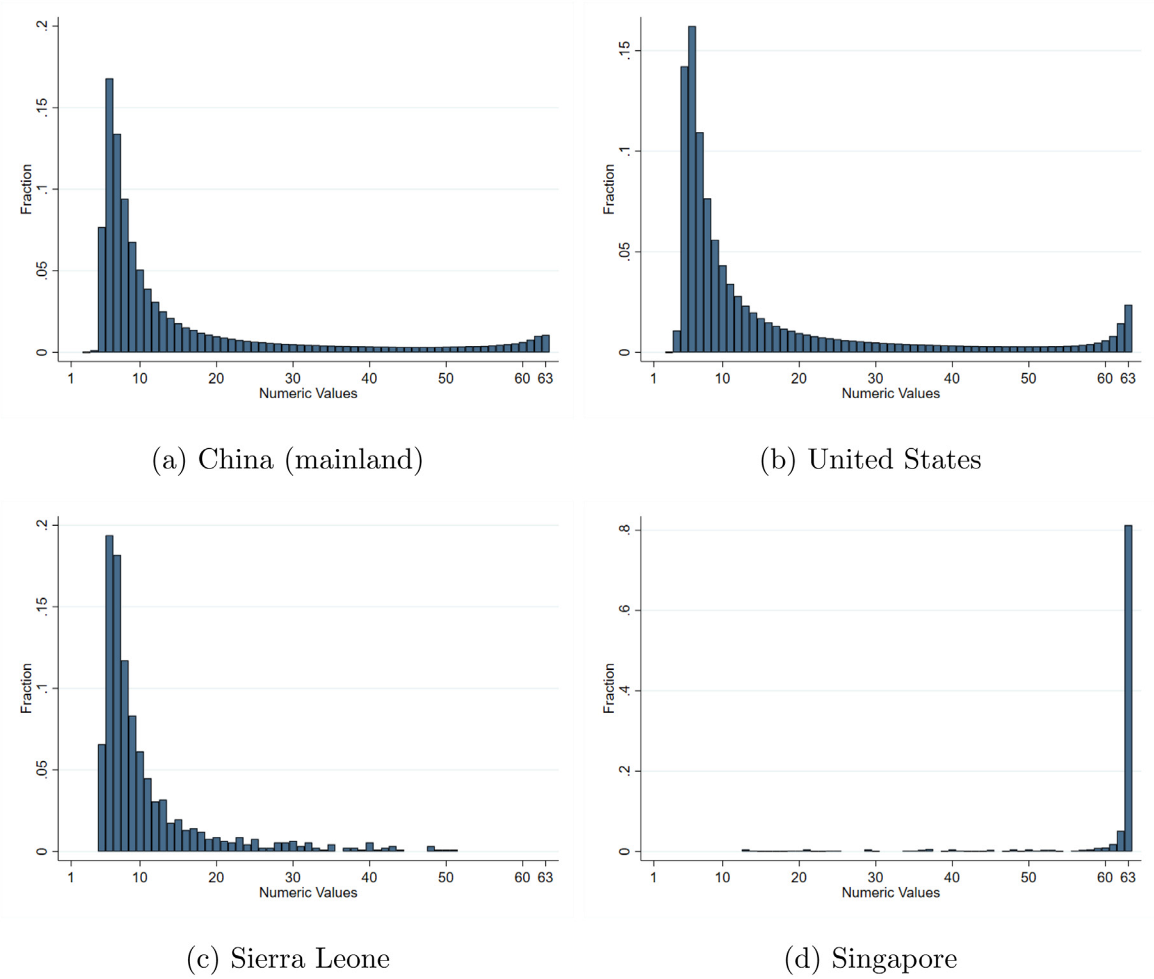

Fig. 6. 2013 DMSP/OLSS Nighttime Lights Distributions.

Table 6
Fraction of nighttime lights in gas flare shapefiles.

ISO code	τ	ISO code	τ	ISO code	τ	ISO code	τ
GNQ	0.75	RUS	0.17	VEN	0.08	AUS	0.01
GAB	0.69	AGO	0.16	ARE	0.06	CHN	0.01
NGA	0.57	SYR	0.15	SAU	0.05	CMR	0.01
LBY	0.36	ECU	0.15	EGY	0.04	MYS	0.00
IRQ	0.32	TKM	0.15	ARG	0.04	USA	0.00
KWT	0.31	TCD	0.14	SDN	0.03	ZAF	0.00
OMN	0.31	BOL	0.13	COL	0.03	PHL	0.00
DZA	0.30	PNG	0.13	CHL	0.03	BRA	0.00
COG	0.26	COD	0.11	PER	0.02	NOR	0.00
QAT	0.26	IRN	0.10	TUN	0.02	TTO	0.00
KAZ	0.24	UZB	0.09	CAN	0.01	...	...
YEM	0.20	IDN	0.09	AZE	0.01	...	...

Notes: τ is the ratio of the sum of nighttime lights in gas flare shapefiles to that in the whole country in DMSP/OLSS data, averaged between 1992–2013.

We consider two possible choices of the location variable. A natural choice is a country's location as measured by its average latitude.[22] The other is a country's average cloud cover that affects the accuracy of night lights.

22 To obtain a country's latitude, we use the latitude of the centroid of its largest contiguous block. For instance, the United States has a few separate bodies of land mass, such as Alaska, Hawaii, and the lower 48 states. We use the centroid of the lower 48 states as the location of the United States.

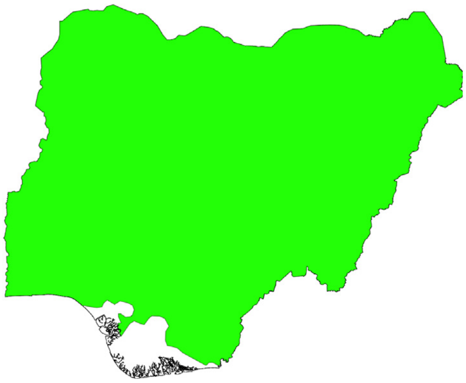

Fig. 7. Gas flares in Nigeria. Uncolored areas indicate regions that have gas flares.

Table 7
Approximated measurement errors in GDP growth rates and covariates.

	RMS difference of GDP growth rate vintages					
	(1)	(2)	(3)	(4)	(5)	(6)
Statistical capacity	−0.0470***		−0.0451***		−0.0468***	
	(0.0150)		(0.0169)		(0.0150)	
Latitude (absolute)		−0.0279**	−0.00436			
		(0.0114)	(0.0178)			
Cloud cover				−0.00114	0.0121	
				(0.0146)	(0.0174)	
Statistical capacity (discrete)						−0.668**
						(0.260)
Cloud cover (discrete)						−0.0493
						(0.392)
Obs	127	162	127	162	127	162
Adjusted R^2	0.0658	0.0299	0.0587	−0.00621	0.0619	0.0296

Notes: Standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Statistical capacity (discrete) is equal to 1 for low capacity countries, 2 for high capacity countries, and 3 for high income countries. Cloud cover (discrete) is equal to 1 for low cloud cover and 2 for high cloud cover.

To get a proxy for measurement errors in GDP growth rates, we use the root mean squared (RMS) difference of vintage GDP growth rates for the sample period between two vintages: July 2014 and October 2019.[23] Larger RMS implies larger revisions to GDP growth rates and larger measurement errors. Since it is difficult to get a proxy for measurement errors in night light growth, we simply consider its standard deviation at the country level to examine its correlation with the location variable.

Column (1) of Table 7 shows statistical capacity is correlated with the accuracy of GDP growth rates. Stronger statistical capacity is associated with smaller GDP growth revision. Columns (2) and (3) show that while latitude is also correlated with the accuracy of GDP growth rates, once we control for statistical capacity, the correlation becomes insignificant. By contrast, columns (4) and (5) show that cloud cover is unrelated with the accuracy of GDP growth rates whether or not statistical capacity is controlled for. Taken together, cloud cover seems a better choice than latitude in terms of satisfying Assumption 2.

In the baseline setup, we discretize statistical capacity and cloud cover. As explained in Section 3.2, we bin statistical capacity into three groups: low capacity countries, high capacity countries, and high income countries, which take the value of 1, 2, and 3, respectively. Since statistical capacity changes slowly over time, such grouping alleviates the concern that statistical capacity itself might be measured with errors. Cloud cover is binned into two groups, equal to 1 for low cloud cover and 2 for high cloud cover. Column (6) of Table 7 shows the result in column (5) holds when we use the discrete version of statistical capacity and cloud cover.

Table 8 shows that the standard deviation of night light growth has a weak correlation with latitude across countries, but a strong one with cloud cover. This suggests that cloud cover better satisfies Assumption 1 than latitude. The result in column (2) holds when we use the discrete version of cloud cover in column (3).

Based on Tables 7 and 8, we choose cloud cover as our location variable.

[23] Both vintages are from the World Development Indicators Database Archive. The July 2014 version is chosen because it is the one immediately after the sample period, while the October 2019 version is one of the latest. We thank the anonymous referee for pointing us in this direction.

Table 8
Variation of night light growth and covariates.

	Std. of night light growth rates		
	(1)	(2)	(3)
Latitude (absolute)	−0.000535		
	(0.000473)		
Cloud cover		0.00185***	
		(0.000577)	
Cloud cover (discrete)			0.0450***
			(0.0157)
Obs	162	162	162
Adjusted R^2	0.00173	0.0543	0.0431

Notes: Standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Cloud cover (discrete) is equal to 1 for low cloud cover and 2 for high cloud cover.

Appendix C. Estimation and simulation

In this section, we provide more details on estimation, conduct simulation exercises to confirm our estimation strategy, and do a number of robustness checks with respect to data and specification.

C.1. Density function approximation

Let $f(x) = h(x)^2$ be a generic density function with support $(-\infty, \infty)$, where $h(x)$ is a continuous function approximated by orthogonal Hermite series $\{p_n(x)\}$:

$$h(x) = \sum_{n=0}^{N} \beta_n p_n(x),$$

where

$$p_n(x) = \sqrt{\frac{1}{\sqrt{\pi}\, n! 2^n}} H_n(x) e^{-\frac{x^2}{2}},$$

and

$$\begin{aligned} H_0(x) &= 1, \\ H_1(x) &= 2x, \\ H_{n+1}(x) &= 2xH_n(x) - 2nH_{n-1}(x). \end{aligned}$$

It follows that

$$\begin{aligned} f(x) = &\sum_{n=0}^{N} \beta_n^2 p_n^2(x) \\ &+ 2 \sum_{n=0}^{N-1} \sum_{m=n+1}^{N} \beta_n \beta_m p_n(x) p_m(x). \end{aligned}$$

In order for $f(x)$ to be a density function, we must have $\int f(x)\, dx = 1$. To satisfy Assumption 6, we need measurement errors of GDP growth to have zero mean for at least one group among low capacity, high capacity, and high income countries. In other words, $\int x f(x)\, dx = 0$ holds for at least one group when $f(x)$ is the density function of GDP growth measurement errors.

In light of the following identities,

$$\int_{-\infty}^{\infty} H_n(x) H_m(x) e^{-x^2}\, dx = \sqrt{\pi} 2^n n! \delta_{nm},$$

we have

$$\int_{-\infty}^{\infty} p_n(x) p_m(x)\, dx = \delta_{nm},$$

and

$$p_{n+1}(x) = \sqrt{\frac{2}{n+1}} x P_n(x) - \sqrt{\frac{n}{n+1}} P_{n-1}.$$

Table 9
Parameterization in simulations.

Group	Y^*	$\sigma(s)$	$\sigma(l)$
1	$Normal(-0.02, 0.03^2)$	0.003	0.03
2	$Normal(0.01, 0.05^2)$	0.003	0.04
3	$Normal(-0.01, 0.02^2)$	0.0028	0.03
4	$Normal(0.015, 0.03^2)$	0.0028	0.04
5	$Normal(0.005, 0.04^2)$	0.0025	0.03
6	$Normal(0, 0.05^2)$	0.0025	0.04

Table 10
Simulation results.

	θ_0		θ_1	
True value	1.5		−0.13	
	Point estimate	Simulation standard error	Point estimate	Simulation standard error
$k_z = 6, k_y = 6, k_{y^*} = 5$	1.483	0.031	−0.131	0.001
$k_z = 7, k_y = 6, k_{y^*} = 5$	1.485	0.029	−0.130	0.001
$k_z = 7, k_y = 7, k_{y^*} = 6$	1.484	0.031	−0.130	0.001

Notes: Point estimates are the average of 400 simulations.

As such,

$$\int_{-\infty}^{\infty} x p_n(x) p_m(x)\, dx = \sqrt{\frac{n+1}{2}}\delta_{n(m-1)}, \qquad m > n,$$

$$\int_{-\infty}^{\infty} x^2 p_n(x) p_m(x)\, dx = \frac{\sqrt{(n+1)(n+2)}}{2}\delta_{n(m-2)} \qquad m > n.$$

The restrictions for $f(x)$ to be a density function with zero mean are therefore

$$\sum_{n=0}^{N} \beta_n^2 = 1, \tag{20}$$

$$\sum_{n=0}^{N-1} \beta_n \beta_{n+1} \sqrt{2(n+1)} = 0. \tag{21}$$

C.2. Estimation details

As mentioned in Section 4.3, the nonparametric densities in the sieve MLE estimator are approximated by finite dimensional parametric representations, where the dimension depends on the sample size. We find that Hermite orthogonal polynomials work well as basis functions with just a few sieve terms.[24] Given the sample size of our data sets, we conduct simulation studies in the next subsection to choose the smoothing parameters in our sieve MLE estimator. With a sample size similar to the DMSP/OLSS sample, our simulation studies show that the estimates are stable with the number of sieve terms used for each density function at about 5. As such we choose 5 for the DMSP/OLSS sample.

C.3. Simulations

We consider a data generating process similar to Eqs. (1) and (2) with a linear night light production function.

There are six equal-sized groups of countries based on statistical capacity ($s_i = 1, 2, 3$) and location ($l_i = 1, 2$). Each group's true GDP growth distribution follows a normal distribution. Measurement errors in both GDP growth and night light growth follow normal distributions where the variances $\sigma(s)$ and $\sigma(l)$ differ for different groups. The linear function $m(\cdot)$ is assumed to have similar coefficients as point estimates in the DMSP/OLSS data. In the simulation, we draw 400 samples with size $N = 3000$. Table 9 shows the parameter details in the simulations.

Table 10 presents the simulation results with different choices of the number of orthogonal Hermite terms. Point estimates are close to the true values.

Appendix D. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jeconom.2021.05.007.

[24] Compared to Hermite polynomials, the drawback for Legendre polynomials is that they have bounded support while Fourier series require many more sieve terms to approximate density functions well.

References

Aruoba, S.B., Diebold, F.X., Nalewaik, J., Schorfheide, F., Song, D., 2016. Improving GDP measurement: A measurement-error perspective. J. Econometrics 191 (2), 384–397.
Bound, J., Brown, C., Mathiowetz, N., 2001. Measurement error in survey data. In: Handbook of Econometrics, vol. 5. Elsevier, pp. 3705–3843.
Bound, J., Krueger, A.B., 1991. The extent of measurement error in longitudinal earnings data: Do two wrongs make a right? J. Labor Econ. 9 (1), 1–24.
Canay, I.A., Santos, A., Shaikh, A.M., 2013. On the testability of identification in some nonparametric models with endogeneity. Econometrica 81 (6), 2535–2559.
Carroll, R.J., Chen, X., Hu, Y., 2010. Identification and estimation of nonlinear models using two samples with nonclassical measurement errors. J. Nonparametr. Stat. 22 (4), 379–399.
Chen, X., 2007. Large sample sieve estimation of semi-nonparametric models. In: Handbook of Econometrics, 6. Elsevier, pp. 5549–5632.
Chen, X., Nordhaus, W.D., 2011. Using luminosity data as a proxy for economic statistics. Proc. Natl. Acad. Sci. 108 (21), 8589–8594.
Donaldson, D., Storeygard, A., 2016. The view from above: Applications of satellite data in economics. J. Econ. Perspect. 30 (4), 171–198.
Dunford, N., Schwartz, J.T., 1971. Linear Operators Part 3: Spectral Operators. vol. 7, JohnWiley & Sons, NewYork.
Elvidge, C.D., Baugh, K.E., Kihn, E.A., Kroehl, H.W., Davis, E.R., 1997. Mapping city lights with nighttime data from the DMSP operational linescan system. Photogramm. Eng. Remote Sens. 63 (6), 727–734.
Elvidge, C.D., Baugh, K., Zhizhin, M., Hsu, F.C., Ghosh, T., 2017. VIIRS night-time lights. Int. J. Remote Sens. 38 (21), 5860–5879.
Elvidge, C.D., Erwin, E.H., Baugh, K.E., Ziskin, D., Tuttle, B.T., Ghosh, T., Sutton, P.C., 2009.
Feng, S., Hu, Y., 2013. Misclassification errors and the underestimation of the US unemployment rate. Amer. Econ. Rev. 103 (2), 1054–1070.
Freyberger, J., 2017. On completeness and consistency in nonparametric instrumental variable models. Econometrica 85 (5), 1629–1644.
Ghosh, T., Powell, R.L., Elvidge, C.D., Baugh, K.E., Sutton, P.C., Anderson, S., 2010. Shedding light on the global distribution of economic activity. Open Geogr. J. 3, 147–160.
Gibson, J., Olivia, S., Boe-Gibson, G., Li, C., 2021. Which night lights data should we use in economics, and where? J. Dev. Econ. 149, 102602.
Griliches, Z., Hausman, J.A., 1986. Errors in variables in panel data. J. Econometrics 31 (1), 93–118.
Henderson, J.V., Storeygard, A., Weil, D.N., 2012. Measuring economic growth from outer space. Amer. Econ. Rev. 102 (2), 994–1028.
Hu, Y., 2021. The Econometrics of Unobservables: Latent Variable and Measurement Error Models and their Applications in Empirical Industrial Organization and Labor Economics. Online Manuscript. Johns Hopkins University.
Hu, Y., Schennach, S.M., 2008. Instrumental variable treatment of nonclassical measurement error models. Econometrica 76 (1), 195–216.
Mattner, L., 1993. Some incomplete but boundedly complete location families. Ann. Statist. 21 (4), 2158–2162.
Pinkovskiy, M., Sala-i-Martin, X., 2016. Lights, camera…income! Illuminating the national accounts-household surveys debate. Q. J. Econ. 131 (2), 579–631.