DETECTING CORRUPTION FROM OUTER SPACE

LEONARDO BACCINI*, YUSHAN HU[†], AND BEN G. LI[‡]

ABSTRACT. Moral hazards in local governance give rise to both corruption and data misreporting. We find that the latter, which is relatively easy to uncover, can be used to detect the former. Our study, using data from Chinese prefectures between 1993 and 2013, demonstrates that the difference between the GDP growth reported by local officials and the GDP growth inferred from nightlight luminosity can predict local corruption. This approach illustrates the effectiveness of integrating remote-sensing technology, non-causal correlation, and economic reasoning to uncover bureaucratic anomalies that originate from similar mechanisms. (*JEL Codes: D73, O47*)

1. INTRODUCTION

The crime of corruption is inherently covert. Economists recently found that new technologies including neural networks (López-Iturriaga and Sanz, 2018), machine learning (Gallego et al., 2021), blockchain (UNODC, 2022), and IT audits (Chalendard et al., 2023) can help expose corruption. Such innovations do not pinpoint specific bureaucrats as suspects but instead identify the type of bureaucrats who are at high risk of corruption. This approach is advantageous because it helps conserve public anticorruption resources in two ways. Firstly, it aids in the prevention of corruption by identifying bureaucrats who are at high risk of engaging in corrupt activities. Successful prevention not only eliminates the economic distortion caused by corruption but also reduces the costs arising from investigations of corruption. Secondly, by pinpointing where corruption may be happening covertly, it assists in concentrating investigative efforts in places where such efforts will be most fruitful. The new anticorruption technologies are not a substitute for direct evidence, due process, or fair trials, but rather complement these essential processes to deliver justice more efficiently.

In this study, we use the discrepancy between the GDP figures presented by local bureaucrats and those inferred from nightlight luminosity to predict local corruption. Satellite-captured nightlight luminosity has been extensively used to develop economic indicators, following the pioneering work of Henderson, Storeygard, and Weil (2012) (hereafter, HSW). Their satellite technique is considered a reliable method to estimate regional economic performance when local officials are either incompetent or untruthful in statistical reporting.

Date: August 29, 2023.

^{*} McGill University.

[†] Renmin University of China.

[‡] University of Massachusetts.

Our findings reveal that uncovering GDP over-reporting with nightlight-based inference can help identify prefectures where corrupt bureaucrats hold office. In our view, the corruption of local bureaucrats is linked to GDP over-reporting through three possible channels. First, opportunistic bureaucrats may lie about both their work performance and financial interests in order to take advantage of the information asymmetry between them and their supervising authorities, a mechanism identified by Martínez (2022) specific to authoritarian regimes. Second, provided that their supervising authorities or the public know about their corruption, corrupt bureaucrats have incentives to present a fake picture of local prosperity in exchange for lighter sanctions. Third, corrupt bureaucrats, preoccupied with self-interests, may neglect their duty to ensure the accuracy of statistical reporting, such that their administrative divisions are more likely to have inaccurate GDP data (biased in either direction). Although these channels differ in the level of intentionality, they all stem from moral hazards that are associated with corrupt behaviors.

This study demonstrates how nightlight-based inference, when combined with economic reasoning, can detect bureaucratic anomalies. Nightlight-based inference has been used in the estimation of government expenditure and public goods provision (e.g., Bleakley and Lin, 2012; Michalopoulos and Papaioannou, 2014). Below we first describe our data, then present our results, and lastly discuss how the detection of corruption using nightlight luminosity contributes to the forensic toolkit available to economists.

2. Data and Results

Our study utilizes data from China, where corruption among government officials is widespread, and the accuracy of GDP reporting is frequently questioned. The use of data from China has two advantages in our context. First, unlike most countries where GDP reporting commences from the provincial or state level, in China, it starts from the prefecture level of government, as noted by Au and Henderson (2006). Second, China operates as a unitary state, where each layer of government adheres to the same anti-corruption laws, practices, and conviction standards as those above it. These two advantages make China a suitable setting for evaluating the effectiveness of the nightlight-based approach we propose for detecting corruption.

Our data span the years from 1993 to 2013. Over the 21-year period, the discrepancy (DISC) between the GDP reported by local bureaucrats and that inferred from nightlight luminosity is defined as:

(1)
$$DISC_{pt} \equiv y_{pt}^b - \tilde{y}_{pt},$$

where $y_{pt}^b = \ln GDP_{pt}^b - \ln GDP_{pt-1}^b$ is the imputed rate of change (i.e., growth rate) of the GDP figures reported by local (b)ureaucrats for prefecture-year duplet pt. The use of logarithmic terms ensures that GDP be comparable over prefectures and time. \tilde{y}_{pt} represents the "fair" GDP growth rate of the duplet, extrapolated from nightlight data,

(2)
$$\tilde{y}_{pt} = \psi x_{pt},$$

where x_{pt} is the duplet-level nightlight measure in logarithm. The data on nightlight luminosity were originally collected by the Defense Meteorological Satellite Program (DMSP) of the US Air Force and first used by HSW to estimate the GDP of developing countries. The upper map in Figure 1 displays the original DMSP-collected nightlight luminosity

 $\mathbf{2}$

over China for the year 2013, and the lower map the duplet-level average nightlight data for the southwestern half of China.¹



FIGURE 1. Nightlight Luminosity: Raw vs. Prefecture Averages

Notes: The upper map displays the original nightlight data over China for the year 2013. The lower map, using the southeastern half of China as an example, shows the prefecture-level averages of nightlight.

¹The demarcation used in the lower map is known as the Hu-line in Chinese geography: the line divides China's territory into two equi-area halves, with the northwestern (southeastern) half accounting stably for 5 percent (95 percent) of Chinese population over time.

The parameter ψ in equation (2) is the inverse elasticity of nightlight growth with respect to GDP growth. The elasticity of nightlight growth with respect to GDP growth is estimated to be between 1.0 and 1.7 (see HSW and Donaldson and Storeygard (2016)). Accordingly, we select ψ from [1/1.0, 1/1.1, 1/1.3, 1/1.5, 1/1.7] and extrapolate \tilde{y}_{pt} and thus $DISC_{pt}$. We use $\psi = 1$ as our benchmark, which stipulates a one-to-one elasticity between nightlight growth and GDP growth, such that $DISC_{pt}$ collapses to the arithmetic difference between the two growth rates and can be interpreted nonparametrically. The smaller ψ is, the less the constructed discrepancy relies on nightlight-based correction.² We intentionally utilize external estimates of ψ reported in the literature instead of conducting our own estimation of ψ , since the elasticity between nightlight and GDP specific to China or its prefectures may itself be influenced by corruption and thus cause endogeneity.

The corruption and appointment data are sourced from the Transparency International China (TIC), which is affiliated with Transparency International, an organization that gathers corruption-related information for academic and policy research purposes. TIC maintains a database of corruption investigations, prosecutions, and convictions of these officials, alongside demographic data such as gender, education, and birth year of prefecture-level head bureaucrats. In China, each prefecture has two head bureaucrats: a mayor who administrates the local government, and a communist party secretary who leads a board of party members overseeing and coordinating the operation of the local government. TIC keeps track of all such head bureaucrats in China (henceforth, *bureaucrats*) over years, marking those convicted of corruption as *corrupt* while reporting the same demographic characteristics for both the corrupt and the non-corrupt.

Micro-level association. We merge the aforementioned $DISC_{pt}$, namely the growth rate discrepancy between the GDP reported by local bureaucrats and the GDP inferred from nightlight data, with the tenure-level data of local bureaucrats through prefecture p and year t, across all prefectures and the years 1993–2013. To examine the association between corruption and GDP discrepancy, we specify the following regression:

(3) $\operatorname{Prob}[Corrupt_{ipt} = 1] = \beta \cdot DISC_{pt} + \mathbb{X}_i \Gamma + \mu_{pv} + \mu_t + \epsilon_{ipt},$

where $Corrupt_{ipt}$ represents whether bureaucrat *i* holds office at prefecture *p* in year *t* is a corrupt one (1) or not (0). β is the parameter of interest. X_i is a vector of individual characteristics including age, gender, education (college-educated or not), and office type (mayor or party secretary). μ_{pv} is province fixed effect, μ_t is year fixed effect, and ϵ_{ipt} is a classic error term clustered by prefecture.

The regression results from specification (3) are reported in Table 1.³ As shown, regardless of the ψ -parameter used, a positive and statistically significant correlation is found between corruption and GDP discrepancy. Notice that we are not seeking to establish a causal effect of $DISC_{ipt}$ on $Corrupt_{pt}$ —a causal effect as such is unlikely to exist—but rather to determine if $DISC_{ipt}$ can be used to predict $Corrupt_{pt}$. Our results

²When ψ is set to zero, $DISC_{pt}$ is equal to the value reported by local bureaucrats y_{nt}^{b} .

³Descriptive statistics for the regression sample are provided in Table A1. A linear probability model is used here, whereas our results from a logit model are similar (available upon request). We prefer the linear probability model because it can process two sets of fixed effects with far less computational resources.

	(1)	(2)	(3)	(4)	(5)
DISC (assuming $\psi = 1/1.0$)	0.0347**				
	(0.0175)				
DISC (assuming $\psi = 1/1.1$)		0.0391**			
		(0.0191)			
DISC (assuming $\psi = 1/1.3$)			0.0472**		
			(0.0222)		
DISC (assuming $\psi = 1/1.5$)				0.0543**	
				(0.0248)	
DISC (assuming $\psi = 1/1.7$)					0.0603**
					(0.0271)
Observations	12,393	12,393	12,393	12,393	12,393
R-squared	0.0414	0.0414	0.0414	0.0415	0.0415

TABLE 1. Micro-level Association between Corruption and GDP Discrepancy

Dependent variable is *Corrupt* (1 or 0). Control variables include age (at the beginning of first citylevel tenure), gender, education level, and position (mayor or party secretary). To save space, only the coefficient of *DISP* is reported. Linear probability model is used to accommodate both province and year fixed effects. Robust standard errors are reported in parentheses, clustered by prefecture. ** p<0.05.

confirm a positive association between them. Such individual-level correlation serves as a micro-foundation for the macro-level prediction we conduct next. Using DISC on a regional or national scale to detect corruption is our ultimate goal. Put differently, in the absence of individual-level corruption data, users of our approach can make predictions about the areas where corruption is most likely to be hidden.

Macro-level detection. Resting on the micro-level specification (3), our macro-level regression specification is

(4)
$$\sum_{it \in p} \mathbb{I}[Corrupt_{ipt} \ge 1] = \delta \cdot \sum_{t \in p} \mathbb{I}[DISC_{pt} > a_t] + \epsilon_p,$$

where δ is the parameter of interest. Specifically, the probabilistic corrupt bureaucrats associated with prefecture p in specification (3) are aggregated in specification (4) into a count of corrupt bureaucrats associated with the prefecture. Since the *DISC* values associated with prefecture p are not additive, the incidents of large $DISC_{pt}$ (defined as being greater than a given percentile a_t in year t) are used instead as the predictor. Characteristics of bureaucrats are excluded rather than aggregated in order to reduce data demand.⁴

⁴By assuming (i) μ_{pv} , μ_t , and ϵ_{ipt} stochastic, independently distributed, and with a mean of zero and (ii) each bureaucrat holding only one prefecture-level tenure, one can obtain specification (4) by aggregating equation (3) conditional on $(X_1, X_2, ..., X_n)$ (here, *n* is the total number of bureaucrats).

(1) (2) (3) (4) (5) (6) (7) (8) (9) (10)(11) Dep variable is count of local corrupt bureaucrats Incidents of high *DISC* (> a = 90th percentile) 0.065*** (0.019)0.065*** Incidents of high *DISC* (> a = 80th percentile) (0.019)Incidents of high *DISC* (> a = 70th percentile) 0.066*** (0.019)0.093*** Incidents of high *DISC* (> a = 60th percentile) (0.022)Incidents of high *DISC* (> a = 50th percentile) 0.071*** (0.027)Incidents of high *DISC* (> a = 40th percentile) 0.067** (0.030)Incidents of high *DISC* (> a = 30th percentile) 0.075** (0.033)Incidents of high *DISC* (> a = 20th percentile) 0.089** (0.037)Incidents of high *DISC* (> a = 10th percentile) 0.067 (0.051)DISC (mean) 3.688** (1.738)Reported GDP growth rate (mean) 2.384 (2.290)Observations 278 278 278 278 278 278 278 278 278 278 278 0.016 R-squared 0.036 0.037 0.035 0.055 0.024 0.017 0.016 0.007 0.013 0.004

TABLE 2.	Macro-level	Association	between	Corruption	and (GDP	Discrepancy
				0 0 0 0 0 - 0			/

Notes: Each observation corresponds to one prefecture. *DISC* is constructed using $\psi = 1$. In columns (1) to (9), an incident of high *DISC* is defined as *DISC*_{pt} being greater than the a-th percentile, where a=90, 80, ..., 10. The larger a is, the stricter is the criterion used for designating high-*DISC* incidents. In column (11), GDP growth rate imputed from reported GDP figures is used as the predictor. Robust standard errors in parentheses. *** p<0.01, ** p<0.05.

LEONARDO BACCINI^{*}, YUSHAN HU^{\dagger} , AND BEN G.

 \mathbf{LI}^{\ddagger}



FIGURE 2. Incidents of High DESC vs. Occurrence of Corruption

Notes: The upper map displays prefecture-level high-*DISC* incidents (high-*DISC* is defined as being above the 50th percentile of the corresponding year): dark gray represents incidents 10 or more, while light gray represents incidents 4 to 9. The lower map displays in black prefectures with at least one corrupt bureaucrat.

The regression results from specification (4) are reported in Table 2. In columns (1) to (9), percentiles from $(a_t =)$ 90th to 10th are respectively used to designate high discrepancy. The lower the percentile that is used, the less likely $DISC_{pt} > a_t$ is associated with GDP-overreporting and thus the less its incidence can predict corruption. As shown, the statistical significance decreases as lower percentiles are used. In column (10), we use average DISC across years as the predictor and find it to have significant predictive power as well. Notice that high reported GDP growth alone (i.e., y_{pt}^b) does not predict corruption (column (11)), as it may reflect true economic growth. In other words, it is not high economic growth itself but data manipulation that covaries with corruption.

The regressor $\sum_{t \in p} \mathbb{I}[DISC_{pt} > a_t]$ in specification (4) can be used in a nonparametric manner. In Figure 2, we compare the map of areas with high $\sum_{t \in p} \mathbb{I}[DISC_{pt} > a_t]$ (with a_t set to the 50th percentile) to the map of identified corruption (i.e., at least one corrupt bureaucrat) during our 21-year sample period. The geographical distributions of the two maps overlap extensively, suggesting a high rate of successful detection.

3. CONCLUSION AND DISCUSSION

In this study, we propose using the discrepancy between the GDP growth reported by local bureaucrats and the GDP growth inferred from nightlight luminosity to detect local corruption. The correlation between GDP-overreporting and corruption receives support from micro-level regression analysis and can be rationalized with three incentive channels all featuring moral hazards in local governance. This approach to detecting corruption, as we illustrate using the case of China, is effective and not data-demanding. Currently, high quality and large quantities of remote-sensing data are becoming available, along with advanced statistical learning techniques. Applying these data and techniques to the detection of corruption can conserve public resources, since redflagging places where corruption may be hidden helps with both prevention of corruption and investigation of unprevented corruption.

This study showcases how studying corruption at the micro-level may help combat corruption at the macro-level. There is a growing body of literature that employs various novel data to test illegal and unethical behaviors predicted by micro-level incentive models, known as "forensic economics" (Fisman and Miguel, 2008; Zitzewitz, 2012). A solid understanding of micro-level behaviors is essential for making effective macro-level public policies. In our view, micro-level corruption studies microfound macro-level anticorruption policies by providing factual bases and operable metrics. By leveraging them, forensic economics can be expanded to deliver anticorruption applications on a larger scale, similar to how modern macroeconomics uses micro-level economic behaviors to inform macro-level economic policies. Such anticorruption applications can incorporate not only the existing forensic economics studies, but also methods and parameters from outside the field, such as the nightlight-GDP elasticities estimated by Henderson, Storeygard, and Weil (2012) utilized in this study.

References

Au, Chun-Chung and J. Vernon Henderson (2006), "Are Chinese Cities Too Small?" Review of Economic Studies, 73, 549–576.

- Bleakley, Hoyt and Jeffrey Lin (2012), "Portage and Path Dependence." Quarterly Journal of Economics, 127, 587-644.
- Chalendard, Cyril, Ana M. Fernandes, Gael Raballand, and Bob Rijkers (2023), "Corruption in Customs." *Quarterly Journal of Economics*, 138, 575–636.
- Donaldson, Dave and Adam Storeygard (2016), "The View from Above: Applications of Satellite Data in Economics." *Journal of Economic Perspectives*, 30, 171–198.
- Fisman, Raymond and Edward Miguel (2008), Economic Gangsters: Corruption, Violence, and the Poverty of Nations. Princeton University Press.
- Gallego, Jorge, Gonzalo Rivero, and Juan Martínez (2021), "Preventing rather than Punishing: An Early Warning Model of Malfeasance in Public Procurement." International Journal of Forecasting, 37, 360–377.
- Henderson, J. Vernon, Adam Storeygard, and David N. Weil (2012), "Measuring Economic Growth from Outer Space." American Economic Review, 102, 994–1028.
- López-Iturriaga, Félix J. and Iván Pastor Sanz (2018), "Predicting Public Corruption with Neural Networks: An Analysis of Spanish Provinces." *Social Indicators Research*, 140, 975–998.
- Martínez, Luis R. (2022), "How Much Should We Trust the Dictator's GDP Growth Estimates?" Journal of Political Economy, 130, 2731–2769.
- Michalopoulos, Stelios and Elias Papaioannou (2014), "National Institutions and Subnational Development in Africa." *Quarterly Journal of Economics*, 129, 151–213.
- UNODC (2022), "Detecting and Investigating Corruption." Published by the United Nations Office on Drugs and Crime, available at https://www.unodc.org/e4j/en/anti-corruption/ module-6/key-issues/detection-mechanisms---auditing-and-reporting.html.
- Zitzewitz, Eric (2012), "Forensic Economics." Journal of Economic Literature, 50, 731-769.

Online Appendix

TABLE A1. Descriptive Statistics for the Micro-level Sample

Variable	Obs	Mean	Std. dev.	Min	Max
<i>Corrupt</i> (corrupt bureaucrat=1)	12,393	0.11	0.32	0	1
DISC (assuming $\psi = 1/1.0$)	12,393	0.0950	0.1852	-1.3788	1.0169
DISC (assuming $\psi = 1/1.1$)	12,393	0.0998	0.1737	-1.3744	1.0149
DISC (assuming $\psi = 1/1.3$)	12,393	0.1071	0.1570	-1.3675	1.0117
DISC (assuming $\psi = 1/1.5$)	12,393	0.1125	0.1457	-1.3626	1.0094
DISC (assuming $\psi = 1/1.7$)	12,393	0.1166	0.1377	-1.3587	1.0077
Age	12,393	50.62	4.37	26	66
Sex (female=1)	12,393	0.04	0.20	0	1
Education (colleged educated=1)	12,393	0.89	0.31	0	1
Position (party secretary=1)	12,393	0.47	0.50	0	1

Notes: Descriptive statistics of the sample used in Table 1 are reported. Each observation is concerned with a bureaucrat at a given prefecture in a given year.