The World Bank Economic Review, 00(0), 2019, 1–19
doi: 10.1093/wber/lhz001
Article

OXFORD

Can Medium-Resolution Satellite Imagery Measure Economic Activity at Small Geographies? Evidence from Landsat in Vietnam

Ran Goldblatt, Kilian Heilmann, and Yonatan Vaizman

Abstract

This study explores the potential and the limits of medium-resolution satellite data as a proxy for economic activity at small geographic units. Using a commune-level dataset from Vietnam, it compares the performance of commonly used nightlight data and higher resolution Landsat imagery, which measures daytime light reflection. The analysis suggests that Landsat outperforms nighttime lights at predicting enterprise counts, employment, and expenditure in simple regression models. A parsimonious combination of the first two moments of the Landsat spectral bands can explain a reasonable share of the variation in economic activity in the cross-section. There is, however, poor prediction power of either satellite measure for changes over time.

JEL classification: E01, R1, O11

Keywords: remote sensing, nighttime lights, Landsat

1. Introduction

The analysis of satellite imagery is now a key methodology in economics and other applied scientific research. Coming straight from impartial satellites, remotely sensed data has the advantage of not being filtered through national data agencies that are potentially inefficient or biased. As Donaldson and Storeygard (2016) lay out its main benefits, remote sensing allows researchers to access information that would otherwise be difficult to obtain due to low state capacity, provides high spatial resolution, and a wide (if not global) geographic coverage. Since the marginal cost of collecting more data is low, repeated consistent samples are often available to researchers to learn about the world. Lastly, satellite imagery ignores administrative boundaries and can therefore be flexibly combined with other data at any geographical unit. Thus, satellites have enabled development economists to study geographic entities that were previously inaccessible because of insufficient data coverage. In future, new commercial satellite

Ran Goldblatt is Chief Scientist at New Light Technologies, Washington, DC; his email address is ran.goldblatt@nltgis.com. Kilian Heilmann (corresponding author) is a Postdoctoral Researcher at the Dornsife Institute for New Economic Thinking (INET) at the University of Southern California, Los Angeles, CA; his email address is heilmann@usc.edu. Yonatan Vaizman is a Senior Researcher at Comcast Applied AI Research, Washington, DC; his email address is yonatan_vaizman@comcast.com.
The authors are grateful to Gordon Hanson, Amit Khandelwal, Klaus Deininger, Fang Xia, Anh Pham for invaluable advice. They also thank participants at the USC Dornsife INET workshops, the Annual Meeting of the Urban Economics Association in Vancouver, the Pacific Conference for Development Economics for helpful comments. A supplementary online appendix is available with this article at *The World Bank Economic Review* website.

projects with continually improving spatial and temporal coverage will only reinforce the importance of remote sensing for academic studies.[1]

Especially, the use of nighttime lights as a proxy for economic activity is important to scholars in economics and other social sciences. Nightlight intensity has been used to approximate economic activity as light is believed to be a normal good with its consumption increasing at higher incomes. Henderson, Storeygard, and Weil (2012) have found strong correlation with GDP at the national level, and researchers have gone on to use nightlights at smaller geographic units. For example, Bleakley and Lin (2012) use nighttime light intensity to measure economic activity in their study on the economic persistence of defunct portage sites, and Alesina, Michalopoulos, and Papaioannou (2016) combines them with subnational ethnolinguistic maps. Similarly, Storeygard (2016) employs nightlights to study intercity transport costs and their impact on income of sub-Saharan African cities.

However, the predictive power of nighttime light at geographies smaller than the national level has recently been disputed in the literature. For example, while Mellander et al. (2015) conclude that nightlight data have a reasonably high correlation with economic activity in the developed-world context of Sweden, Bickenbach et al. (2016) find severe parameter instability between regions in India, Brazil, and the United States, and cast doubt that the correlation between nightlights and GDP at the country level carries over to the subnational level. Chen and Nordhaus (2011) find similar issues when scaling down from national to grid cell levels. The reasons for this are threefold: The common nightlight data produced by the National Oceanic and Atmosphere Administration (NOAA) are only available at a geospatial resolution of 30 arc seconds (about 1km). This resolution might be too coarse to capture the characteristics of small geographical units. In addition, nighttime lights have a tendency to extend into neighboring regions, a phenomenon called the blooming effect (Small, Pozzi, and Elvidge 2005; Abrahams, Oram, and Lozano-Gracia 2018), which complicates the identification of the actual source of lights. Lastly, nighttime light data are saturated at a certain threshold of light intensity which is often exceeded in very bright urban cores, thus making the analysis of nighttime light data at small geographies difficult. The degree to which these shortcomings of the data are relevant for studies in developing counties is however largely unknown, as economic data at small geographies are hard to find in the relevant context.

This paper presents an alternative to nightlights and uses Landsat imagery to predict economic activity. It is in the tradition of other papers that have explored the potential of other remotely sensed data products that are not based on nightlights. A variety of studies use high-resolution imagery and convolutional neural network models to extract object features that are then used to predict economic outcomes such as poverty (Jean et al. 2016; Engstrom, Hersh, and Newhouse 2017; Babenko et al. 2017; Perez et al. 2017). The drawback to these approaches is that they tend to use computationally intensive algorithms and rely on expensive proprietary satellite imagery. While these studies have made remarkable advances in the use of remote sensing in economics, applied researchers interested in measuring economic variables remotely might lack the skills or resources to use such methods or data.

This study therefore follows a different route and relies on medium-resolution imagery and simple regression models to predict economic outcomes from the sky. It looks at imagery from the Landsat program, a public domain satellite product that has been used to survey the earth over the last 30 years, and investigate its prediction power for economic activity. In contrast to nightlights, Landsat satellites measure light reflectance during daytime at different wavelengths of the spectral band. Having a much finer spatial (30 m) and temporal resolution (16 days) than nightlights, Landsat imagery has the potential to detect very different correlates of economic activity and income on the ground. Yet unlike high-resolution imagery that often operates at scales of less than one meter, the medium-resolution Landsat images do not allow for feasible object recognition. The empirical exercise explores the small-scale properties of Landsat imagery

1 Private satellite ventures likes Planet and Digital Globe have recently pushed the limits to spatial resolution of up to 0.5 meters/pixel.

to predict the distribution of economic activity and compare its performance to nightlights in Vietnam. A finely-geocoded nationwide dataset on enterprise counts, employment, and expenditure makes it possible to test the potential and limits of both nightlight and Landsat imagery at small geographic scales. The results indicate that Landsat outperforms nighttime lights at predicting measures of economic activity and consumption in parsimoniously specified regression models in the aggregate. Using only two simple Landsat spectral indices that are designed to capture certain land use categories can explain 25–35 percent of the variation in formal employment and enterprise counts as well as expenditure in the cross-section. More flexible specifications making use of more characteristics of the distribution of the raw Landsat data can improve this fit to 40–50 percent. Employing least absolute shrinkage and selection operator (LASSO) regularization and machine-learning type cross-validation techniques suggests that these relationships are stable and not just fitting statistical noise. In contrast to the cross-sectional results, the remotely sensed data have virtually no prediction power for changes in the time series.

The study adds to the literature on remote sensing by showing that parsimonious linear models of medium-resolution satellite data can capture the spatial variation of economic activity to a fair degree and are vastly superior to simple specifications of nightlight data in the context of a developing country. This provides a viable and inexpensive alternative for applied researchers to study areas with inferior data availability. The findings further substantiate concerns about the predictive power of nighttime lights at small geographies raised in previous studies. The negative result in the time-series cautions against the use of remotely sensed data in simple models to predict economic growth over time, as cross-section and time-series parameters are highly unstable in the setting of Vietnam.

The paper proceeds as follows: section 2 describes the economic and satellite data used in the study, section 3 compares the predictive power of nightlights and Landsat in the cross-section while section 4 does the same in the time series. Section 5 concludes and discusses the results.

2. Data and Descriptive Statistics

This paper uses different satellite imagery products to predict commune-level economic data from Vietnam. This section introduces the sources of the satellite measures and the economic data.

Economic Data

This study uses economic data at the production commune/ward level in Vietnam for the years between 2004 and 2012. In Vietnam, around 9,000 communes and wards make up the third-level administrative subdivision after districts and provinces, and cover the whole country. In cities, these are called wards, and in more rural areas they are called communes. For simplicity these are referred to as communes in the following. These units have a large variation in size and cover areas between 0.05 and 1,567 square kilometers, with a median area of 14.7 square kilometers. Communes therefore represent very different entities. While in cities, communes tend to be small and represent neighborhoods, there are several large communes that are sparsely populated and often sit along the western border with Laos. Even though the exact decision process of outlining commune borders is unknown, the formation of these units is highly endogenous. Unfortunately, the data do not come with population data from which a mechanism (such as a target population size within a commune) could be inferred. Figure 1 provides a graphical introduction to the shape of the commune units in relation to district boundaries. Because communes are too small to sensibly plot them for the whole country, two map inlets display the heterogeneity of their boundaries. While in rural areas, boundaries can be large with diameters of more than 10 kilometers, they can be very small in urban areas.

Figure 1. Administrative Boundaries of Vietnam

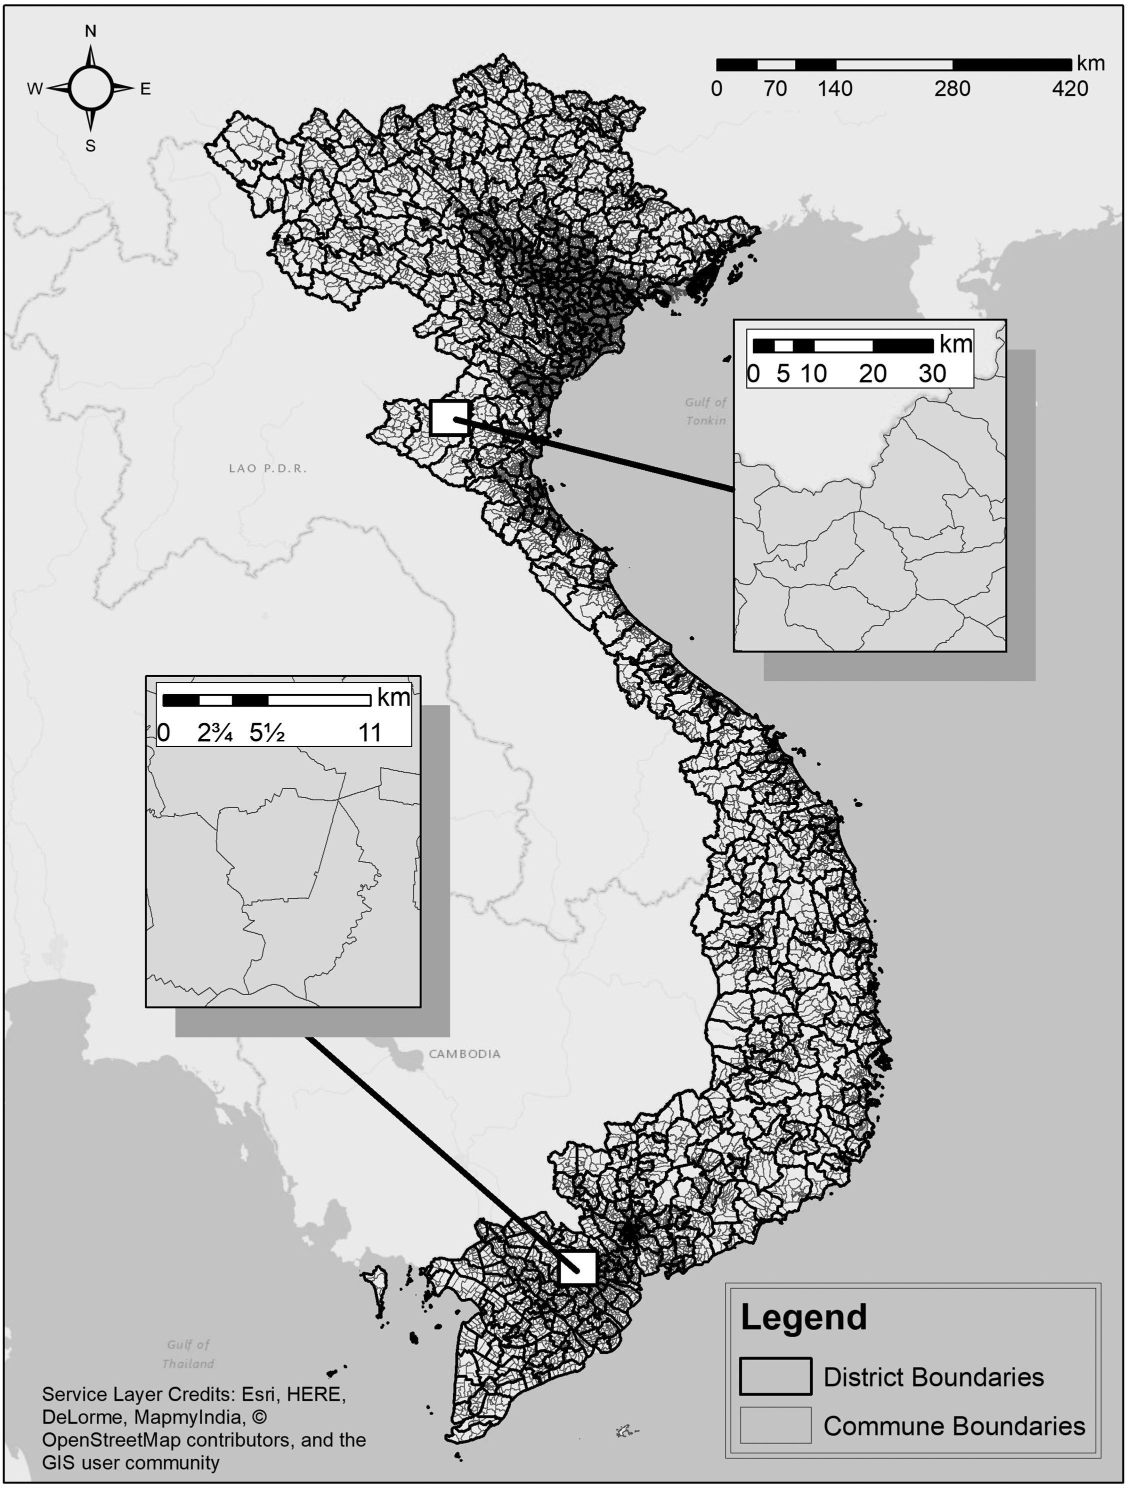

This study uses geocoded economic data for communes from the Enterprise Survey conducted annually by the Vietnam General Statistics Office.[2] This survey covers characteristics of firms such as employment, revenue, profits, and industry classification. The firms are geographically identified at the commune level. These data are aggregated to the commune in order to work with the total count of enterprises and counts of employment at this small geographic scale. Economic activity of each commune is then approximated by the number of workers and the number of enterprises employing them. The data come with several quality concerns, which are discussed here. At first, they are limited by the fact that only formal employment is recorded. In a country with a potentially very high share of the informal economy (Cling, Razafindrakoto, and Roubaus 2011), these numbers can severely underestimate the actual employment numbers, and results should therefore be interpreted as measures of the formal economy. Second, the data exclude self-operating farmers and only account for employment in agricultural enterprises. Further, the sampling scheme is designed to fully survey enterprises with more than 10 employees only, while a random sample is used for smaller firms, and is thus not fully representative. There are also inconsistencies over time as boundaries occasionally change.[3] Lastly, an issue with the geocoding arises due to the fact that the enterprise data are not based on plant location but rather on firm location. This potentially exaggerates measures of economic activity at the headquarters location and underestimates it at the branch location.

This study further employs real per capita expenditure data (measured in 2010 US dollars) from the Vietnam Household Living Standards Survey (VHLSS). This survey is conducted bi-annually and has several advantages: It follows the same communes over time and is not affected by boundary changes. It is generally believed to be nationally representative and has been used to compute official poverty statistics. However, there are certain limitations due to the survey design. While the household survey maintains a representative sample at the province level, it might not be so at the commune level (Hansen and Le 2013) and might undercount migrant workers (Pincus and Sender 2008). The household survey has a smaller geographical coverage than the enterprise data and includes about 3,000 communes. The expenditure data also suffer from small sample size as only three households are randomly drawn to be surveyed within each enumeration area.[4] Nevertheless, the sampling issue should not be correlated with the remote sensing measure and is therefore unlikely to bias the results.

Figure S3 in the supplementary online appendix plots the distribution of the three variables that will serve as the main measures of economic activity together with commune size on a log scale. While commune size, the number of employees, and per capita expenditure approximately follow a log-normal distribution, the number of enterprises is highly skewed and has the majority of its mass in the left tail. Summary statistics are tabulated in the supplementary online appendix.

Satellite Imagery

This study uses remote sensing imagery from the Defense Meteorological Program Operational Line-Scan System (DMSP-OLS) nighttime light dataset and the Landsat program. This section introduces the Landsat program and describes the processing of the data.

2 The authors thank Klaus Deininger at the World Bank and Fang Xia at the University of International Business and Economics, Beijing, for providing access and invaluable guidance relating to the data.

3 This study deals with this issue by aggregating up communes that have split or merged to create boundaries consistent over time.

4 For robustness checks regarding the low sample size, the supplementary online appendix (available with this article at The World Bank Economic Review website) also analyzes measures of incomes from the same survey that are based on a higher number of respondents.

Landsat Data

The Landsat program is managed by the United States Geological Survey (USGS) and consists of several satellites that capture global imagery at frequent intervals. Unlike satellites that measure nightlights, the Landsat satellites record daytime reflection of sunlight at different spectral bands. These bands include wavelengths of visible and invisible light, and thus provide rich information on objects on the ground. Depending on their structure, objects on earth reflect different portions of the electromagnetic spectrum, and therefore form a "spectral signature". Landsat data are organized into individual georeferenced and time-coded images and cover all of the earth except for the polar regions.

The Landsat program dates back to the early 1970s and has launched several satellites since then. The currently active satellites (Landsat 7 and Landsat 8) provide a spatial resolution of 30 × 30 meters, and their output can therefore be classified as "medium-resolution" imagery. While Landsat imagery lacks the detail of commercial mapping services such as Google Maps, the spatial resolution is fine enough to visually recognize certain features like cities and other land use. Besides visible light that can be recognized by the human eye, the satellites also capture nonvisible light such as infrared and thermal heat. While widely used in environmental sciences, the vast Landsat dataset has only recently been introduced into the social sciences, mostly in urban studies. Several papers have examined the predictive power of Landsat imagery for population counts and urban boundaries, such as Goldblatt et al. (2016). There is very little usage of Landsat in economics despite the fact that Landsat imagery is available easily at no cost.[5]

This paper uses Landsat 7, which launched in 1999 and is still operating to produce data.[6] Landsat 7 records eight different spectral bands and has a temporal resolution (the time until the satellite revisits a certain position on earth) of 16 days. Table 1 provides an overview of the recorded bands and their resolution. The analysis uses annual simple composite images of Vietnam. The simple composite algorithm corrects for the disturbance by cloud-coverage that is perennial in tropical regions and stitches together cloud-free images collected at different times. This comes at the loss of temporal accuracy as the data is not sourced from a single snapshot in time. In specific, a standard top-of-atmosphere (TOA) calibration is employed on all USGS Landsat 7 Raw Scenes in one year with less than 10 percent cloud coverage and then use the median of each pixel that satisfies this restriction. Several characteristics of the band distribution within each commune are extracted. For the empirical analysis, the mean, median, and standard deviation, as well as the 25th and 75th percentile of all bands listed in table 1 are calculated.

Table 1. Bands of Landsat 7

Bands	Wavelength	Resolution
Band 1 - Blue	0.45–0.52 μm	30 m
Band 2 - Green	0.52–0.60 μm	30 m
Band 3 - Red	0.63–0.69 μm	30 m
Band 4 - Near infrared (NIR)	0.77–0.90 μm	30 m
Band 5 - Shortwave infrared (SWIR) 1	1.55–1.75 μm	30 m
Band 6 - Thermal	10.40–12.50 μm	60 m
Band 7 - Shortwave infrared (SWIR) 2	2.09–2.35 μm	30 m
Band 8 - Panchromatic	0.52–.90 μm	15 m

Source: USGS Landsat.
Note: Landsat records data for Band 6 at two different sensitivities. Since they are highly correlated we proceed with only the less sensitive setting.

5 The major exception is Burchfield et al.'s (2006) study employing Landsat 5 imagery to measure urban sprawl.
6 The newest satellite, Landsat 8, collects data for additional bands, but was only launched in 2013.

Nightlights Data

For nighttime lights, the stable light band of the DMSP-OLS is used. In contrast to daytime Landsat imagery, these satellites measure light emitted from the globe at night. The dataset has a spatial resolution of 30 arc-seconds which, depending on the distance to the equator, is about 1 kilometer. Each pixel is coded with an integer value between zero (no light) and 63 (maximum light). This top coding due to saturation is an issue in bright city areas that easily hit the maximum sensitivity of the satellite sensor. While nightlights are in principle measured every 24 hours, the raw data require elaborate ex post processing, and datasets are typically released as annual composites.[7] The paper uses the stable lights product that removes unstable light sources such as moonlight, clouds, fires, and gas flares that create large outliers in the data (Baugh et al. 2010). Annual images are extracted to calculate the sum of light (SOL) of each commune. This commonly used measure simply adds up all pixel values within a commune.

3. Predictive Power of Satellite Imagery in the Cross-Section

This section presents the methodology to correlate satellite data with economic activity data and report the empirical results in the cross-section.

Econometric Setup

This study starts by exploring the predictive power of satellite imagery in the cross-section. It estimates simple linear ordinary least squares (OLS) regressions to predict economic variables using the remotely sensed satellite data for each sample year separately:[8]

$$\log y_i = \alpha + \beta\ satellitedata_i + \epsilon_i \tag{1}$$

where y_i denotes the economic outcome of commune i, and vector *satellitedata* contains commune-specific statistics from the remotely sensed satellite products. The parameters of interest are β and the R-squared measure as an indicator of the predictive power of remotely sensed data for economic outcomes.

Later specifications control for the area of the communes, which turns out to be an important predictor for economic outcomes. Commune area is not random, and accounting for the endogenous choice of commune boundaries is important. Figure S1 in the supplementary online appendix depicts the strong negative relationship between the number of enterprises and employees, expenditure, and commune size. On average, small communes have higher economic activity with estimated elasticities of –.52 (enterprises) and .65 (employment) respectively. The same negative relationship holds for per capita expenditure with an albeit smaller elasticity of –.22.

Empirical Results

Nighttime Lights

The nighttime lights serve as the baseline of the study. The analysis is based on regressions of the measures of economic activity (employment, number of enterprises, expenditure) on nightlight data from DMSP-OLS. This exercise focuses on one single year in 2012. Since commune boundaries change, the shapefile does not perfectly match the economic data provided for all years. The highest matching rate is for the year 2012, and therefore the analysis proceeds with this year.

7 The Visible Infrared Imaging Radiometer Suite (VIIRS) data product released by NASA now provides monthly releases from 2012 onward.

8 The supplementary online appendix explores the sensitivity of the results to alternative estimation specifications. Poisson models for count data yield measures of predictive power very similar to OLS.

This paper follow the sum of lights (SOL) approach that is commonly applied for indicating light intensity. The sum of lights approach counts up all nightlight sensor values (coded from 0 to 63) of pixels that fall within the area of a commune. To account for nonlinearities in the relationship, log-log models are estimated.[9] In the case of the nightlight approach with a singular regressor, the coefficient β can therefore be interpreted as the elasticity of the economic variable with respect to the total light emitted.

Table 2, panel A reports the results from the simple linear regression model in (1). In column (1), the estimated elasticity of the number of enterprises with respect to nightlights is 0.393 in the year 2012, indicating a strong relationship between economic activity and nightlights even at small geographies. The coefficient is significant at the 99.9 percent confidence level and measured with high precision. The R-squared measure of 0.09 however indicates a rather low predictive power of the nightlights in the cross-section. The employment regression yields similar results. The elasticity of the number of formal employees per commune with respect to the sum of lights is 0.563 and again measured with a

Table 2. Regression Results in the Cross-Section

Dep. Var.	Enterprises	Enterprises	Employment	Employment	Expenditure	Expenditure
Panel A: Nightlights						
Log SOL	0.393***	0.589***	0.563***	0.806***	0.0655***	0.126***
	(0.0148)	(0.0194)	(0.0197)	(0.0254)	(0.00779)	(0.00844)
Log Area		−0.631***		−0.784***		−0.225***
		(0.0146)		(0.0180)		(0.00762)
Observations	7643	7643	7643	7643	2629	2629
Adjusted R^2	0.090	0.348	0.103	0.325	0.023	0.328
Panel B: Landsat Spectral Indices						
Average NDBI	6.191***	5.817***	7.227***	6.678***	1.043***	0.752***
	(0.212)	(0.208)	(0.266)	(0.262)	(0.118)	(0.110)
Average NDVI	−4.129***	−2.926***	−4.894***	−3.131***	−1.732***	−0.916***
	(0.124)	(0.133)	(0.164)	(0.184)	(0.0784)	(0.0799)
Log Area		−0.219***		−0.321***		−0.148***
		(0.0137)		(0.0193)		(0.00822)
Observations	8493	8493	8493	8493	2990	2990
Adjusted R^2	0.341	0.362	0.257	0.282	0.274	0.349
Panel C: Combination						
Average NDBI	6.543***	6.099***	7.534***	6.887***	1.289***	1.193***
	(0.222)	(0.222)	(0.276)	(0.272)	(0.122)	(0.123)
Average NDVI	−3.678***	−3.815***	−4.191***	−4.392***	−1.405***	−1.429***
	(0.138)	(0.139)	(0.179)	(0.178)	(0.0856)	(0.0855)
Log SOL		0.364***		0.530***		0.0566***
		(0.0127)		(0.0176)		(0.00682)
Observations	7643	7643	7643	7643	2629	2629
Adjusted R^2	0.330	0.406	0.242	0.332	0.253	0.270

Source: Authors' calculations based on Enterprise Survey, VHLSS, Landsat, and DMSP-OLS data.
Notes: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. Standard errors are in parentheses.

9 The log-log specification drops communes with zero nightlights, of which there are 850. Using a log-linear prediction model does not alter the main results (see robustness checks in the supplementary online appendix).

small standard error. Again the nightlights explain only a rather small share of the between variation in employment ($R^2 = 0.103$). The elasticity of per capita expenditure with nightlights is much smaller and estimated with a β of 0.126. This indicates that (at least in the development country setting of Vietnam) nightlights are more responsive to differences in production but less so to consumption. The R-squared measure of expenditure is even lower at $R^2 = 0.023$, suggesting that nightlights are not very predictive of overall cross-section consumption patterns in the setting of Vietnam. This finding is in contrast to other studies that find higher prediction power for small geographic units of up to 40–50 percent in Africa (Jean et al. 2016) and of up to 60 percent in Sri Lanka (Engstrom, Hersh, and Newhouse 2017).

Controlling for potentially endogenous commune sizes confirms the strong positive relationship between nightlights and economic activity. Column (2) includes the log area (in square kilometers) as a regressor. The coefficient is strongly negative and indicates an elasticity of the number enterprise with respect to commune size of –0.631. The coefficient on log SOL increases to 0.589 and is still strongly significant. This is a natural consequence of the negative correlation between commune size and economic activity as outlined before. Similarly, in column (4) the employment elasticity increases to 0.806. For expenditure, the elasticity increases only slightly from 0.065 to 0.126. The considerable increase in the R-squared measure in all regressions indicates the predictive power of commune size independently of nightlights and confirms that commune boundaries are endogenous.

Landsat Spectral Indices

The focus is now on the question whether Landsat band values have similar predictive power for ground truth data of economic activity as nightlights. Since Landsat bands measure the reflectance of light of a certain wavelength and are difficult to interpret numerically, the analyses first proceed by using spectral indices derived from these bands. These measure certain land-use patterns and can be easily interpreted. The regressions use the normalized difference built-up index (NDBI) and the normalized difference vegetation index (NDVI) which are nonlinear combinations of two Landsat bands each.[10] These indices are defined as

$$NDBI = \frac{NIR - SWIR}{NIR + SWIR}, \quad \text{and} \quad NDVI = \frac{NIR - Red}{NIR + Red}$$

where *Red* corresponds to the Landsat band 3 (red light), NIR is the Near Infrared measurement of band 4, and SWIR is the value of the Shortwave Infrared band 5. These indices are designed to measure urban areas and vegetation by capturing typical spectral signatures of these features. Their values range between –1 and 1, and a higher index value corresponds to more vegetation and built-up area presence respectively. For example, the NDVI measure is designed to capture live green vegetation on the ground. For photosynthesis, live plants absorb visible light (low wavelengths) but reflect infrared light (higher wavelengths), thus a higher value of $NIR - Red$ indicates presence of vegetation. The NDVI has been used in the environmental sciences to distinguish vegetation from other land uses. Similarly, NDBI is designed to detect land cover typical of urban areas (Zha, Gao, and Ni 2003).

These indices are chosen for several reasons: (1) They are simple transformations of the raw Landsat data that can be easily calculated by applied researchers. (2) They directly correspond to land uses that are man-made (cities, agriculture) and are plausibly correlated with economic activity. (3) They have been commonly used in other fields and can in principle be replicated with other satellite products that measure spectral light such as Sentinel-2. All measures follow a distribution that can be approximated by a normal distribution, although there are thick tails in the right (NDBI) and the left (nightlights, NDVI) end of the distributions (see density figures in the supplementary online appendix).

10 A third commonly used indicator is the normalized difference water index (NDWI) measuring water bodies. In this setting, NDVI and NDWI are highly negatively correlated. The latter is excluded to avoid issues of multicollinearity.

The analysis proceeds by estimating equation (1) with the two Landsat indices above. Since these indices are already normalized and well bounded, there is no need to apply the natural logarithm. Table 2, panel B summarizes the results. In column (1), both indices are statistically significant in explaining differences in the (log) number of enterprises between communes. While a higher measure of built-up areas is positively correlated, the coefficient on NDVI is negative and indicates that the presence of vegetation predicts fewer enterprise in the cross-section. This is not surprising, as high NDBI and low NDVI values indicate the presence of cities. NDVI is often used as a measure of agricultural productivity, but as self-employed farmers are not covered by the enterprise survey, the vegetation measure does not reflect this. The R-squared measure of the model is 0.341, suggesting a rather high prediction power with the use of only two indices. This is about four times larger than in the nightlights model.

The coefficients in the regression on employment are largely similar, indicating that employment is higher in more urban places. The Landsat indices explain 25.7 percent of the variation in employment and thus about 2.5 times as much as the nightlight approach (10.3 percent). The coefficients in the expenditure model are much smaller in absolute value, but still highly significant at the 0.1 percent level. Again, the difference in predictive power of both remote sensing approaches is striking. The simple Landsat model explains about 10 times as much of the variation as the sum of light measure ($R^2 = 0.274$ versus $R^2 = 0.023$). Figure 2 summarizes the predictive power of the models by plotting the actual values versus the predicted values for all three economic outcome variables.

When controlling for commune size, the coefficients on NDBI and NDVI are reduced in absolute value, but remain highly statistically significant. This indicates some correlation between commune area and vegetation and built-up area respectively, highlighting the endogenous nature of commune boundaries. However, the correlation is less pronounced than in the nightlight regression. The increase in R-squared

Figure 2. Scatterplots: Predicted vs Actual Values

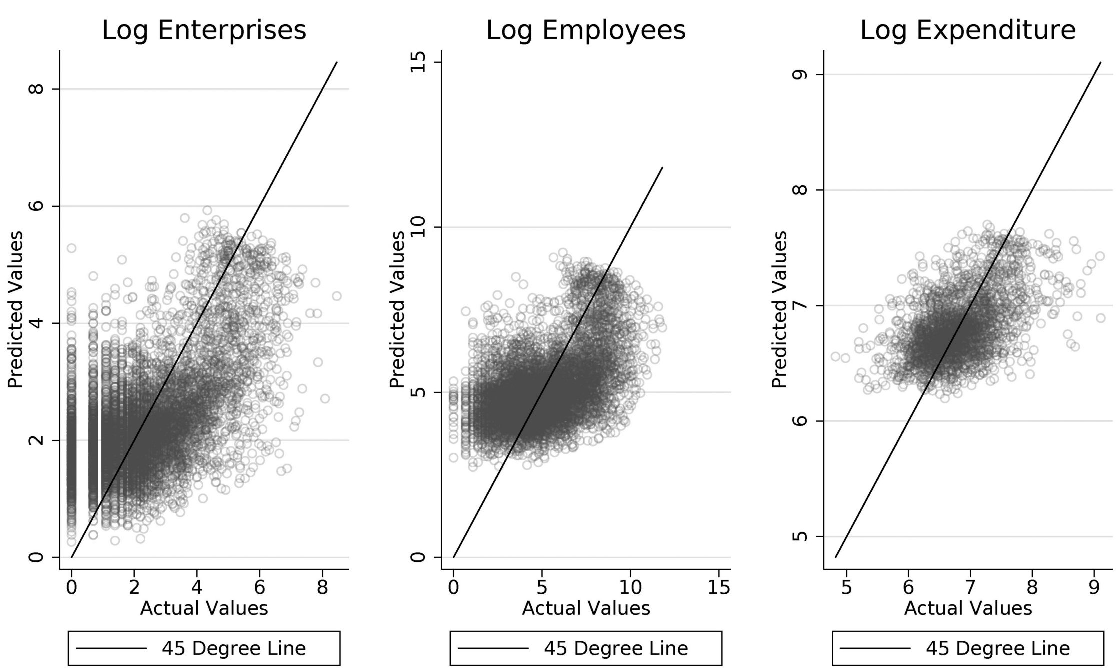

Note: These scatterplots show the predicted versus the actual values from the simple model with two Landsat spectral indices (NDBI and NDVI) for log enterprise, log employment, and log per capita expenditure in the cross-section of the year 2012. These prediction models create some overestimation of low-activity communes and tend to underestimate high-activity neighborhoods.

measures after controlling for commune size is smaller than for nightlights, suggesting that the Landsat measure is rather independent of the factors that drive commune size. In case a researcher is not comfortable using commune size as a predictor, for example in out-of-sample prediction in a different country, the results suggest that Landsat is the more powerful correlate of economic activity.

Combining Nightlights and Landsat

The focus is now on exploring combinations of nighttime light data and Landsat imagery as predictors for the economic outcome variables. Due to their different observation period during the day, the two satellites might pick up very different correlates of economic activity. Combining them as independent sources of variation could increase the predictive power. Initial analysis shows that at the commune level, nighttime lights are positively correlated with NDBI ($\rho = 0.17$) but negatively with NDVI ($\rho = -0.05$). This would be a natural conclusion if nighttime lights are indicative of urban settlements that are picked up by the NDBI measure whereas rural areas with high share of vegetation emit less light.

Table 2, panel C summarizes the regression results for the combinations of Landsat indices and nighttime light. Independent of the Landsat indices, the log sum of light measure remains a strong predictor of economic activity with a p-value of less than 0.001. Compared to the base regressions of the two measures in panels A and B, the coefficients on all three predictors are very similar and nested in each other's confidence band. This suggests that nightlights and Landsat pick up very different correlates of economic activity. Naturally, model fit measures increase when incorporating both satellite datasets. Percentage increases in R-squared are 23 percent for enterprises, 37 percent for employment, and 6.7 percent for expenditure. Combining the two data sources thus improves prediction accuracy for economic differences between communes.

Exploiting all Landsat Spectral Bands

Next, the predictive power of the full set of spectral Landsat bands in the cross-section is explored. The study regresses a "kitchen sink" specification that includes all Landsat bands as explanatory variables and then compares the fit with the results from the regressions with the simple indices only. Band averages as well as other characteristics (such as the standard deviation, median, interquartile range) of the band distributions within a geographical unit are used as predictors. While the Landsat indices such as NDBI provide an easy to interpret measure of ground characteristics correlated with economic outcomes, a researcher who is interested in remote sensing economic activity needs to form a prior belief on which spectral signatures to use as predictors. In contrast, using all spectral Landsat bands as predictive variables allows for an agnostic and flexible (yet difficult to interpret) way of recovering the statistical relationship between measures of economic activity and remotely sensed data.

However, the large number of potential variables poses the risk of overfitting the data and of mistaking noise as a valuable signal. To alleviate this issue, two kinds of exercises are performed to guide the econometric approach: LASSO techniques (Tibshirani 1996; Belloni, Chernozhukov, and Hansen 2014) to restrict the variable space and cross-validation methods from the machine learning literature. The sample is divided into a training and a testing dataset to judge out-of-sample validity of the estimated parameters. The algorithm randomly attributes 70 percent of the sample to a training dataset, estimates the coefficients, and then examines the out-of-sample fit for the remaining 30 percent testing dataset. This exercise is repeated 500 times to calculate the cross-validated R-squared as the average out-of-sample R-squared of these 500 draws.

Table 3 compares the regression results for the Landsat indices (column 1) with the ones of more flexible models. Column (2) first only uses the mean band values of each commune and then augments them with the standard deviation (column 3). Skewness of the distribution of each band is accounted for by including medians (column 4) and the interquartile range (25th and 75th percentile, column 5). Comparing specifications (1) and (2) for enterprise, using only the band means as predictors ($R^2 = 0.416$)

Table 3. Landsat Kitchen Sink Regressions

Variables included	Indices (1)	Means (2)	+Std. dev. (3)	+Median (4)	+Interquartile range (5)
Dependent variable: Log(enterprises per commune)					
Observations	8493	8493	8493	8493	8493
Adjusted R-squared	0.341	0.416	0.493	0.505	0.512
Crossvalidated R-squared	0.340	0.416	0.492	0.503	0.508
Variables selected by LASSO	2/2	8/8	16/16	23/24	37/40
Dependent variable: Log(employment per commune)					
Observations	8493	8493	8493	8493	8493
Adjusted R-squared	0.257	0.319	0.403	0.413	0.419
Crossvalidated R-squared	0.257	0.318	0.401	0.410	0.414
Variables selected by LASSO	2/2	8/8	16/16	22/24	39/40
Dependent variable: Log(per capita expenditure)					
Observations	2990	2990	2990	2990	2990
Adjusted R-squared	0.274	0.356	0.397	0.404	0.413
Crossvalidated R-squared	0.274	0.320	0.378	0.388	0.396
Variables selected by LASSO	2/2	8/8	16/16	24/24	35/40

Source: Authors' calculations based on Enterprise Survey, VHLSS, Landsat, and DMSP-OLS data.
Note: Crossvalidated R-squared is the average R-squared from 500 replications of splitting the data into a 70 percent training sample to estimate the model parameters and a 30 percent testing dataset to calculate the out-of-sample R^2.

yields already a better fit than NDBI and NDVI only ($R^2 = 0.340$), indicating that other band values than those to calculate the indices provide valuable information. The LASSO regression picks all of the band means, and the cross-validated R-squared is reduced only marginally compared to the full in-sample R-squared. Adding the standard deviations as predictors further increases both in-sample and out-of-sample R-squared measures ($R^2 = 0.492$), while adding further information about the distribution yields only very marginal improvements in the fit.

For employment and per capita expenditure, the results are similar. Using band means and standard deviations yields higher prediction power than the Landsat indices, while adding further distribution characteristics beyond the first two moments yields only small improvements in the R-squared measure. Throughout the analysis, the LASSO technique tends to select almost all potential regressors as informative. This is the case despite the very high correlation between certain Landsat bands. The cross-validation exercise yields very little evidence of overfitting and only in the expenditure regression are there small differences between in-sample and out-of-sample fit. This suggests that in the cross-section, the parameters of the Landsat bands are highly stable throughout the whole country for predicting economic activity.[11]

Exploring Heterogeneity of the Prediction Power

The analysis next explores sources of heterogeneity in the prediction power of both satellite products and addresses measurement errors that are likely correlated with the satellite measures. As stated above, the original motivation to introduce Landsat imagery was due to the expected difficulty of nightlights to capture very small but highly productive units. This section lends evidence to this by performing additional analysis of the prediction power of Landsat and nightlights by dividing the data into subsamples.

11 The study confirms an earlier result that SOL is an independent predictor of economic activity in unreported regressions. For each model specification, the LASSO approach picks the sum of lights measure as a predictor. Inclusion of this regressor can push the cross-validated R-squared measures of column (5) to 0.5499, 0.4538, and 0.4077 for enterprises, employment, and expenditure respectively.

First, the sample is divided by the share of formal employment. Since the data only measure economic activity of formal enterprises, there is concern that the share of informal employment will be correlated with the satellite data and thus will introduce a nonclassical measurement error. To check this, the paper follows McCaig and Pavcnik (2015) and calculates the share of informal workers at the district level using the 2009 Population and Housing Census.[12] The sample is then split into deciles of formal employment, and separate regressions for log employment at the commune level are run for each subsample.[13] Figure 3 shows the resulting R-squared measures for three specifications: (1) Sum of Lights, (2) a specification with only the two Landsat spectral indices (NDBI and NDVI), (3) the first two moments of all Landsat spectral bands. In the upper panel of fig. 3 it can be seen that, as expected, the Landsat measures perform extraordinarily well in communes with a high share of formal employment. These communes tend to be cities with smaller area size. The flexible Landsat specification dominates the nightlight approach at all

Figure 3. Heterogeneity of the Prediction Power

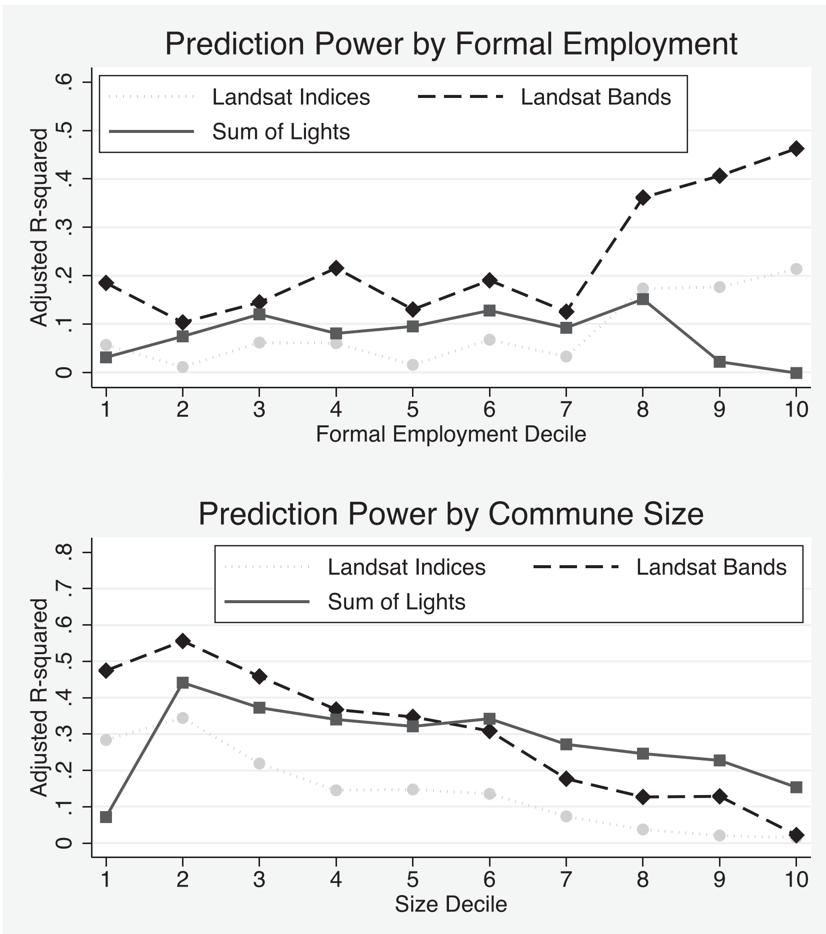

12 McCaig and Pavcnik (2015) define the share of informal employment by the share of all workers either being self-employed or working in family and farm businesses.

13 Due to the uneven number of communes within a district, the decile bins do not contain exactly the same number of communes.

deciles, while the simple indices are superior to the SOL measure only for communes with a high share of formal employment.

Next, the sample is split up into deciles of commune size, and again the study regresses individual models for each decile. The lower panel of fig. 3 shows the resulting R-squared measures according to commune size. As expected, the sum of lights approach does extremely poorly at predicting economic activity of very small communes. The prediction power then increases sharply for the next decile and then decays almost monotonously with size. In contrast, both Landsat specifications have high prediction power for the smallest communes. One potential reason is that Landsat does well at capturing the built-up area. There is similar to decay of prediction power with increasing commune size where the full Landsat band specification dominates the sum of lights approach for the lower half of the sample. The simple Landsat indices specification, however, performs less well than the nightlights beyond the smallest communes.

The analysis concludes that, consistent with earlier studies, nightlights perform poorly at predicting economic activity at very small geographic units. This drives the overall prediction power of typical nightlight regressions down, suggesting that the nightlight-economic activity relationship is unstable. The prediction of Landsat instead is poor for very large communes. For practical purposes, the results suggest that Landsat is a good predictor if researchers are interested in studying formal growth in urban areas, while nightlights are superior for the purpose of measuring economic activity in rural areas. Regardless of which setting is more important for the practitioner, both products can be a useful complement to each other and, if available, should both be used increase prediction power for economic activity.

4. Evaluating Changes in the Time Series

The paper next examines the predictive power of satellite imagery in the time series and looks at changes in economic activity over time.

Econometric Setup

The study continues to use a simple linear prediction model and estimates differences in the economic outcome variables on changes in the remotely sensed satellite data. For that the estimating equation in (1) is differenced to arrive at

$$\Delta \log y_{i,t} = \alpha + \beta \ \Delta satellitedata_{i,t} + \nu_{i,t} \tag{2}$$

where the operator Δ denotes the change between t and $t{-}1$ and $\nu_{i,\,t} = \epsilon_{i,\,t} - \epsilon_{i,\,t-1}$ is the potentially serially correlated error term. The key questions of interest are whether the parameter β is stable between the cross-sectional regressions in section 3 and the analysis in the time series and whether the reduced variation in the satellite measure has enough predictive power for economic activity at the commune level.

The data contain the number of enterprises and employment for the years 2004 to 2012, and expenditure for the years 2004, 2008, 2010, and 2012. However, inconsistencies in the commune boundaries and changes in the sampling frame of the surveys lead to problems in the comparability of the data. The analysis therefore explores the predictive power for changes of different horizons. Potential serial correlation is dealt with by estimating equation (2) both in a panel framework as well as in long differences of different length. The analysis focuses on medium-run changes of enterprise counts and employment from 2004 to 2012. For the change in expenditure, results are reported for two time frames that correspond to the sampling frameworks of the VHLSS: a four-year period between 2004 and 2008 (based on the 1999 population census) and a shorter horizon from 2012 to 2014 (which both were based on the 2009 population census).

Empirical Results

Nightlights and Landsat Indices

Table 4 summarizes the results for the different specifications of equation (2) for the combined nighttime lights and the simple Landsat indices. In contrast to the cross-section regressions, the results show very low predictive power and parameter instability for predicting changes in the economic outcome variables of employment and enterprise counts. For the percentage change in enterprise counts for the whole period from 2004 to 2012 (column 1), the estimated elasticity with respect to nighttime lights is –0.0356 which is statistically significant. This finding differs vastly from the cross-sectional regression where the elasticity was positive. The employment regression in column (2) yields a similarly small SOL coefficient. This suggests that the enterprise growth is negatively correlated with increases in nighttime light output, which is counterintuitive. The analysis focuses now on the simple Landsat indices and assesses the changes in the mean NDBI and NDVI as predictors for changes in the economic outcome variables. The estimated coefficients for both measures are positive and highly significant. Thus an increase in urban areas and vegetation is associated with expansion of enterprises and employment at the commune level. Overall, the model performs poorly at capturing the overall variation over time and yields very little predictive power with R-squared measures of close to zero.

To control for serial correlation, a long-difference regression is estimated for the same period in columns (3) and (4). This yields positive coefficients on nighttime light intensity, although this parameter is imprecisely measured in the enterprise regression. Again there are strong positive effects for NDBI and a negative effect of NDVI in the enterprise regression, but an insignificantly positive effect of NDVI in the employment model. The sign changes in these regressions suggest the presence of large measurement error. Again, there is very low predictive power with R-squared statistics of at most 0.017. In summary, the time series models perform poorly, and the parameter instability suggests that the model cannot nail down the relationship between changes in satellite measures and economic activity well.

The expenditure data is limited by the lower sample size and the shorter time period of consistent data. The results indicate negative coefficients for the SOL measure, which are imprecisely estimated. While the coefficients on NDBI are largely similar in both the 2010–2012 and 2004–2008 regressions and estimated to be around 0.3, they are only statistically significant in the former. The NDVI coefficients are statistically not significant and fluctuate. Similar to the production measure regressions, the models for expenditure can explain virtually no part of the overall variation in the data.

Table 4. Regression Results for Changes over Time

	Panel 2004–2012		Long Differences 2004–2012		2010–2012	2004–2008
	Enterprises	Employment	Enterprises	Employment	Expenditure	Expenditure
Change in SOL	−0.0356***	−0.0410***	0.000496	0.114***	−0.00263	−0.00283
	(0.00311)	(0.00509)	(0.0128)	(0.0231)	(0.0108)	(0.0130)
Change in NDBI	0.492***	0.685***	1.049***	1.523***	0.304*	0.333
	(0.0432)	(0.0708)	(0.161)	(0.290)	(0.123)	(0.215)
Change in NDVI	0.218***	0.469***	−0.891***	0.390	0.106	−0.163
	(0.0243)	(0.0398)	(0.132)	(0.237)	(0.112)	(0.133)
Observations	28768	28768	5308	5307	2614	2324
Adjusted R^2	0.010	0.009	0.017	0.009	0.001	0.001

Source: Authors' calculations based on Enterprise Survey, VHLSS, Landsat, and DMSP-OLS data.
Notes: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. Standard errors are in parentheses.

In conclusion, the analysis finds large parameter instability not only between the cross-section and the time series regressions, but also within different periods of the time series. The conclusion is that regardless of the time period used, neither changes in nightlights nor Landsat indices individually can predict changes of economic outcome variables in a useful way. The results confirm the cross-sectional results that the NDBI measure is a strong correlate of economic activity, even though it explains only a marginal share of its variation over time.

Spectral Landsat Bands

This section examines the predictive power of the spectral Landsat bands over time. Table 5 summarizes the findings. This exercise restricts attention to the first two band moments that showed the highest predictive power in the cross-section. This table only reports estimates on coefficients that were significant at the 5 percent level in any of the specifications.

Table 5. Long Differences Regression

	Change in enterprises (1)	Change in Employment (2)	Change in Expenditure (3)
Change in average B1	0.0388	0.0741***	−0.00323
	(0.0272)	(0.0146)	(0.0126)
Change in average B3	−0.0897**	−0.0890***	0.0180
	(0.0333)	(0.0182)	(0.0173)
Change in average B4	−0.0512***	−0.0498***	0.00699
	(0.0154)	(0.00765)	(0.00672)
Change in average B5	0.109***	0.0717***	−0.0154
	(0.0192)	(0.0103)	(0.0106)
Change in average B7	−0.0996***	−0.0556***	0.0201
	(0.0232)	(0.0126)	(0.0144)
Change in average B8	0.0326	0.0349*	−0.0184
	(0.0359)	(0.0171)	(0.0159)
Change in standard deviation B2	−0.0937*	−0.0851***	−0.0180
	(0.0447)	(0.0247)	(0.0248)
Change in standard deviation B3	0.0781*	0.0901***	0.00771
	(0.0333)	(0.0183)	(0.0177)
Change in standard deviation B4	0.0178	0.0192***	0.0104
	(0.00978)	(0.00561)	(0.00535)
Change in standard deviation B5	−0.0412*	−0.0552***	−0.0224*
	(0.0190)	(0.0107)	(0.00999)
Change in standard deviation B7	0.0151	0.0320**	0.0271*
	(0.0215)	(0.0122)	(0.0115)
Change in standard deviation B8	0.0422**	0.000126	−0.00390
	(0.0143)	(0.00798)	(0.00827)
Observations	5794	5795	2677
Adjusted R^2	0.015	0.056	0.005

Source: Authors' calculations based on Enterprise Survey, VHLSS, Landsat, and DMSP-OLS data.
Notes: Table only reports coefficients that were significant in at least one specification. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. Standard errors are in parentheses.

Column (1) reports the results of a linear regression of the number of enterprises on changes in the raw band moments. The R-squared from this augmented regression is still very low at $R^2 = 0.015$ although several band moments are highly significant predictors. The employment regression yields largely similar results: All the coefficient estimates have the same sign as in the enterprise regression although the statistical significance varies. The prediction power is considerably higher for this variable, and the full Landsat model can explain 5.6 percent of the squared variation in log employment over time. This is, of course, still a much smaller share than in the cross-section. For expenditure, the model performs poorly and has barely any prediction power. Only two of the many coefficients are measured with some statistical precision while other estimates are often very different from the enterprise and employment regressions.

The analysis suggests that even when augmenting the Landsat models flexibly with moments of all spectral bands, the prediction power of the simple OLS model in the time series is very low. While Landsat performed well at capturing the spatial variation of economic activity in the cross-section, changes over time cannot be easily modeled using simple variables from either satellite measure. There are several possibilities that could lead to this result. Measurement error over time in either the economic or satellite data could increase the noise-to-signal ratio in the regression. Another speculative reason is that while land use patterns and economic activity are strongly correlated in the long term, land use change might lag behind short-term economic growth and decline. For example, an increase in production might only manifest itself in recognizable changes on the ground through construction of factories and roads after several years. Likewise, a potential asymmetric response of land use between growth and decline could not be properly captured in a linear model.

5. Discussion and Conclusion

This paper has introduced remotely sensed imagery data from the Landsat program and evaluated its usefulness for prediction of economic activity. Unlike the commonly used nightlight data, Landsat imagery comes at a much higher spatial resolution, is measured more frequently, and provides data on a multitude of spectral bands during daytime. While nighttime lights from the DMSP-OLS measures human activity in the form of light consumption, Landsat imagery captures a bigger picture of the earth shaped by both nature and man. Thus, there is potential for detecting a variety of relevant features on the ground that correlate with socioeconomic data.

Using small-scale economic data, the analysis shows that Landsat imagery can act as a strong predictor for counts of enterprises and employment as well as expenditure in the cross-section in the context of Vietnamese communes. Simple combinations of Landsat bands that were developed to detect urban areas and vegetation already have reasonable predictive power, while flexible combinations of band means and standard deviations can explain a large share of the differences in production and consumption measures between communes. Cross-validation exercises and LASSO regressions indicate strong parameter stability and do not suggest overfitting due to the large number of potential parameters.

Comparing the Landsat measures to the often-used nighttime light approach, the results indicate that Landsat outperforms nightlights at predicting economic outcomes in the cross-section aggregate. This is especially true in very small areas and in communes with a high share of formal employment. These units tend to be urban areas, and this confirms concerns that the DMSP-OLS nightlight data is too coarse and saturated to be applied on a very small geographic scale. However, nightlights perform reasonably well at larger commune sizes. This suggests that Landsat and nightlights are complements rather than substitutes for applied researchers interested in capturing the spatial variation of economic activity. Given their very different nature of data collection, including both satellite measures into prediction models can improve their precision. This is especially attractive as both nightlights and Landsat data are freely accessible.

When looking at changes over time, neither nightlights nor Landsat has much prediction power in the simple models. Although some coefficients of Landsat bands are significant in the regression, they can only explain a small share of the variation over time and the R-squared of the linear models is close to zero. Likewise, the results show large differences between the nightlight coefficients in the cross-section and time-series regressions. This casts doubt on the usefulness of simple Landsat or nightlight models for predicting economic growth at very small entities and confirms Bickenbach et al.'s (2016) findings of parameter instability in the SOL approach.

The analysis showed that the strength of such simple models lies in capturing stable features on the ground that correlate with economic activity but that they are unable to capture changes over time. This does not rule out that more sophisticated models and more advanced signal extraction can further improve the value of Landsat or nightlight data in the time series. This is an active research area that is constantly expanding the understanding of remote sensing tools for measuring economic growth.

References

Abrahams, A., C. Oram, and N. Lozano-Gracia. 2018. "Deblurring DMSP Nighttime Lights: A New Method Using Gaussian Filters and Frequencies of Illumination." *Remote Sensing of Environment* 210: 242–58.

Alesina, A., S. Michalopoulos, and E. Papaioannou. 2016. "Ethnic Inequality." *Journal of Political Economy* 124 (2): 428–88.

Babenko, B., J. Hersh, D. Newhouse, A. Ramakrishnan, and T. Swartz. 2017. "Poverty Mapping using Convolutional Neural Networks Trained on High and Medium Resolution Satellite Images, with an Application in Mexico." preprint arXiv:1711.06323.

Baugh, K., C. D. Elvidge, T. Ghosh, and D. Ziskin. 2010. "Development of a 2009 Stable Lights Product Using DM-SPOLS Data." *Proceedings of the Asia-Pacific Advanced Network* 30: 114–30.

Belloni, A., V. Chernozhukov, and C. Hansen. 2014. "High-Dimensional Methods and Inference on Structural and Treatment Effects." *Journal of Economic Perspectives* 28 (2): 29–50.

Bickenbach, F., E. Bode, P. Nunnenkamp, and M. Söder. 2016. "Night Lights and Regional GDP." *Review of World Economics* 152 (2): 425–47.

Bleakley, H., and J. Lin. 2012. "Portage and Path Dependence." *Quarterly Journal of Economics* 127 (2): 587–644.

Burchfield, M., H. G. Overman, D. Puga, and M. A. Turner. 2006. "Causes of Sprawl: A Portrait from Space." *Quarterly Journal of Economics* 121 (2): 587–633.

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): 8589–94.

Cling, J.-P., M. Razafindrakoto, and F. Roubaus. 2011. "The Informal Sector in Vietnam." Report of the International Labour Organization. Geneva, Switzerland.

Donaldson, D., and A. Storeygard. 2016. "The View from Above: Applications of Satellite Data in Economics." *Journal of Economic Perspectives* 30 (4): 171–98.

Engstrom, R., J. S. Hersh, and D. L. Newhouse. 2017. "Poverty from space: using high-resolution satellite imagery for estimating economic well-being (English)." Policy Research Working Paper No. WPS 8284, Washington, DC: World Bank Group

Goldblatt, R., W. You, G. Hanson, and A. K. Khandelwal. 2016. "Detecting the Boundaries of Urban Areas in India: A Dataset for Pixel-Based Image Classification in Google Earth Engine." *Remote Sensing* 8, (8), 634–61.

Hansen, H., and T. D. Le. 2013. "The Importance of Being Surveyed: The Representativeness and Impact of the Vietnam Household Living Standards Surveys." Unpublished Manuscript.

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

Jean, N., M. Burke, M. Xie, W. M. Davis, D. B. Lobell, and S. Ermon. 2016. "Combining Satellite Imagery and Machine Learning to Predict Poverty." *Science* 353 (6301): 790–94.

McCaig, B., and N. Pavcnik. 2015. "Informal Employment in a Growing and Globalizing Low-Income Country." *American Economic Review* 105 (5): 545–50.

Mellander, C., J. Lobo, K. Stolarick, and Z. Matheson. 2015. "Night-Time Light Data: A Good Proxy Measure for Economic Activity?" *PloS One* 10 (10): e0139779.

Perez, A., C. Yeh, G. Azzari, M. Burke, D. Lobell, and S. Ermon. 2017. “Poverty Prediction with Public Landsat 7 Satellite Imagery and Machine Learning.” preprint arXiv:1711.03654.

Pincus, J., and J. Sender. 2008. “Quantifying Poverty in Viet Nam: who Counts?” *Journal of Vietnamese Studies* 3 (1): 108–50.

Small, C., F. Pozzi, and C. D. Elvidge. 2005. “Spatial Analysis of Global Urban Extent from DMSP-OLS Night Lights.” *Remote Sensing of Environment* 96 (3): 277–91.

Storeygard, A. 2016. “Farther on Down the Road: Transport Costs, Trade and Urban Growth in Sub-Saharan Africa.” *Review of Economic Studies* 83 (3): 1263–95.

Tibshirani, R. 1996. “Regression Shrinkage and Selection via the Lasso.” *Journal of the Royal Statistical Society*. Series B (Methodological) 58 (1): 267–88.

Zha, Y., J. Gao, and S. Ni. 2003. “Use of Normalized Difference Built-Up Index in Automatically Mapping Urban Areas from TM Imagery.” *International Journal of Remote Sensing* 24 (3): 583–94.