# Measuring Economic Growth with Many Signals

Yingyao Hu *            Jiaxiong Yao †

June 29, 2026

**Abstract**

Measuring economic activity with many signals has become possible in the age of big data. We provide a general framework for measuring a latent variable with many signals, extending existing two- and three-signal models to $K$ signals, with closed-form MSE-optimal and BLUE combinations and convergence results as $K \to \infty$. We also show nonlinear transformations of observables can generate additional informative measurements, proving MSE convergence to the oracle bound via an orthonormal polynomial sieve. Applied to GDP growth measurement, satellite signals prove more informative over longer horizons, and China, India, and Indonesia's official growth statistics appear systematically biased or oversmoothed.

JEL classification: *C01, E01, O47*

Keywords: *Measurement error model, many signals, GDP growth, nighttime lights, signal extraction*

*Department of Economics, Johns Hopkins University. Email: yhu@jhu.edu.

†International Monetary Fund, 700 19th St NW, Washington, DC 20431. Email: jyao@imf.org. The views expressed in this paper are those of the authors and do not necessarily reflect those of the International Monetary Fund, its Board of Directors or the countries they represent.

# 1 Introduction

The proliferation of nontraditional data has dramatically expanded the menu of measurements available for tracking economic activity. Remote sensing now provides a rich set of independent signals correlated with a country's true GDP growth, including satellite-recorded nighttime luminosity, nitrogen dioxide concentrations, urban land cover, greenness, and built-up areas. Beyond remote sensing, an even broader array of indicators captures different facets of economic activity: greenhouse gas emissions, electricity consumption, automated ship-tracking, card transactions, online search and other digital footprints, mobile-phone records, and text- and news-based sentiment indices, among numerous others. The theoretical and empirical challenge is no longer one of data scarcity but of data integration: how should these many noisy and partially overlapping signals be combined, alongside official GDP statistics, to recover the best possible estimate of true economic growth?

The seminal paper of Henderson, Storeygard, and Weil (2012) (HSW henceforth) introduced a two-signal model that combines official GDP growth with nighttime luminosity to estimate true GDP growth. However, the model is under-identified and HSW must calibrate the signal-to-noise ratio of official GDP for advanced economies. Civelli, Gaduh, and Yousuf (2024) (CGY henceforth) resolve this by adding a third signal, urban land cover, which provides additional moment conditions to exactly identify the model.

In this paper, we develop a general framework for measuring a latent variable with $K \geq 3$ signals, generalizing the HSW and CGY frameworks along two dimensions—from three signals to an arbitrary $K$ and incorporating nonlinearities through additionally generated measurements—and apply the framework to estimating economic growth.

First, we prove identification under fairly mild assumptions about measurement errors of the $K$ signals. We then provide closed-form solutions to optimal linear combinations of the measurements to approximate the latent variable. Distinguishing between the mean squared error (MSE) optimal linear combination and the best linear unbiased estimator (BLUE) combination, we show that the CGY's three-signal model is a special case of the MSE-optimal linear combination. We establish convergence results as $K \to \infty$, characterizing when additional signals help and when they do not.

Second, we incorporate nonlinearities in signals. By constructing nonlinear functions of the observables, we generate additional measurements that extract the nonlinearities. We prove that a polynomial sieve of generated measurements converges to the oracle bound—the minimum MSE achievable from observables—and characterize the fundamental information limits of generated versus independent measurements. Generated measurements that capture nonlinearity and independent signals play complementary roles. Intuitively, independent signals could drive the oracle bound toward zero because each independent signal adds new information not spanned by the existing measurements. Generated measurements close the gap between the linear-only MSE and the oracle bound by extracting nonlinear information that linear projections miss, but the total information content of generated measurements is bounded.

We apply this general framework to estimating true economic growth using official GDP growth, a range of satellite data, and traditional measurements. We consider two specifications that differ in the signals they admit and the error structure they allow. The first pairs official GDP with five satellite signals alone—nighttime lights, $NO_2$, urban land cover, NDVI, and NDBI—whose measurement errors are plausibly uncorrelated with that of official GDP. This clean-error structure isolates how much each satellite signal contributes on its own and how its usefulness scales with the time horizon. The second specification combines official GDP with two satellite signals (nighttime lights and $NO_2$) and two traditional energy statistics ($CO_2$ emissions and electricity consumption), whose measurement errors are allowed to correlate with that of official GDP. This specification asks how far traditional measurements can lower the estimation error once their correlated-error structure is modeled explicitly. In each specification, the estimated signal-to-noise ratios allow us to determine the informativeness of each signal, while the optimal weights demonstrate the overall usefulness of each signal and their nonlinear content.

Our paper makes several contributions relative to the existing literature. First, we extend the literature on measuring economic growth with satellite data. Most of the existing literature focuses on the under-identification problem in the two-signal model of HSW, which requires additional assumption or information. Hu and Yao (2022) imposes additional structures of the measurement errors in official GDP growth and nighttime lights while CGY adds an additional signal. Beyer, Hu, and Yao (2022) leverages heteroskedasticity of measurement errors for identification. In this paper, we build on CGY's three-signal result and generalize it to a $K$-signal model ($K \geq 3$), providing a new framework to analyze an arbitrarily large set of signals.

Second, the measurement error literature focuses on identification of the distribution of the latent variable $x^*$ and the relationship between the measurements and the latent variable, described by the conditional distribution $f(x_k|x^*)$. For example, Kotlarski (1965) and Rao (1992) show that two measurements with classical measurement errors can identify the distribution of the latent variable and the conditional distribution $f(x_k|x^*)$. This result is the crucial step in the identification of the nonlinear regression models with measurement errors in the style of Li (2002) and Schennach (2004). The seminal works in Hu (2008) and Hu and Schennach (2008) show that these distributions of interest can be identified from three measurements when they are independent conditional on the latent variable.[1] This paper takes a different approach. Instead of identifying the full distribution of $x^*$, we provide an estimate of the true value in each observation through an optimal linear combination of signals. Under the sufficient conditions for convergence, this identification is at the observational level instead of at the distributional level.

Third, our results are indirectly related to the many or weak instruments literature (Staiger and Stock (1997) and Stock and Wright (2000)). If we consider $x_1$ as the dependent variable and $x_2$ as an endogenous regressor in our setup, the additional measurements can be considered as instruments when the variance-covariance matrix of the errors is diagonal. In that case, the larger the coefficients $\gamma$ are, the stronger the correlation between the endogenous variable and the instruments is. This paper focuses on approximation to the latent variable, while the instrument variable literature usually focuses on the asymptotic properties of the estimator of the structural parameter. The setting in this paper can also be considered as a generalization of the principal component analysis (Bai and Ng (2002) and Stock and Watson (2002)). We focus on a single latent factor but with a general correlation among the error terms with a cross-sectional sample, while

[1]We refer to Hu (2017) and Schennach (2019) for a recent survey.

factor models focus more on the number of latent factors and the model prediction with a large $N$ and large $T$ panel.

In the empirical application, a central finding is that the signal-to-noise ratio of satellite data rises substantially as the horizon lengthens from annual frequency to long differences. At annual frequency, measurement noise dominates the satellite signals, so the optimal combination assigns most weight to official GDP as the least-noisy measure; among the satellite signals, nighttime lights are by far the most informative, with $NO_2$ and urban land cover contributing far less. Over a fifteen-year long difference, however, measurement noise washes out in long-run averages while the cumulative latent signal grows with the horizon, and the satellite signals overtake official GDP. Here urban land cover becomes the most informative signal, followed by nighttime lights. The reversal reflects what each signal captures: nighttime lights record both the spatial extent and the intensive-margin brightness of activity, making them responsive to annual fluctuations, whereas urban land cover tracks only the slow structural expansion of built-up area, which is largely invisible year to year but accumulates into the dominant signal over long horizons.

Combining satellite data with traditional measurements lowers the estimation error substantially and produces our best estimates. Admitting $\mathrm{CO}_2$ emissions and electricity consumption alongside nighttime lights and $NO_2$ reduces the minimum mean squared error of the annual panel by roughly four-fifths relative to the satellite-only specification. The two types of measurements play complementary roles: the satellite signals supply the orthogonality conditions that identify the structural parameters, while the energy statistics—mechanically tied to year-on-year economic activity—carry the richest predictive content, with signal-to-noise ratios an order of magnitude above the satellite signals at annual frequency. Adding nonlinear transformations of the signals yields only marginal further gains, concentrated almost entirely in the squared nighttime light term.

Applying the optimal combination to individual countries reveals systematic discrepancies between official and estimated growth for some countries. China and India's estimated growth lies systematically below official statistics, consistent with documented over-reporting in fast-growing emerging economies, while Brazil's estimated growth lies above its official figures, as its emissions and electricity use grew faster than its reported output would imply. The estimates also expose excessive smoothness in some official series: Indonesia's official growth is the smoothest in our sample, yet the satellite and energy signals imply year-on-year volatility more than double the official figure. These findings are consistent with the previous literature Nakamura, Steinsson, and Liu (2016), Martinez (2022), Hu and Yao (2022).

As an external check on these country-level discrepancies, we compare the model's disagreement with official GDP to the revisions that countries subsequently make to their own growth figures across vintages of the IMF World Economic Outlook, which provide an independent, model-free proxy for the reliability of official statistics. The magnitude of the model's correction lines up closely with the magnitude of these revisions: the countries whose official growth we depart from most are also those whose figures are revised most. The relationship holds for magnitude but not direction, since vintage revisions capture how much official growth is later corrected rather than whether it was originally over- or under-stated.

The rest of the paper is organized as follows. Section 2 sets up the many-measurement model, proves its identification, provides the optimal linear approximations, and discusses scale invariance and signal-to-noise ratios. Section 3 establishes convergence results and estimation theory. Section 4 extends the framework to incorporate nonlinear measurements, proves convergence to the oracle bound, and characterizes the complementarity between generated and independent measurements. Section 5 presents simulations demonstrating the value of nonlinear measurements and many independent signals. Section 6 applies the framework to measuring economic growth, comparing two specifications—one pairing official GDP with satellite signals alone and the other adding traditional energy statistics with correlated measurement errors—and extends each with a polynomial sieve that captures nonlinearity. Section 7 concludes.

# 2 Many-Measurement Model

## 2.1 The Model

We consider a cross-sectional i.i.d. sample, which includes observations of a vector of measurements $x = (x_1, x_2, ..., x_K)^T$ with $K \geq 3$. These observables are generally-defined measurements of a scalar latent variable $x^*$ satisfying:

$$x_k \quad = \quad \gamma_k x^* + \varepsilon_k, \quad cov(x^*, \varepsilon_k) = 0, \quad k = 1, 2, ..., K. \tag{1}$$

The measurement error, defined as $x_k - x^*$, is classical if $\gamma_k = 1$ and the error term $\varepsilon_k$ is independent of the true value $x^*$. The uncorrelatedness in Equation (1) is not restrictive if we consider Equation (1) as linear projection equations under the condition that the observables have finite variances. $\varepsilon_k$ is allowed to have non-zero mean and may be correlated with each other.

## 2.2 Identification

Our analysis requires that some of the measurements should contain information about the latent variable so that some $\gamma$'s need to be nonzero. Since $x^*$ is not observed anywhere, we also need normalization conditions. We assume as follows:

**Assumption 1.** *(Three measurements) There are three measurements containing information on $x^*$ and mutually uncorrelated measurement errors, i.e., $\gamma_1 \neq 0$, $\gamma_2 \neq 0$, $\gamma_3 \neq 0$ and*

$$cov(\varepsilon_1, \varepsilon_2) = cov(\varepsilon_1, \varepsilon_3) = cov(\varepsilon_2, \varepsilon_3) = 0.$$

*In addition, one of the three measurement errors is uncorrelated with all other measurement errors. Without loss of generality, let $\varepsilon_1$ satisfy*

$$cov(\varepsilon_1, \varepsilon_k) \quad = \quad 0, \quad k = 4, ..., K.$$

Identification of $\gamma_k$ requires additional restrictions:

**Assumption 2.** *1) (Normalization) One of $\gamma_1$, $\gamma_2$ and $\gamma_3$ is known, and one of $E[\varepsilon_1]$, $E[\varepsilon_2]$, and $E[\varepsilon_3]$ is known. Without loss of generality, let $\gamma_1 = 1$ and $E[\varepsilon_1] = 0$.*
*2) (Full rank) The variance-covariance matrix $V(\varepsilon)$ has full rank.*

Assumptions 1 and 2.1 provide enough moments to solve for all the $\gamma_k$ from the second order moments given that there are $K$ uncorrelation restrictions and $K$ unknowns, including $(\gamma_2, ..., \gamma_K)$ and $var(x^*)$. Note that $\gamma$ and $\sigma^2_{x^*}$ can be solved from the $K$ restrictions as follows:

$$\begin{aligned}
\gamma_2 &= \frac{cov(x_2, x_3)}{cov(x_1, x_3)} \\
\gamma_k &= \frac{cov(x_k, x_1)}{\sigma^2_{x^*}}, \quad k = 3, ..., K. \\
\sigma^2_{x^*} &= \frac{cov(x_1, x_2)cov(x_1, x_3)}{cov(x_2, x_3)}.
\end{aligned}$$

Given the normalization of one error mean ($E[\varepsilon_1] = 0$), all the first order moments can be estimated given $\gamma_k$ as $E(\varepsilon_k) = E(x_k) - \gamma_k E(x^*) = E(x_k) - \gamma_k E(x_1)$.

Assumption 2.2 rules out redundant measurements. Notice that if one wants to add an additional measurement $X_{K+1}$ into the model, the only requirement is that its error term has to be uncorrelated with the error in the first measurement, i.e., $cov(\varepsilon_1, \varepsilon_{K+1}) = 0$. Other than that, the additional measurement can be correlated with other measurements or be totally irrelevant to the latent variable of interest $x^*$.

### 2.3 Optimal linear combinations

We propose two linear combinations of the measurements to approximate the latent $x^*$ that seek to minimize the mean squared error.[2] The first is the MSE-optimal combination which allows for tradeoff between bias and variance, and the second is the best linear unbiased estimator (BLUE) combination.

Consider a linear combination that combines all measures of economic activity,

$$\hat{x}^*(\lambda) \equiv \lambda_0 + \lambda_1^T x.$$

The MSE-optimal linear combination minimizes the mean squared error over all $\lambda = (\lambda_0, \lambda_1)$:

$$\lambda^{\text{MSE}} = \arg\min_{\lambda} E[\hat{x}^*(\lambda) - x^*]^2.$$

The BLUE combination requires that the linear approximation has the conditional mean equal to the latent variable, i.e., $E[\hat{x}^*(\lambda)|x^*] = x^*$, if the errors are conditional mean independent of $x^*$, i.e., $E[\varepsilon|x^*] = E[\varepsilon]$. In other words, the linear approximation is conditionally unbiased. This implies that $\lambda$ must satisfy[3]

$$\lambda \in \Lambda := \{\lambda : \lambda_1^T \gamma = 1; \lambda_0 = E[x^*] - \lambda_1^T E[x]\}.$$

[2] We use the terms "optimal linear combination" and "predictor" interchangeably throughout the paper, both referring to a linear combination of the signals that approximates the latent variable.

[3] See Supplemental Appendix.

The resulting BLUE combination is

$$\lambda^{\text{BLUE}} \quad = \quad \arg\min_{\lambda\in\Lambda} E[\hat{x}^*(\lambda) - x^*]^2.$$

Both minimization problems lead to closed-form solutions. We summarize the results as follows:

**Proposition 1.** *If Assumptions 1 and 2 hold, then:*

(i) MSE-optimal (bias-allowing). *The linear combination minimizing $E[\hat{x}^*(\lambda) - x^*]^2$ over all $\lambda = (\lambda_0, \lambda_1)$ is*

$$\hat{x}^*(\lambda^{\text{MSE}}) \quad = \quad \frac{\sigma_{x^*}^2}{1 + \sigma_{x^*}^2\, \gamma^T V^{-1}\gamma}\, \gamma^T V^{-1}(x - E[x]) + E[x_1], \tag{2}$$

*with the minimum MSE*

$$E[\hat{x}^*(\lambda^{\text{MSE}}) - x^*]^2 \quad = \quad \frac{\sigma_{x^*}^2}{1 + \sigma_{x^*}^2\, \gamma^T V^{-1}\gamma}$$

*and the optimal weights*

$$\begin{aligned} \lambda_1^{\text{MSE}} \quad &= \quad \frac{\sigma_{x^*}^2\, V^{-1}\gamma}{1 + \sigma_{x^*}^2 \gamma^T V^{-1}\gamma}, \\ \lambda_0^{\text{MSE}} \quad &= \quad -\,(\lambda_1^{\text{MSE}})^T\, E[x] + E[x_1] \end{aligned}$$

(ii) BLUE (unbiased). *The linear combination minimizing $E[\hat{x}^*(\lambda) - x^*]^2$ over $\lambda \in \Lambda$ is*

$$\hat{x}^*(\lambda^{\text{BLUE}}) \quad = \quad \frac{1}{\gamma^T V^{-1}\gamma}\gamma^T V^{-1}(x - E[x]) + E[x_1] \tag{3}$$

*with the minimum MSE*

$$E[\hat{x}^*(\lambda^{\text{BLUE}}) - x^*]^2 \quad = \quad \frac{1}{\gamma^T V^{-1}\gamma}$$

*and the optimal weights*

$$\begin{aligned} \lambda_1^{\text{BLUE}} \quad &= \quad \frac{1}{\gamma^T V^{-1}\gamma} V^{-1}\gamma, \\ \lambda_0^{\text{BLUE}} \quad &= \quad -\,(\lambda_1^{\text{BLUE}})^T\, E[x] + E[x_1]. \end{aligned}$$

*In both cases, the error variance matrix is*

$$V = \begin{pmatrix} \sigma_1^2 & 0 & 0 & 0 & \dots & 0 \\ 0 & \sigma_2^2 & 0 & \sigma_{24} & \dots & \sigma_{2K} \\ 0 & 0 & \sigma_3^2 & \sigma_{34} & \dots & \sigma_{3K} \\ 0 & \sigma_{42} & \sigma_{43} & \sigma_4^2 & \dots & \sigma_{4K} \\ \vdots & \vdots & \vdots & \vdots & \dots & \vdots \\ 0 & \sigma_{K2} & \sigma_{K3} & \sigma_{K4} & \dots & \sigma_K^2 \end{pmatrix},$$

*where*

$$\begin{aligned} \sigma_{jk} &= cov(x_j, x_k) - \gamma_j \gamma_k \sigma_{x^*}^2 \\ \sigma_k^2 &= \sigma_{kk} \\ \gamma_2 &= \frac{cov(x_2, x_3)}{cov(x_1, x_3)} \\ \gamma_k &= \frac{cov(x_k, x_1)}{\sigma_{x^*}^2} \quad k = 3, \dots, K \\ \sigma_{x^*}^2 &= \frac{cov(x_1, x_2) cov(x_1, x_3)}{cov(x_2, x_3)}. \end{aligned}$$

**Proof**: See Supplemental Appendix.

This proposition provides closed-form solutions for both the MSE-optimal and the BLUE linear approximations of the latent variable. Both can be directly estimated using the first and second order moments of observed variables. The MSE-optimal linear combination dominates the BLUE in mean squared error because it allows a small amount of bias in exchange for a larger reduction in variance:

$$\begin{aligned} E[\hat{x}^*(\lambda^{\text{MSE}}) - x^*]^2 &= \frac{\sigma_{x^*}^2}{1 + \sigma_{x^*}^2 \gamma^T V^{-1} \gamma} = \frac{1}{1/\sigma_{x^*}^2 + \gamma^T V^{-1} \gamma} \\ &\leq \frac{1}{\gamma^T V^{-1} \gamma} = E[\hat{x}^*(\lambda^{\text{BLUE}}) - x^*]^2, \end{aligned} \tag{4}$$

with equality in the limit $\sigma_{x^*}^2 \to \infty$. The gap $\sigma_{x^*}^2 / [\gamma^T V^{-1} \gamma (1 + \sigma_{x^*}^2 \gamma^T V^{-1} \gamma)]$ is small when the signals are informative ($\gamma^T V^{-1} \gamma$ is large) or when the latent variance is large, and it is largest when signals are noisy relative to $\sigma_{x^*}^2$. The BLUE remains useful when conditional unbiasedness is required for interpretation or for downstream inference.

Two remarks are in order.

**Remark 1:** *CGY Three-Signal Formula as a Special Case.*
Civelli, Gaduh, and Yousuf (2024) propose a three-signal model with independent errors. It is a special case in Proposition 1 with $K = 3$ and diagonal $V = \text{diag}(\sigma_1^2, \sigma_2^2, \sigma_3^2)$. The closed-form

MSE-optimal weight $\lambda_1^{\text{MSE}} = \frac{\sigma_{x^*}^2 V^{-1}\gamma}{1+\sigma_{x^*}^2\gamma^T V^{-1}\gamma}$. The first entry with diagonal $V$ and $\gamma_1 = 1$ is

$$\lambda_{1,1}^{\text{MSE}} = \frac{\sigma_{x^*}^2 (V^{-1}\gamma)_1}{1+\sigma_{x^*}^2\gamma^T V^{-1}\gamma} = \frac{\sigma_{x^*}^2/\sigma_1^2}{1+\sigma_{x^*}^2\left(1/\sigma_1^2+\gamma_2^2/\sigma_2^2+\gamma_3^2/\sigma_3^2\right)} = \frac{1}{1+\dfrac{\sigma_1^2}{\sigma_{x^*}^2}+\dfrac{\sigma_1^2}{\sigma_2^2}\gamma_2^2+\dfrac{\sigma_1^2}{\sigma_3^2}\gamma_3^2},$$

preciously what CGY has derived.[4] Proposition 1's BLUE is its unbiased counterpart:

$$\lambda_{1,1}^{\text{BLUE}} = \frac{(V^{-1}\gamma)_1}{\gamma^T V^{-1}\gamma} = \frac{1/\sigma_1^2}{1/\sigma_1^2+\gamma_2^2/\sigma_2^2+\gamma_3^2/\sigma_3^2} = \frac{1}{1+\dfrac{\sigma_1^2}{\sigma_2^2}\gamma_2^2+\dfrac{\sigma_1^2}{\sigma_3^2}\gamma_3^2}.$$

The MSE-optimal and BLUE combinations trade off bias and variance via the term $\sigma_1^2/\sigma_{x^*}^2$, and their weights on the first measurement coincide as $\sigma_{x^*}^2 \to \infty$.

**Remark 2:** *Scale Invariance and Signal-to-Noise Ratio.*
In the case where $V$ is diagonal, the optimal weight on measurement $x_k$ is $\lambda_{1,k}^{\text{BLUE}} = \frac{\gamma_k/\sigma_k^2}{\sum_j \gamma_j^2/\sigma_j^2}$ and the minimum MSE is $(\sum_k \gamma_k^2/\sigma_k^2)^{-1}$. This expression reveals that what determines the contribution of each measurement is its signal-to-noise ratio $\gamma_k^2/\sigma_k^2$, not $\gamma_k$ alone. If measurement $x_k$ contains less information on $x^*$, i.e., a small $\gamma_k$, the optimal weight is smaller. If the measurement contains more noise, i.e., a larger $\sigma_k^2$, the optimal weight is also smaller. If one puts a white noise as an additional measurement, its corresponding $\gamma_{K+1} = 0$ implies that its optimal weight is zero. Therefore, adding useless measurements will not change the optimal linear approximation.

An important property of the minimum MSE formula is its scale invariance. Replacing any measurement $x_k$ with $c \cdot x_k$ for a nonzero constant $c$ simultaneously rescales $\gamma_k$ by $c$ and the corresponding entries of $V$ by $c^2$, leaving $\gamma^T V^{-1}\gamma$ unchanged. The minimum MSE, optimal weights, and the approximation $\hat{x}^*(\lambda^{\text{BLUE}})$ are therefore invariant to the units or scale of each measurement. This invariance has a practical consequence: Assumption 2 normalizes $\gamma_1 = 1$, which fixes the scale of the first measurement. For the remaining measurements, only the signal-to-noise ratio $\gamma_k^2/\sigma_k^2$ matters, not the separate values of $\gamma_k$ and $\sigma_k^2$.

# 3 Convergence and Estimation

In this section, we discuss the approximation property of the BLUE linear combination when the number of measurements goes to infinity. Since the MSE-optimal linear combination dominates the BLUE in mean squared error and their mean squared error only differ by a term $1/\sigma_{x^*}^2$ (see equation (4)), their convergence properties are the same.

We consider the unitary decomposition of $V$ as follows:

$$V = U^T D U$$

where $U^T U = I$ and $D := diag\{\varsigma_1^2, \varsigma_2^2, ..., \varsigma_K^2\}$ is a diagonal matrix containing all the eigenvalues of

[4] See equation (11) in Civelli, Gaduh, and Yousuf (2024).

$V$ with $\varsigma_1^2 \geq \varsigma_2^2 \geq ... \geq \varsigma_K^2 > 0$.

### 3.1 Convergence

As the number of measurements $K$ diverges, we investigate the limit of the minimum mean squared errors. We assume

**Assumption 3.** *As $K \to \infty$, the smallest eigenvalue of $V$, i.e., $\varsigma_K^2$ satisfies for $\beta \geq 0$,*

$$\varsigma_K^2 = O(K^{-\beta}).$$

If the error variance represents the accuracy of the measurement procedure, the diminishing error variance implies that researchers may obtain more measurements from more accurate procedures.

In a big data scenario, we may obtain more measurements from different aspects that have not been considered before. In the case of many measurements, it is likely to have measurements containing less information, i.e., a small coefficient $\gamma_k$. To describe such a scenario, we assume

**Assumption 4.** *As $k \to \infty$, the coefficient $\gamma_k$ satisfies for $\alpha \geq 0$*

$$\gamma_k = O(k^{-\alpha}).$$

In the case $\alpha = 0$, it means that all the measurements are comparable in the sense that they contain comparable information on the latent variable of interest even when we obtain more and more measurements.

The asymptotic property of the expected minimum mean squared error can be summarized as follows:

**Proposition 2.** *If Assumptions 1, 2, 3, and 4 hold, then*

$$\begin{aligned}
\lim_{K\to\infty} E[\hat{x}^*(\lambda^{\text{BLUE}}) - x^*]^2 &= O\left(\frac{K^{-\beta}}{\sum_{k=1}^K k^{-2\alpha}}\right) \\
&= \begin{cases} 0 & \beta > 0 \\ 0 & \beta = 0; \frac{1}{2} \geq \alpha \geq 0 \\ O(1) & \beta = 0; \alpha \geq \frac{1}{2} \end{cases}
\end{aligned}$$

**Proof**: See Supplemental Appendix.

The magnitude of $\alpha$ describes how information on the latent variable $x^*$ diminishes in an additional measurement when there are many measurements. When $\alpha$ is small, more measurements bring in more information so that better approximation can be achieved and the true values will be reached in the limit. When $\alpha$ is large, i.e., $\alpha \geq \frac{1}{2}$, the coefficients decline so fast to zero that more measurements do not help further reduce the mean squared error. That implies that any linear combination of the measurements cannot reveal the latent variable even if the number of measurements is large.

Note that since multiplying a measurement $x_k$ by a constant $c$ rescales $\gamma_k$ by $c$ while leaving $E[\hat{x}^*(\lambda^{\text{BLUE}}) - x^*]^2 = (\gamma^T V^{-1}\gamma)^{-1}$ unchanged, Assumptions 3 and 4 implicitly require a consistent normalization of the measurements. Throughout this section, we adopt the convention that the measurements are normalized so that $\text{Var}(x_k) = O(1)$ as $k \to \infty$. Under this convention, $\gamma_k = O(k^{-\alpha})$ is equivalent to the signal-to-noise ratio $\gamma_k^2/\sigma_k^2$ being $O(k^{-2\alpha})$, since $\sigma_k^2 = \text{Var}(x_k) - \gamma_k^2\sigma_{x^*}^2 \to O(1)$. The orthonormal polynomial basis introduced in Section 4 provides one such normalization.

### 3.2 Estimation

Given the closed-form solution, the estimator of $\lambda^{\text{BLUE}}$ is a straightforward plug-in estimator through replacing population means and variances with their sample counterparts from a random sample of $\{x\}$ with sample size $N$. Let $\hat{\gamma}_k$, $\hat{\sigma}_k^2$, and $\hat{\lambda}_N^{\text{BLUE}} := (\hat{\lambda}_{0,N}, \hat{\lambda}_{1,N}, ..., \hat{\lambda}_{K,N})^T$ stand for the estimators for $\gamma_k$, $\sigma_k^2$, and $\lambda^{\text{BLUE}}$, respectively. For example, we have

$$\hat{\gamma}_2 = \frac{\frac{1}{N}\sum_{i=1}^{N} x_{2,i}x_{3,i} - \bar{x}_2\bar{x}_3}{\frac{1}{N}\sum_{i=1}^{N} x_{1,i}x_{3,i} - \bar{x}_1\bar{x}_3}.$$

Therefore, we suggest a plug-in estimator as follows:

$$\hat{\lambda}_{k,N} \equiv G_k(\hat{\mu})$$

where $\hat{\mu}$ is a $(\frac{K\times(K+3)}{2})$ by 1 vector of all the sample averages $\frac{1}{N}\sum_{i=1}^{N} x_{j,i}x_{k,i}$ for $j \geq k$ and $j, k = 0, 1, ..., K$. Notice that $\lambda_k = G_k(\mu)$ with $\mu = E\hat{\mu}$. Similarly, we may define $\lambda_0 = G_0(\mu)$. Given that function $G_k$ are known real analytic functions, we define the remainder term in a Taylor expansion as follows:

$$e_{k,N} \quad := \quad G_k(\hat{\mu}) - G_k(\mu) - \frac{\partial G_k(\mu)}{\partial \mu^T}(\hat{\mu} - \mu).$$

We assume

**Assumption 5.** *As $K \to \infty$ and $N \to \infty$, 1) $\max_{0\leq k\leq K} E|x_k|^4 < \infty$; 2) for each $j$ and $k$, $E\left|x_j x_k - E(x_j x_k)\right|^4 \leq \infty$; and 3) $\max_{0\leq k\leq K}(E|e_{k,N}|^4)^{1/4} = o\left(\sqrt{\frac{K^{3.5}\ln K}{N}}\right)$.*

This assumption guarantees the mean square convergence of the linear approximation. Assumption 5.3 is the minimum requirement for the remainder term to be dominated by the leading linear term.

The asymptotics of the proposed linear approximation is summarized as follows:

**Proposition 3.** *Suppose that Assumptions 1, 2, 3, 4 and 5 hold. Then,*

$$E|\hat{x}^*(\hat{\lambda}_N^{\text{BLUE}}) - x^*|^2 \quad = \quad O\left(\frac{K^{6.5}\ln K}{N}\right) + O\left(\frac{K^{-\beta}}{\sum_{k=1}^{K} k^{-2\alpha}}\right).$$

*Therefore, if* $\lim_{K\to\infty;N\to\infty}\frac{K^{6.5}\ln K}{N}=0$, *then the optimal linear approximation satisfies*

$$\lim_{N\to\infty} E[\hat{x}^*(\hat{\lambda}_N^{\text{BLUE}})-x^*]^2 \quad = \quad \begin{cases} 0 & \beta>0 \\ 0 & \beta=0; \frac{1}{2}\geq\alpha\geq 0 \\ O(1) & \beta=0; \alpha\geq\frac{1}{2} \end{cases}$$

**Proof**: See Supplemental Appendix.

This result implies that many measurements may reveal the true values under appropriate assumptions. But the number of measurements needs to be smaller than the sample size in the limit. Such a speed is slower than the popular restrictions imposed on the number of instruments in the many instruments literature, i.e., $O(N)$. In the weak instrument literature, a popular restriction on the coefficients on the instrument is $O(N^{-1/2})$. In the many measurements setting in this paper, the coefficient on the latent variable can only diminish at a much slower rate if the measurement error variance does not vanish.

Proposition 3 describes the convergence behavior of the estimated optimal linear approximation. When the information in the measurements diminishes too fast, i.e., $\alpha\geq\frac{1}{2}$, any linear combination of the measurements can't reveal the latent variable even if the number of measurements goes to infinity. If the information in the measurements diminishes slowly enough, i.e., $\alpha<\frac{1}{2}$, more measurements bring in more useful information so that the estimated optimal linear approximation will converge to the true values in the limit.

# 4 Nonlinear Measurements

Sections 2 and 3 establish the optimal predictor for $K$ linear measurements and its convergence properties. However, the true relationship between an observable $x_k$ and the latent variable $x^*$ may be nonlinear. Since Equation (1) is a linear projection, any nonlinear dependence on $x^*$ is absorbed into the projection residual $\varepsilon_k$. In this section, we show how to extract this nonlinear information by constructing additional measurements from the observables.

## 4.1 How to Generate an Additional Measurement?

Suppose that the conditional expectations of $x_2$ and $x_3$ given $x^*$ are nonlinear:

$$E[x_k|x^*]=m_k(x^*), \quad k=2,3,$$

where $m_k$ is a possibly nonlinear function. The linear projection $x_k=\gamma_k x^*+\varepsilon_k$ captures only the best linear approximation, and the projection residual takes the form

$$\varepsilon_k=\big(m_k(x^*)-\gamma_k x^*\big)+u_k,$$

where $u_k=x_k-m_k(x^*)$ is the pure measurement noise satisfying $E[u_k|x^*]=0$. The first component, $m_k(x^*)-\gamma_k x^*$, is the nonlinear signal — the part of $x_k$'s dependence on $x^*$ that the linear projection does not capture. This information is still present in the observables and can be recovered through nonlinear transformations.

Consider a function $h(x_2, x_3)$. Since $x_2$ and $x_3$ both depend on $x^*$, the composite $x_a \equiv h(x_2, x_3)$ will generally also depend on $x^*$. Its linear projection onto $x^*$,

$$x_a = \gamma_a x^* + \varepsilon_a, \quad \gamma_a = \frac{\text{cov}(h(x_2, x_3), x^*)}{\sigma^2_{x^*}},$$

may yield a nonzero $\gamma_a$ that captures information about $x^*$ distinct from what $\gamma_2$ and $\gamma_3$ already provide.

As a concrete example, consider the function:

$$h(x_2, x_3) = \left(\frac{x_2}{\gamma_2} - \frac{x_3}{\gamma_3}\right)^\delta.$$

Substituting $x_k = \gamma_k x^* + \varepsilon_k$ gives

$$\frac{x_2}{\gamma_2} - \frac{x_3}{\gamma_3} = \frac{\varepsilon_2}{\gamma_2} - \frac{\varepsilon_3}{\gamma_3},$$

so that $h(x_2, x_3)$ is a nonlinear function of the projection residuals $\varepsilon_2$ and $\varepsilon_3$. Because these residuals contain nonlinear functions of $x^*$ (i.e., $m_k(x^*) - \gamma_k x^*$), the transformation $h$ extracts the nonlinear signal. Specifically, raising the residual difference to a power $\delta \neq 1$ generates higher-order moments of $x^*$ that, after projection, yield a nonzero $\gamma_a$ whenever $x^*$ has a non-symmetric distribution and the nonlinear components $m_2$ and $m_3$ differ in shape (i.e., $(m_2(x^*) - \gamma_2 x^*)/\gamma_2 \neq (m_3(x^*) - \gamma_3 x^*)/\gamma_3$).

In order for the generated measurement $x_a$ to fit into the framework of Proposition 1, we need to ensure that its projection error $\varepsilon_a$ is uncorrelated with $\varepsilon_1$. This requires strengthening the assumption on $\varepsilon_1$ so that the measurement error $\varepsilon_1$ is conditional mean independent of $x^*$, $\varepsilon_2$, and $\varepsilon_3$, i.e.,

$$E[\varepsilon_1 | x^*, \varepsilon_2, \varepsilon_3] = E[\varepsilon_1]. \tag{5}$$

This assumption implies that the measurement error $\varepsilon_1$ is mean independent of $x_2$ and $x_3$. We can then consider a generated measurement

$$x_a \quad \equiv \quad h(x_2, x_3) := \gamma_a x^* + \varepsilon_a$$

with $\gamma_a$ being the projection coefficient and $\varepsilon_a$ being the projection error, i.e.,

$$\text{cov}(\varepsilon_a, x^*) = 0.$$

The user-specific function $h$ should satisfy $Ex_a^2 < \infty$, which guarantees the existence of the linear projection.

To see why $\text{cov}(\varepsilon_1, \varepsilon_a) = 0$, note that $x_a = h(x_2, x_3)$ is a function of $x_2$ and $x_3$, which are themselves functions of $x^*$, $\varepsilon_2$, and $\varepsilon_3$. Therefore $x_a = \tilde{h}(x^*, \varepsilon_2, \varepsilon_3)$ for some function $\tilde{h}$, and Equation (5) implies

$$E[\varepsilon_1 \cdot \tilde{h}(x^*, \varepsilon_2, \varepsilon_3)] = E\big[E[\varepsilon_1 | x^*, \varepsilon_2, \varepsilon_3] \cdot \tilde{h}(x^*, \varepsilon_2, \varepsilon_3)\big] = E[\varepsilon_1] \cdot E[\tilde{h}(x^*, \varepsilon_2, \varepsilon_3)],$$

so that $\text{cov}(\varepsilon_1, x_a) = 0$ and hence $\text{cov}(\varepsilon_1, \varepsilon_a) = 0$. The error variance-covariance matrix for the augmented system $(x_1, x_2, x_3, x_a)$ therefore has the structure

$$V((\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_a)^T) = \begin{pmatrix} \sigma_1^2 & 0 & 0 & 0 \\ 0 & \sigma_2^2 & 0 & \sigma_{24} \\ 0 & 0 & \sigma_3^2 & \sigma_{34} \\ 0 & \sigma_{42} & \sigma_{43} & \sigma_4^2 \end{pmatrix}.$$

The zeros in the first row and column follow from Equation (5). The off-diagonal entries $\sigma_{24}$ and $\sigma_{34}$ are generically nonzero because $x_a$ is constructed from $x_2$ and $x_3$, so the projection errors $\varepsilon_a$, $\varepsilon_2$, and $\varepsilon_3$ will typically be correlated. The flexible error covariance structure of $V$ in Proposition 1 accommodates these correlations.

We can then generate many additional measurements with various functions $h(\cdot, \cdot)$. Each function may capture different nonlinear features of the relationship between the observables and $x^*$, and the optimal linear combination in Proposition 1 automatically assigns the appropriate weight to each generated measurement. Therefore, the results developed above apply as follows.

**Corollary 1.** *If Assumptions 1, 2, and Equation (5) hold with $x = (x_1, x_2, x_3)^T$, then the results in Proposition 1 hold with $x = (x_1, x_2, x_3, x_a)^T$, where $x_a = h(x_2, x_3)$ satisfies $Ex_a^2 < \infty$.*

## 4.2 Generating a Sequence of Measurements

Corollary 1 shows that a single nonlinear function $h(x_2, x_3)$ can be added to the measurement system. A natural question is whether one can construct a systematic sequence of generated measurements $h_1, h_2, h_3, \ldots$ that progressively extracts all the nonlinear information contained in $(x_2, x_3)$. We show that an appropriate choice of basis functions achieves this, linking the construction above to the convergence framework of Section 3.

**Polynomial sieve.** Given observables $x_2$ and $x_3$, consider the bivariate polynomial basis of increasing degree, excluding the constant and linear terms already in the model:

$$\begin{aligned} &\text{Degree 2:} \quad h_1 = x_2^2, \quad h_2 = x_2 x_3, \quad h_3 = x_3^2, \\ &\text{Degree 3:} \quad h_4 = x_2^3, \quad h_5 = x_2^2 x_3, \quad h_6 = x_2 x_3^2, \quad h_7 = x_3^3, \\ &\quad \vdots \\ &\text{Degree } d: \quad x_2^j x_3^{d-j}, \quad j = 0, \ldots, d. \end{aligned}$$

Through degree $d$, the number of generated measurements is $K_d = (d+1)(d+2)/2 - 3$. By Corollary 1, each monomial $h_k(x_2, x_3)$ satisfies the measurement equation

$$h_k = \gamma_{a,k} x^* + \varepsilon_{a,k}, \quad \text{cov}(\varepsilon_1, \varepsilon_{a,k}) = 0,$$

with projection coefficient $\gamma_{a,k} = \text{cov}(h_k, x^*)/\sigma_{x^*}^2$. The augmented system $\{x_1, x_2, x_3, h_1, \ldots, h_{K_d}\}$ satisfies the framework of Proposition 1.

The key convergence property follows from the completeness of polynomials in $L^2$: under mild regularity conditions on the distribution of $(x_2, x_3)$, any square-integrable function of $(x_2, x_3)$ can be approximated arbitrarily well by linear combinations of the monomials $\{h_k\}_{k\geq 1}$, so the conditional expectation $E[x^* \mid x_2, x_3]$ lies in the closure of the linear span of $\{1, x_2, x_3, h_1, h_2, \ldots\}$. Because the MSE-optimal predictor in Proposition 1 is the $L^2$ projection of $x^*$ onto the span of the measurements, the sieve drives this projection to the best predictor that uses $x_1$ linearly and $(x_2, x_3)$ without functional-form restriction. To state the limit, we define the *oracle bound*

$$O \;:=\; \mathrm{Var}(x^* \mid x_1, x_2, x_3), \tag{6}$$

the minimum mean squared error attainable from $(x_1, x_2, x_3)$ by any predictor, linear or nonlinear, and assume the conditional mean is affine in the reference signal,

$$E[x^* \mid x_1, x_2, x_3] \;=\; \beta\, x_1 + g(x_2, x_3). \tag{7}$$

**Proposition 4.** *Suppose Assumptions 1 and 2 and the conditional-mean-independence condition* (5) *hold, that polynomials are dense in* $L^2(x_2, x_3)$*, and that the affine restriction* (7) *holds, with* $0 < O < \sigma^2_{x^*}$*. Then, along the polynomial sieve* $\{h_1, \ldots, h_{K_d}\}$ *of degree at most d,*

$$\lim_{d\to\infty} E\big[\hat{x}^*(\lambda^{\mathrm{MSE}}_{K_d}) - x^*\big]^2 \;=\; O, \tag{8}$$

$$\lim_{d\to\infty} E\big[\hat{x}^*(\lambda^{\mathrm{BLUE}}_{K_d}) - x^*\big]^2 \;=\; \frac{\sigma^2_{x^*}\, O}{\sigma^2_{x^*} - O} \;>\; O. \tag{9}$$

*The MSE-optimal predictor attains the oracle bound, while the BLUE converges to a strictly larger limit; the two coincide only as* $O \to 0$ *(equivalently* $\sigma^2_{x^*} \to \infty$*).*
*A sufficient condition for* (7) *is that, conditional on* $(x_2, x_3)$*, the pair* $(x^*, \varepsilon_1)$ *is jointly Gaussian, with* $x^* \mid (x_2, x_3) \sim \mathcal{N}\big(E[x^* \mid x_2, x_3], \mathrm{Var}(x^* \mid x_2, x_3)\big)$*,* $\mathrm{Var}(x^* \mid x_2, x_3)$ *constant in* $(x_2, x_3)$*, and* $\varepsilon_1 \perp (x^*, x_2, x_3)$ *with* $\varepsilon_1 \sim \mathcal{N}(0, \sigma_1^2)$*.*

**Proof**: See Supplemental Appendix.

The MSE-optimal limit (8) is the oracle bound—the minimum MSE achievable from the three observables by any predictor, linear or nonlinear. The gain from generated measurements—the gap between the linear-only MSE $(\gamma^T V^{-1}\gamma)^{-1}$ and the oracle bound—represents the information about $x^*$ contained in the nonlinear dependence of $(x_2, x_3)$ on $x^*$. The two limits in (8)–(9) differ because the unbiasedness restriction $\lambda_1^T\gamma = 1$ prevents the BLUE from trading bias for variance: by the dominance relation (4) the two risks differ in precision by the constant $1/\sigma^2_{x^*}$, so when the limiting risk is positive the BLUE converges to a strictly larger value, in contrast to the convergence-to-zero regime of Section 3 where the gap is negligible.

### 4.3 When Does Nonlinearity Matter?

The oracle bound equals the linear-only MSE when $(x^*, x_2, x_3)$ are jointly Gaussian and the measurement equations are linear. In such a case, generated measurements provide no improvement.

However, the gap is positive in two economically relevant scenarios:

(i) *Nonlinear measurement equations*: The conditional mean $E[x_k|x^*] = m_k(x^*)$ is nonlinear. For example, satellite nighttime light sensors saturate at high luminosity levels, making nighttime light growth a concave function of true GDP growth. Similarly, logarithmic relationships between emissions and economic activity introduce curvature that the linear projection misses.

(ii) *Non-Gaussian latent variable*: Even with linear measurement equations, if $x^*$ has a non-Gaussian distribution—for instance, if the cross-country distribution of GDP growth is right-skewed—then $E[x^*|x_k]$ is nonlinear.

We illustrate these two channels with concrete examples.

**Example 1** (Powers of a single signal)**.** *The simplest sequence uses powers of one observable. Suppose $\varepsilon_3$ is independent of $x^*$ (not merely uncorrelated) with $E[\varepsilon_3] = 0$, and let $E[x^*] = 0$ for simplicity. Define $h_k = x_3^k$ for $k = 2, 3, \ldots$. For $k = 2$:*

$$x_3^2 = \gamma_3^2(x^*)^2 + 2\gamma_3 x^* \varepsilon_3 + \varepsilon_3^2.$$

*Since $\varepsilon_3$ is independent of $x^*$:*

$$\gamma_{a,2} = \frac{cov(x_3^2, x^*)}{\sigma_{x^*}^2} = \frac{\gamma_3^2\, E[(x^*)^3]}{\sigma_{x^*}^2} = \gamma_3^2 \kappa_3 \sigma_{x^*}, \tag{10}$$

*where $\kappa_3 = E[(x^*)^3]/\sigma_{x^*}^3$ is the skewness of $x^*$. Thus the squared signal $x_3^2$ is informative about $x^*$—in the sense that $\gamma_{a,2} \neq 0$—if and only if $x^*$ is asymmetrically distributed ($\kappa_3 \neq 0$). This illustrates channel (ii): non-Gaussianity of the latent variable is what makes even-powered transformations useful.*

**Example 2** (Powers of residual differences)**.** *A sequence that directly extends the single generated measurement in Corollary 1 is*

$$h_k(x_2, x_3) = \left(\frac{x_2}{\gamma_2} - \frac{x_3}{\gamma_3}\right)^k, \quad k = 2, 3, 4, \ldots$$

*As shown earlier, $x_2/\gamma_2 - x_3/\gamma_3 = \varepsilon_2/\gamma_2 - \varepsilon_3/\gamma_3$. When the measurement equations are nonlinear, the residuals $\varepsilon_k = (m_k(x^*) - \gamma_k x^*) + u_k$ contain the nonlinear signal $m_k(x^*) - \gamma_k x^*$. Define the* nonlinear discrepancy

$$\delta(x^*) := \frac{m_2(x^*)}{\gamma_2} - \frac{m_3(x^*)}{\gamma_3},$$

*which captures the difference in the nonlinear shapes of the two measurement equations. Then*

$$\frac{x_2}{\gamma_2} - \frac{x_3}{\gamma_3} = \delta(x^*) + \nu, \qquad \nu := \frac{u_2}{\gamma_2} - \frac{u_3}{\gamma_3},$$

*and $h_k = (\delta(x^*) + \nu)^k$ captures the $k$-th moment of the nonlinear discrepancy. Successive powers extract progressively finer features of $\delta(x^*)$. By the density of polynomials, as $k \to \infty$ the* *span of $\{h_k\}$ includes all $L^2$ functions of $\delta(x^*) + \nu$, so the optimal linear predictor converges to $E[x^*|x_1, x_2, x_3, \delta(x^*) + \nu]$.*

*This sequence is informative whenever $\delta(x^*) \not\equiv 0$, i.e., whenever the two measurement equations have different nonlinear shapes. For instance, if nighttime light growth saturates at high GDP growth (concave $m_2$) while $NO_2$ density grows approximately linearly ($m_3 \approx \gamma_3 x^*$), then $\delta(x^*) = m_2(x^*)/\gamma_2 - x^* \neq 0$ and the power series captures information about $x^*$ that neither $x_2$ nor $x_3$ provides through their linear projections. This illustrates channel (i).*

### 4.4 Implications for Convergence

We now connect the generated measurements to the convergence framework of Section 3. A subtlety arises: for raw monomials $h_k = x_2^j x_3^{d-j}$, the projection coefficient $\gamma_{a,k}$ need not decay with $k$—it can grow, because the variance of higher-degree monomials grows. Moreover, any measurement is scalable: replacing $h_k$ with $c \cdot h_k$ multiplies $\gamma_{a,k}$ by $c$, so $\gamma_{a,k}$ can be made to satisfy any desired decay rate through rescaling. As discussed in Section 2, the projection coefficient $\gamma_{a,k}$ alone does not characterize how much information $h_k$ carries about $x^*$; the signal-to-noise ratio $\gamma_{a,k}^2/\sigma_{a,k}^2$ is the correct scale-invariant measure.

**Orthonormal basis and the law-of-total-variance bound.** To make the connection with Proposition 2 precise, we adopt a canonical normalization. Let $\{\phi_k(x_2, x_3)\}_{k\geq 1}$ be polynomials orthonormal with respect to the distribution of $(x_2, x_3)$:

$$E[\phi_k] = 0, \quad E[\phi_k \phi_l] = \delta_{kl}.$$

Each $\phi_k$ has unit variance, so its projection coefficient $\gamma_{a,k} = E[\phi_k \cdot x^*]/\sigma_{x^*}^2$ directly reflects informativeness without normalization ambiguity. The conditional expectation has the orthogonal expansion

$$E[x^*|x_2, x_3] - E[x^*] = \sum_{k=1}^{\infty} \alpha_k \, \phi_k(x_2, x_3), \qquad \alpha_k = E[x^* \phi_k] = \gamma_{a,k} \sigma_{x^*}^2.$$

Two arguments combine to bound the sum of squared coefficients. First, Parseval's identity applied to the orthogonal expansion gives

$$\sum_{k=1}^{\infty} \alpha_k^2 = \mathrm{Var}\big(E[x^*|x_2, x_3]\big). \tag{11}$$

Second, the law of total variance,

$$\sigma_{x^*}^2 = \mathrm{Var}\big(E[x^*|x_2, x_3]\big) + E\big[\mathrm{Var}(x^*|x_2, x_3)\big],$$

implies $\text{Var}(E[x^*|x_2, x_3]) \leq \sigma^2_{x^*}$ since the second term is non-negative. Combining the two, and using $\alpha_k = \gamma_{a,k}\sigma^2_{x^*}$,

$$\sum_{k=1}^{\infty} \gamma^2_{a,k} \;=\; \frac{\sum_{k=1}^{\infty} \alpha_k^2}{\sigma^4_{x^*}} \;=\; \frac{\text{Var}\big(E[x^*|x_2, x_3]\big)}{\sigma^4_{x^*}} \;\leq\; \frac{\sigma^2_{x^*}}{\sigma^4_{x^*}} \;=\; \frac{1}{\sigma^2_{x^*}} \;<\; \infty. \tag{12}$$

This has three important consequences:

1. *The coefficients decay*: $\gamma_{a,k} \to 0$ as $k \to \infty$ in the orthonormal basis. The rate of decay depends on the smoothness of $E[x^*|x_2, x_3]$ as a function of $(x_2, x_3)$: exponential decay for analytic conditional expectations, polynomial decay $\gamma_{a,k} = O(k^{-s/2})$ for $s$-times differentiable ones.

2. *Total information is bounded*: The sum $\sum \gamma^2_{a,k} < \infty$ means that the generated measurements collectively contain a finite amount of information about $x^*$, regardless of how many are generated. This is because all generated measurements are functions of the same two observables $(x_2, x_3)$.

3. *MSE converges to a positive constant*: Unlike independent signals from different sources, generated measurements cannot drive MSE to zero.

**Connection to Propositions 2–3.** In the orthonormal basis, $\text{Var}(\phi_k) = 1$ implies that the error variances $\sigma^2_{a,k} = 1 - \gamma^2_{a,k}\sigma^2_{x^*} \to 1$ as $k \to \infty$, so they are bounded away from zero ($\beta = 0$ in Assumption 3). The convergence of $\sum \gamma^2_{a,k}$ requires $\gamma_{a,k} = o(k^{-1/2})$, which places the generated measurements in the regime $\alpha \geq 1/2$ of Proposition 2. In this regime, the MSE converges to a positive constant—consistent with the oracle bound $\text{Var}(x^*|x_1, x_2, x_3) > 0$.

This contrasts with independent signals from genuinely different sources, where $\sum \gamma_k^2$ can diverge: each new signal brings information about $x^*$ that is not contained in the previous signals. In that case, $\alpha < 1/2$ is achievable, and Proposition 2 guarantees MSE $\to 0$.

**Remark 3:** *Complementarity of generated and independent measurements.*
Generated measurements and independent signals play complementary roles. Independent signals drive the oracle bound toward zero as $K \to \infty$, because each independent signal adds new information not spanned by the existing measurements. Generated measurements close the gap between the linear-only MSE $(\gamma^T V^{-1} \gamma)^{-1}$ and the oracle bound $\text{Var}(x^*|x_1, \ldots, x_K)$ at each finite $K$, by extracting nonlinear information that linear projections miss. The law-of-total-variance bound $\sum \gamma^2_{a,k} < \infty$ quantifies the total nonlinear information available from each pair of observables. The most effective approach combines both: many independent signals augmented with their nonlinear interactions.

# 5 Simulations

We present three simulations, each with $N = 5{,}000$ observations, to illustrate the convergence theory developed in Sections 3 and 4. The first two simulate an expanding sequence of independent signals and show how the rate of SNR decay determines whether MSE converges to zero or to a positive constant, as stated in Proposition 2. The third simulates a setting where the base signals are *nonlinear* in the latent variable; adding polynomial moments of these signals progressively reduces the MSE but the reduction is bounded, consistent with the law-of-total-variance bound of Section 4.4.

### 5.1 Simulation 1: Independent Signals with Fast SNR Decay

The data-generating process (DGP) is

$$x_k \;=\; \gamma_k\, x^* + \varepsilon_k, \qquad \gamma_k = k^{-1/2}, \quad \varepsilon_k \overset{iid}{\sim} N(0,1), \quad k = 1, 2, \ldots, K,$$

with $x^* \sim N(0,1)$ and $\sigma_k^2 = 1$ for all $k$. Since the projection coefficients satisfy $\gamma_k = k^{-1/2}$, the cumulative SNR is

$$\sum_{k=1}^{K} \frac{\gamma_k^2}{\sigma_k^2} = \sum_{k=1}^{K} \frac{1}{k} = H_K \to \infty,$$

where $H_K$ is the $K$-th harmonic number. The MSE of the optimal unbiased predictor (Proposition 1) therefore equals $1/H_K$, which converges to zero as $K \to \infty$, confirming the $\alpha < 1/2$ case of Proposition 2.

Because the cross-covariance identification of Proposition 1 requires at least three mutually-uncorrelated measurements (Assumption 1), the empirical MSE is reported starting from $K = 3$. Panel (a) of Figure 1 plots the theoretical MSE $1/H_K$ together with the empirical MSE from the simulation. Both decline smoothly toward zero, and the empirical curve closely tracks the theoretical one across the full range of $K$ tested.

### 5.2 Simulation 2: Independent Signals with Slow SNR Decay

We keep the same DGP structure but set

$$\gamma_k = \frac{1}{k}, \qquad \varepsilon_k \overset{iid}{\sim} N(0,1).$$

The cumulative SNR now sums to

$$\sum_{k=1}^{K} \frac{\gamma_k^2}{\sigma_k^2} = \sum_{k=1}^{K} \frac{1}{k^2} \xrightarrow{K\to\infty} \frac{\pi^2}{6},$$

by the Basel result. The limiting MSE is $6/\pi^2 \approx 0.608 > 0$, placing this DGP firmly in the $\alpha \geq 1/2$ regime of Proposition 2 where the MSE is bounded away from zero no matter how many signals are added.

Panel (b) of Figure 1 illustrates this: starting from $K = 3$, the theoretical MSE is 0.73 and the empirical MSE is 0.74; both decelerate as $K$ grows and asymptote to the theoretical limit $6/\pi^2 \approx 0.608$ (dashed line).

Figure 1: Simulation Results: MSE as a Function of the Number of Signals $K$

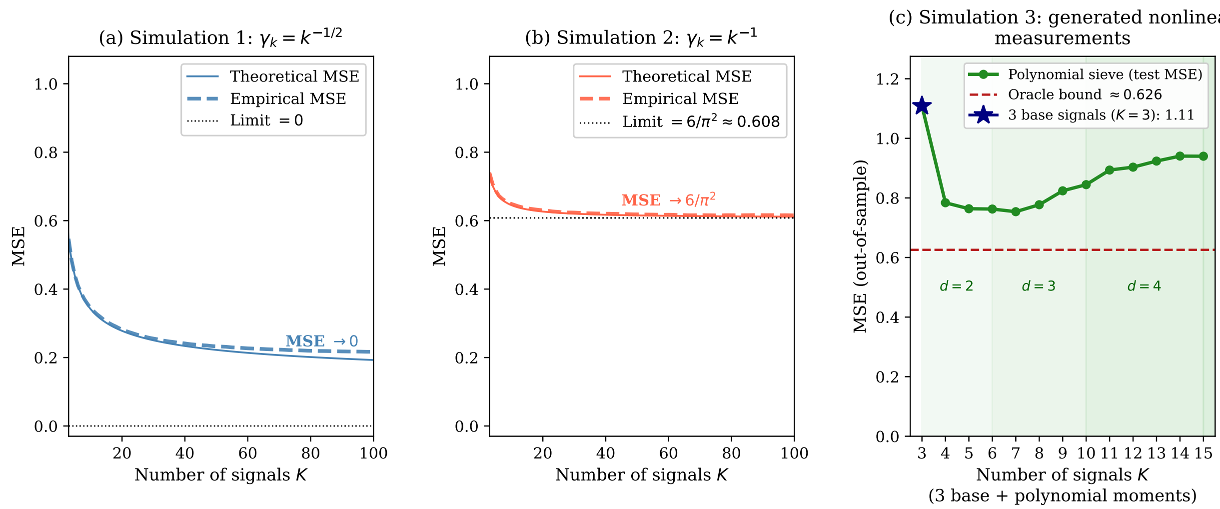


*Notes:* This figure presents the three simulation results, each with $N = 5{,}000$ observations and 80/20 train–test split throughout. In panels (a) and (b) the empirical MSE curve is averaged over 50 independent Monte Carlo replications to suppress test-set sampling noise at large $K$. For panels (a)–(b), the theoretical MSE is the population formula $1/\left(\sum_{j=1}^{K} \gamma_j^2\right)$ while the empirical MSE follows the paper's identification strategy in Section 3. For Panel (c), when $K > 3$, polynomial monomials $x_2^j x_3^l$ of total degree $d = j + l$ are added as generated measurements, with the $x_2$-dominant terms (highest $j$) added first within each degree. The oracle bound explicitly uses $\exp(2x_3) \approx x^*$ and $9x_2^2 \approx x^*$ as additional signals.

## 5.3 Simulation 3: Generated Nonlinear Measurements

The third simulation illustrates the role of nonlinear generated measurements when the base signals are themselves nonlinear functions of the latent variable. The DGP is

$$
\begin{aligned}
x_1 &= x^* + \varepsilon_1, & \varepsilon_1 &\sim N(0,\, 2.25),\\
x_2 &= \tfrac{1}{3}\sqrt{x^*} + \varepsilon_2, & \varepsilon_2 &\sim N(0,\, 0.0225),\\
x_3 &= 0.5\log(x^*) + \varepsilon_3, & \varepsilon_3 &\sim N(0,\, 0.09),
\end{aligned}
$$

with $x^* \sim \text{LogNormal}(0, 1)$, so that $x^*$ is always positive and right-skewed ($\text{Var}(x^*) \approx 4.30$ in sample). The noise standard deviations are calibrated such that the three signals have comparable signal-to-noise ratios: the signal components have standard deviations of approximately 2.07, 0.20, and 0.50 for $x_1$, $x_2$, $x_3$, while the noise standard deviations are 1.5, 0.15, and 0.30.

Here $x_2$ and $x_3$ are nonlinear in $x^*$: $x_2$ is a square-root transformation and $x_3$ is a logarithmic transformation. As discussed in Section 4, these nonlinear relationships mean that the linear predictor using only $(x_1, x_2, x_3)$ is sub-optimal; richer polynomial functions of the observed signals contain additional information about $x^*$.

Starting from the three base signals, we incrementally add polynomial monomials $x_2^j x_3^l$ of total degree $d = j + l$ for $d = 2, 3, 4$, with the highest power of $x_2$ added first within each degree. The structural parameters $(\gamma_k, \mathbf{V})$ are estimated purely from cross-covariances of observables following Section 3, and all MSE values are evaluated out-of-sample using an 80/20 train–test split.

Panel (c) of Figure 1 reports the results. With only the three base signals ($K = 3$), the test MSE is 1.11 (about 25.8% of $\text{Var}(x^*)$). Adding the first degree-2 moment $x_2^2$—which satisfies $9x_2^2 \approx x^*$ and therefore provides an approximately linear signal for $x^*$—immediately reduces the MSE to 0.78, achieving a 29% improvement. Further polynomial moments fluctuate around 0.8–1.0, consistent with the law-of-total-variance bound: the total information that polynomial functions of $(x_2, x_3)$ can contribute to predicting $x^*$ is finite, bounded by $\sum_k \gamma_{a,k}^2 \leq 1/\sigma_{x^*}^2$ (equation (12)). The dashed oracle-bound line ($\approx 0.63$), computed from the BLUE combination in Proposition 1 that explicitly uses $\exp(2x_3) \approx x^*$ and $9x_2^2 \approx x^*$ as key generated signals, represents the asymptotic information floor of polynomial sieves in this DGP.

The contrast with Simulation 1 is instructive: in Simulation 1, independent signals from different sources drive the MSE to zero because each new signal provides fresh, uncorrelated information about $x^*$. In Simulation 3, all generated moments are functions of the same three observables $(x_1, x_2, x_3)$, so the law-of-total-variance bound caps the total information and the MSE cannot reach zero regardless of how many polynomial terms are added.

# 6 Measuring Economic Growth With Many Measurements

Official GDP growth provides a useful summary of economic activity, but its accuracy varies with the size of the informal economy, statistical capacity, the frequency of revisions, and political incentives to misreport. To improve on it, we assemble eight country-year measurements: official real GDP growth ($x_1$), nighttime light log growth ($x_2$), $NO_2$ tropospheric-column log growth ($x_3$), urban land-cover share log change ($x_4$), NDVI level change ($x_5$), NDBI level change ($x_6$), $CO_2$ emissions log growth ($x_7$), and electricity-consumption log growth ($x_8$), whose sources are described in Section 6.1. We consider two complementary specifications to combine them: (1) pairing official GDP with satellite signals alone, and (2) adding traditional measurements. In other words, the first six measurements enter Specification 1, while measurements 1–3, 7, and 8 enter Specification 2. The first specification isolates the informational content of satellite measurements. By contrast, the second specification exploits the rich information content of traditional measurements with satellite measurements playing an identification role. For each specification, we examine signal contributions at both annual frequency and long horizons, and assess whether nonlinear transformations of the signals add further content.

## 6.1 Data Sources

**Real GDP growth ($x_1$).** Annual real GDP growth rates and real GDP levels at purchasing-power parity (constant 2017 international dollars) come from the World Bank's World Development Indicators. We use the growth-rate series at annual frequency and transform it to log growth via $x_{1,it} = 100 \ln(1 + \text{growth}_{it}/100)$.

**Nighttime light log growth ($x_2$).** We use the Simulated VIIRS Nighttime Light (SVNL) dataset of Chen, Wang, Zhang, Shen, and Chen (2024), a global annual harmonized 500 m product spanning 1992–2023 constructed by a U-Net super-resolution neural network (NTLSRU-Net) that performs cross-sensor calibration between DMSP/OLS imagery (1992–2013) and NPP-VIIRS imagery (2012–2023), with Landsat NDVI as an auxiliary input to reduce saturation and blooming in DMSP data. The SVNL product is internally consistent across the DMSP–VIIRS transition by construction, avoiding the discontinuity that arises in raw-concatenated DMSP–VIIRS series. We aggregate pixel values to country-year totals over each country's administrative boundary and compute annual log growth.

**$NO_2$ tropospheric-column log growth ($x_3$).** We use the OMI/Aura $NO_2$ Cloud-Screened Total and Tropospheric Column L3 Global Gridded $0.25° \times 0.25°$ V3 product (OMNO2d) from NASA's Goddard Earth Sciences Data and Information Services Center (GES DISC). The product reports daily global grids of tropospheric $NO_2$ vertical column density in molecules/cm$^2$, retrieved from the Ozone Monitoring Instrument aboard the Aura satellite, which entered operation in October 2004. We download all daily HDF-EOS5 files (7,270 files covering January 2005 through December 2024) via the NASA Earthdata API. For each 0.25° grid cell we compute the arithmetic mean of valid daily cloud-screened observations (`ColumnAmountNO2TropCloudScreened`) within each calendar year. Cells are assigned to countries by intersecting their centers with Natural Earth Admin-0 polygons, and country-year values are computed as the area-weighted mean across assigned cells with weights equal to the true spherical-cap area, $A_{ij} = R^2 \, \Delta\lambda \, (\sin\phi_{n,j} - \sin\phi_{s,j})$. We then compute annual log growth.

**Urban land-cover share log change ($x_4$).** We use the Copernicus Climate Change Service (C3S) Global Land Cover product (`C3S-LC-L4-LCCS-Map-300m`), the operational continuation of ESA-CCI Land Cover, at 300 m spatial resolution and annual frequency 2000–2022. Following Civelli, Gaduh, and Yousuf (2024), we measure urban land cover by its area-weighted share of the country's land area rather than its total area in km$^2$. For each country-year, the urban share is

$$\text{share}_{i,t} \;=\; \frac{\sum_{p\in\text{country}_i} \mathbf{1}[\text{LCCS}_{p,t} = 190] \cdot a_p}{\sum_{p\in\text{country}_i} \mathbf{1}[\text{LCCS}_{p,t} \neq 210, \neq 0] \cdot a_p},$$

where $a_p$ is the true spherical-cap area of pixel $p$, the numerator sums pixel areas with LCCS class 190 (urban), and the denominator sums all non-water non-no-data pixel areas. Aggregation is performed on Google Earth Engine over FAO GAUL 2015 level-0 polygons. The signal $x_4$ is the year-on-year log change in this share in percent, $x_{4,it} = 100 \cdot (\ln \text{share}_{i,t} - \ln \text{share}_{i,t-1})$; the long-difference variable is $x_4^{\text{LD}} = 100 \cdot (\ln \overline{\text{share}}_{2020\text{–}22} - \ln \overline{\text{share}}_{2006\text{–}08})$.

**NDVI level change ($x_5$) and NDBI level change ($x_6$).** NDVI is constructed from MODIS Terra Vegetation Indices, `MODIS/061/MOD13A1`, a 16-day global composite at 500 m. We retain pixels flagged "good" or "marginal" by the summary quality assurance band (`SummaryQA`) and average the scaled NDVI annually within each country. NDBI is constructed from MODIS Terra Surface Reflectance, `MODIS/061/MOD09A1`, an 8-day composite at 500 m. For each cloud-screened pixel (`StateQA` bits 0–2 and 10 clear), we compute NDBI = (SWIR − NIR)/(SWIR + NIR) from bands `sur_refl_b06` and `sur_refl_b02`, average annually within each country, and retain the country-year level. Both indices are aggregated on Google Earth Engine via area-weighted reduction over FAO GAUL 2015 polygons. Because NDVI and NDBI are bounded ratios that can be negative, we use level changes ($x_{5,it} = \text{NDVI}_{i,t} - \text{NDVI}_{i,t-1}$ and analogously for NDBI) rather than log differences.

**$CO_2$ emissions log growth ($x_7$).** $CO_2$ emissions data come from Our World in Data, which packages the Global Carbon Project series (Friedlingstein, Jones, O'Sullivan, Andrew, Hauck, Peters, Peters, Pongratz, Sitch, Le Quéré, et al. (2019)) of fossil-fuel and industry emissions at country level. We use total annual $CO_2$ emissions in million tonnes and compute annual log growth.

**Electricity consumption log growth ($x_8$).** Electricity consumption (annual, country level, terawatt-hours) comes from the Our World in Data energy dataset, which compiles data from Ember and the U.S. Energy Information Administration. We compute annual log growth.

**Country boundaries and harmonization.** Where consistent with the underlying data product, we use FAO Global Administrative Unit Layers (GAUL) 2015 level-0 polygons; otherwise we use Natural Earth Admin-0 polygons. Country identifiers are harmonized to ISO 3166-1 alpha-3 codes.

## 6.2 GDP Paired with Satellite Signals

In this specification, we pair official GDP with five satellite signals only – nighttime lights, $NO_2$, urban land cover, NDVI, and NDBI (Model A). These satellite measurements have been shown to reflect economic activity under different circumstances in the literature[5], and their measurement errors are plausibly uncorrelated with that in official GDP. We assess which signals carry growth information, how their usefulness scales with the time horizon, and what their mechanisms are.

Model A ($K = 6$) maintains $\varepsilon_k \perp \varepsilon_{\text{GDP}}$ for $k = 2, \dots, 6$ and uses GDP as the fully-uncorrelated reference of Assumption 1. The only structural off-diagonal in $V$ is $\sigma_{5,6}$, accommodating the NDVI–NDBI shared MODIS calibration error.[6] We leave the detailed setup of the model – measurement equations, summary statistics, the structural-zero pattern of the error variance matrix, and identification – to Appendix A. Throughout we use the 1%/99%-trimmed sample. After 1%/99% trimming, the panel contains $N = 2{,}203$ observations on 146 countries (2006–2022) and the long-difference cross-section contains 124 countries.

Table 1 reports the estimated structural parameters $\gamma_k$ and $\sigma_k^2$, the signal-to-noise ratio $\text{SNR}_k = \gamma_k^2 \widehat{\sigma}_{y^*}^2 / \sigma_k^2$, and the effective BLUE contribution $\lambda_k \gamma_k$ for both the panel and the long-difference. Since the BLUE combination requires $\sum_k \lambda_k \gamma_k = 1$ for unbiasedness, the product weight $\lambda_k \gamma_k$ reflects the contribution of each measure.

In the panel, official GDP dominates: its BLUE contribution is $\lambda_1 \gamma_1 = 0.869$, with all five satellite signals contributing only 0.131 in aggregate. At annual frequency, the signal-to-noise ratios of the satellite signals are uniformly low – nighttime lights' $\text{SNR}_2 = 0.101$ is the highest

[5] See Donaldson and Storeygard (2016) for a survey of satellite data in economics.

[6] NDVI = (NIR − Red)/(NIR + Red) and NDBI = (SWIR − NIR)/(SWIR + NIR), where NIR, Red, and SWIR denote near-infrared, red-band, and short-wave infrared surface reflectances respectively. Both indices are computed from the same MODIS surface-reflectance product, so their calibration errors share a common component.

Table 1: Specification 1: Estimation Results – Panel vs. Long-Difference

| | Description | Panel (2006–2022) | | | | Long-difference | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | $\gamma_k$ | $\sigma_k^2$ | SNR | $\lambda_k\gamma_k$ | $\gamma_k$ | $\sigma_k^2$ | SNR | $\lambda_k\gamma_k$ |
| $x1$ | GDP | 1.000 | 15.556 | 0.929 | 0.869 | 1.000 | 886.16 | 0.263 | 0.083 |
| $x2$ | NL | 1.416 | 286.58 | 0.101 | 0.095 | 1.928 | 990.01 | 0.875 | 0.276 |
| $x3$ | $NO_2$ | 0.2331 | 61.139 | 0.013 | 0.012 | 0.3637 | 246.07 | 0.125 | 0.040 |
| $x4$ | Urban land cover | 0.1214 | 9.727 | 0.022 | 0.021 | 1.356 | 246.47 | 1.739 | 0.549 |
| $x5$ | NDVI | −0.000156 | 0.000177 | 0.002 | 0.003 | −0.000144 | 0.000113 | 0.042 | 0.027 |
| $x6$ | NDBI | −0.000042 | 0.000318 | 0.000 | 0.000 | −0.000178 | 0.000185 | 0.040 | 0.026 |
| $\widehat{\sigma}^2_{y^*}$ | | | 14.455 | | | | 232.97 | | |
| Min. MSE | | | 13.523 | | | | 73.528 | | |
| $N$ | | | 2203 | | | | 124 | | |

*Notes:* This table presents the Specification 1 estimation results in the panel and in the long-difference cross-section, using the 1%/99%-trimmed sample. The panel consists of 2,203 country-year observations on 146 countries over 2006–2022, and the long-difference consists of a 124-country cross-section comparing 2006–08 averages with 2020–22 averages. Identification uses GDP as the fully-uncorrelated reference, and $V$ is diagonal apart from $\sigma_{5,6}$. The effective contributions $\lambda_k\gamma_k$ sum to one by unbiasedness.

among them – so the BLUE rightly assigns most weight to official GDP as the least-noisy available measure. In the long-difference, the balance shifts sharply. Measurement noise washes out in long-run averages while the cumulative latent signal grows with the horizon, causing the SNRs of the satellite signals to rise substantially. Official GDP's BLUE share falls to 0.083 while the satellite signals collectively account for the remaining 0.917, with urban land cover alone reaching $\lambda_4\gamma_4 = 0.549$ – the highest effective contribution of any signal.

Among the satellite signals, nighttime lights lead in the panel ($\text{SNR}_2 = 0.101$, $\lambda_2\gamma_2 = 0.095$) while urban land cover leads in the long-difference ($\text{SNR}_4 = 1.74$, $\lambda_4\gamma_4 = 0.549$, versus nighttime lights' $\lambda_2\gamma_2 = 0.276$). The ordering reflects what each signal captures. Nighttime lights record both the spatial extent of economic activity and its intensive margin – the brightness of existing lights within illuminated areas – making them responsive to annual fluctuations in output. Urban land cover tracks only the spatial expansion of built-up area; growth in economic intensity within existing pixels is invisible to it. Because spatial expansion is a slow structural process, urban land cover has a low panel SNR (0.022) but becomes the most informative signal over a fifteen-year span ($\text{SNR}_4 = 1.74$), while nighttime lights retain a high long-difference SNR (0.88) through their intensive-margin channel.

$NO_2$, NDVI, and NDBI contribute little in either setting. Tropospheric $NO_2$ density is confounded by wind, photolysis, and seasonal mixing that add idiosyncratic noise, yielding SNRs of 0.013 in the panel and 0.125 in the long-difference and effective BLUE contributions of 0.012 and 0.040 respectively. NDVI and NDBI have the lowest SNRs of any signal – 0.002 and 0.0001 in the panel – because as 500 m MODIS products they aggregate over large pixels and miss the within-pixel intensive-margin variation that growth would amplify; their long-difference SNRs of 0.042 and 0.040 remain well below those of nighttime lights and urban land cover.

Although $NO_2$ contributes little to the BLUE prediction directly, it plays an indispensable identification role in Specification 1. In both specifications, the nighttime light loading $\gamma_{\text{NL}}$ is recovered from the ratio Cov(NL, $NO_2$)/Cov(GDP, $NO_2$), which requires a non-trivial covariance between $NO_2$ and GDP. Table 2 shows that $NO_2$'s correlation with GDP is 0.085 in the panel and rises to 0.306 in the long-difference, large enough to anchor the identification. Urban land cover, by contrast, has a comparable panel correlation with GDP (0.108) but its long-difference correlation falls to 0.054 – near zero – making it an unreliable identifying signal precisely at the horizons where satellite signals are most valuable.

Table 2: Correlations of Satellite Signals with Official GDP

| | Panel | | Long-difference | |
|---|---|---|---|---|
| Signal | Corr. with GDP | $N$ | Corr. with GDP | $N$ |
| NL | +0.161 | 2,996 | +0.211 | 175 |
| $NO_2$ | +0.085 | 2,713 | +0.306 | 156 |
| Urban land cover | +0.108 | 2,423 | +0.054 | 139 |
| NDVI | −0.047 | 2,683 | +0.005 | 153 |
| NDBI | +0.009 | 3,190 | +0.073 | 186 |

*Notes:* This table reports the correlations between each satellite signal and official GDP, in the panel and in the long-difference cross-section, using the 1%/99%-trimmed sample. The panel consists of country-year observations over 2006–2022, and the long-difference is the country-level difference of 2020–22 averages minus 2006–08 averages. GDP, NL, $NO_2$, and urban land cover are expressed as percent (log) changes, while NDVI and NDBI are unitless level changes. The number of observations $N$ varies across signals due to data availability.

Next, we assess whether nonlinear transformations of the three informative satellite signals add further content. Dropping NDVI and NDBI – both uninformative at annual frequency – we extend $(x_1, x_2, x_3, x_4) = (\text{GDP, NL, NO}_2, \text{Urban})$ with three squares and three pairwise cross terms of $(x_2, x_3, x_4)$, producing a $K = 10$ polynomial sieve (Model A$'$). On the same $N = 2{,}329$ panel (148 countries), the BLUE MSE falls by 5.4% from 11.51 at $K = 4$ to 10.88 at $K = 10$. Among the six generated measurements, only $(\text{NL})^2$ contributes meaningfully, with $\lambda_5\gamma_5 = 0.014$ and SNR 0.008 – a small but visible share, while the other five generated terms each contribute under 0.005. Because the nighttime light data we use is free of the top-coding that saturates early generations of satellite imagery in bright areas, this curvature reflects a genuine nonlinear feature of the light–output relationship rather than a sensor artifact. The full $K = 10$ table is in Appendix A.

## 6.3 GDP Paired with Satellite Signals and Traditional Measurements

In this specification, we combine official GDP with two satellite signals (nighttime lights and $NO_2$) and two traditional energy statistics ($CO_2$ emissions and electricity consumption) whose measurement errors are correlated with that of official GDP. This specification assesses how far traditional measurements – once admitted into the model with their correlated-error structure – can lower the BLUE MSE in the annual panel and at long horizons.

Model B ($K = 5$) accommodates the correlated errors with $\sigma_{1,4} \neq 0$, $\sigma_{1,5} \neq 0$, and $\sigma_{4,5} \neq 0$. We use NL as the reference signal, whose measurement error is generated by a separate satellite system and is uncorrelated with every other error in the model. After 1%/99% trimming on $x_2$–

$x_5$, the panel contains 2,526 country-year observations on 163 countries and the long-difference cross-section contains 145 countries. Appendix A reports the full setup.

Table 3: Specification 2: Estimation Results – Panel vs. Long-Difference

| | | Panel (2006–2022) | | | | Long-difference | | | |
|---|---|---|---|---|---|---|---|---|---|
| | Description | $\gamma_k$ | $\sigma_k^2$ | SNR | $\lambda_k\gamma_k$ | $\gamma_k$ | $\sigma_k^2$ | SNR | $\lambda_k\gamma_k$ |
| $x1$ | GDP | 1.000 | 13.499 | 0.622 | 0.277 | 1.000 | 849.86 | 0.164 | 0.008 |
| $x2$ | NL | 1.689 | 300.56 | 0.080 | 0.028 | 3.325 | 346.52 | 4.436 | 0.899 |
| $x3$ | $NO_2$ | 0.4244 | 59.015 | 0.026 | 0.009 | 0.3737 | 269.60 | 0.072 | 0.015 |
| $x4$ | $CO_2$ | 1.995 | 46.599 | 0.717 | 0.330 | 2.215 | 1887.98 | 0.361 | 0.045 |
| $x5$ | Electricity | 1.401 | 21.275 | 0.775 | 0.357 | 1.914 | 1575.76 | 0.323 | 0.033 |
| $\widehat{\sigma}^2_{y^*}$ | | 8.399 | | | | 139.00 | | | |
| Min. MSE | | 2.938 | | | | 28.183 | | | |
| $N$ | | 2526 | | | | 145 | | | |

*Notes:* This table presents the Specification 2 estimation results in the panel and in the long-difference cross-section, using the 1%/99%-trimmed sample. The panel consists of 2,526 country-year observations on 163 countries over 2006–2022, and the long-difference consists of a 145-country cross-section comparing 2006–08 averages with 2020–22 averages. Identification uses NL as the reference signal, and $V$ has the structural off-diagonals $\sigma_{1,4}$, $\sigma_{1,5}$, and $\sigma_{4,5}$.

Adding the two energy-statistics signals to GDP, nighttime lights, and $NO_2$ lowers the BLUE minimum MSE substantially. The panel MSE falls from 13.52 in Specification 1 to 2.94 in Specification 2, a 78% reduction; the latent variance also adjusts, from $\widehat{\sigma}^2_{y^*} = 14.46$ to 8.40, with part of the difference reflecting the larger sample (163 vs. 146 countries) and part reflecting the structural off-diagonals $\sigma_{1,4}$ and $\sigma_{1,5}$ absorbing covariance that would otherwise enter $\sigma^2_{y^*}$. Of the four non-GDP signals, the two energy measurements have the highest panel SNRs by a wide margin: electricity's is 0.78 and $CO_2$'s is 0.72, against NL's 0.08 and $NO_2$'s 0.026. Effective contributions follow the same ordering, with $\lambda_5\gamma_5 = 0.357$ for electricity, $\lambda_4\gamma_4 = 0.330$ for $CO_2$, 0.028 for NL, and 0.009 for $NO_2$. Energy use is mechanically tied to year-on-year economic activity, while the satellite signals aggregate spatial information that captures only a fraction of the underlying flow.

The two satellite signals play an identification role rather than a predictive one. NL ($x_2$) and $NO_2$ ($x_3$) carry the orthogonality conditions that pin down $\gamma_2$, $\sigma^2_{y^*}$, and $\gamma_k$ for $k \geq 3$. Once identification is achieved, the structural off-diagonals $\sigma_{1,4} = -2.11$ and $\sigma_{1,5} = -1.76$ characterize how the measurement errors in GDP, $CO_2$, and electricity comove: when official GDP overstates growth, $CO_2$ and electricity look low relative to the latent factor, and the BLUE reallocates weight away from GDP and toward the energy signals. Without the two satellite signals, neither $\gamma$ nor the off-diagonals would be identifiable in the system.

Nonlinearity in the satellite and energy signals adds little beyond a single channel. Extending Specification 2 with second-order generated measurements – four squares and six pairwise cross terms of $(x_2, x_3, x_4, x_5)$ – produces a $K = 15$ polynomial sieve (Model B′) that lowers the BLUE MSE by a further 3%, from 2.94 to 2.85. Among the ten generated measurements, only $(\text{NL})^2$ contributes meaningfully, with $\lambda_6\gamma_6 = 0.089$ and SNR 0.21, an order of magnitude above any other generated term. The prominence of $(\mathrm{NL})^2$ again reflects the nonlinearity in the light–output relationship. Squared terms in $NO_2$, $CO_2$, and electricity each contribute under 0.005, and cross terms under 0.01. The full polynomial-sieve table is in Appendix A.

## 6.4 Country-Level Estimates

Armed with two specifications that incorporate satellite signals (Models A and A′) and traditional measurements (Models B and B′), we examine official GDP growth against the improved estimates. We focus on Model A, B, and B′, as Model B shows a substantial reduction in the BLUE MSE while the improvement from adding nonlinearity in Model A is similar to that in Model B.

Figure 2 plots the comparison for six stylized emerging economies. Table 4 reports the mean of each series and its volatility, measured as the within-country mean absolute year-on-year change in growth. Model A tracks official GDP closely at annual frequency, as the large BLUE weight on official GDP leaves limited room for the satellite signals to pull the estimate away. Models B and B′ move more independently: the energy signals can shift the BLUE substantially when $CO_2$ and electricity growth diverge from reported output, and they sometimes deviate meaningfully from official statistics. B and B′ are broadly similar to each other, with B′'s nonlinear term introducing modest additional adjustments.

Figure 2: Official GDP vs. BLUE Estimates Under the Three Specifications, 2006–2022

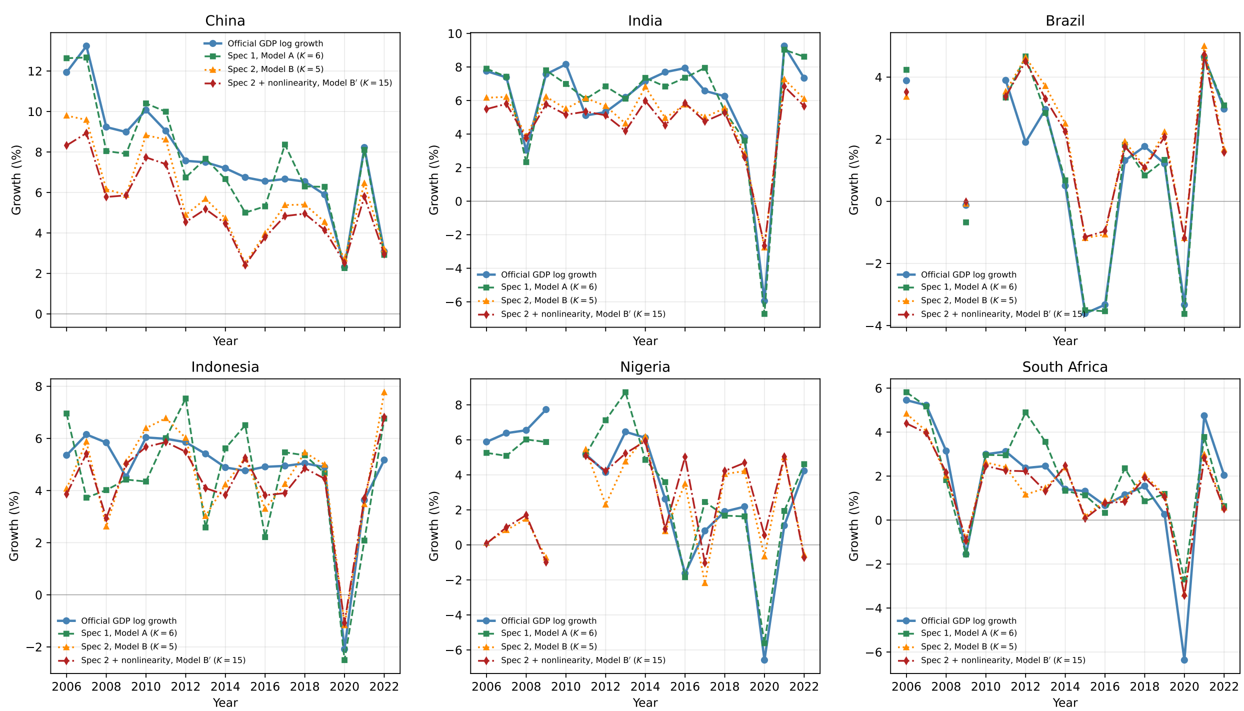


*Notes:* This figure plots official GDP log growth against the three BLUE estimates for six emerging economies over 2006–2022, using the 1%/99%-trimmed sample. The blue solid line is official GDP log growth, the green dashed line is Specification 1 (Model A, $K = 6$), the orange dotted line is Specification 2 (Model B, $K = 5$), and the red dash-dotted line is Specification 2 with nonlinearity (Model B′, $K = 15$). Years dropped by the trimming appear as breaks in the lines rather than being connected across: Brazil drops 2007, 2008, and 2010, whose tropospheric $NO_2$ growth lies in the extreme 1%/99% tails, and Nigeria drops 2010, whose $CO_2$ growth exceeds the 99% cutoff.

Table 4: Country-Level Mean Growth and Volatility, 2006–2022

| Country | Mean growth (%) | | | | Volatility, mean $\|\Delta g\|$ (pp) | | | |
|---|---|---|---|---|---|---|---|---|
| | Official | A ($K$=6) | B ($K$=5) | B′ ($K$=15) | Official | A ($K$=6) | B ($K$=5) | B′ ($K$=15) |
| China | 7.69 | 7.49 | 5.78 | 5.27 | 1.60 | 2.19 | 1.71 | 1.46 |
| India | 5.91 | 5.93 | 5.06 | 4.67 | 2.92 | 3.00 | 2.08 | 1.94 |
| Brazil | 1.05 | 1.15 | 1.87 | 1.78 | 2.91 | 3.08 | 2.46 | 2.35 |
| Indonesia | 4.78 | 4.46 | 4.56 | 4.35 | 1.24 | 2.55 | 2.18 | 1.67 |
| Nigeria | 3.32 | 3.53 | 2.16 | 2.56 | 2.62 | 2.58 | 3.54 | 3.22 |
| South Africa | 1.76 | 2.03 | 1.50 | 1.46 | 2.31 | 2.24 | 1.89 | 1.83 |

*Notes:* This table reports the country-level means and volatility of official GDP log growth and of the three BLUE predictors, computed across the available trimmed years for each country. Volatility is defined as the mean absolute year-on-year change in growth, $\frac{1}{T_i-1}\sum_t |g_{i,t} - g_{i,t-1}|$, in percentage points. This measures the typical magnitude of year-on-year shifts in the growth rate, not the dispersion of the growth level itself.

China and India's BLUE estimates lie systematically below official growth under Models B and B′. China's mean falls from official 7.69 to 5.78 under Model B and to 5.27 under Model B′, a downward correction of 1.9 to 2.4 percentage points. India's mean falls from official 5.91 to 5.06 and 4.67, a downward correction of 0.9 to 1.2 percentage points. These corrections are consistent with Nakamura, Steinsson, and Liu (2016), Hu and Yao (2022), and Martinez (2022), who document systematic over-reporting of growth in fast-growing emerging economies. Brazil moves in the opposite direction: its mean rises from official 1.05 to 1.87 under Model B and 1.78 under Model B′, an upward adjustment of about 0.7 to 0.8 pp. $CO_2$ and electricity grew more than the latent factor would predict given Brazil's reported GDP, so the BLUE reads the Brazilian series as understated over our sample window.

Volatility, measured as the mean absolute year-on-year change in growth $\frac{1}{T_i-1}\sum_t |g_{i,t} - g_{i,t-1}|$, reveals a sharper anomaly for Indonesia. Indonesia's official series has the lowest volatility among the six countries at 1.24 percentage points – a striking smoothness for a commodity-exporting emerging economy whose growth path one would expect to display year-on-year shocks from oil-price swings, monetary cycles, and the global financial crisis. Specification 1, which gives Indonesia's satellite signals their largest share of the BLUE, implies a volatility of 2.55 pp, more than double the official figure. The satellite signals tell a story of year-on-year activity that the official statistics filter out. Specifications 2 and 2-with-nonlinearity, which lean on the smoother energy signals, imply intermediate volatilities of 2.18 and 1.67 pp – still well above the official 1.24. The pattern is consistent with the long-standing observation that Indonesia's quarterly GDP is among the smoothest in our sample.

For the four other countries, Model B compresses volatility relative to official statistics. India's official volatility of 2.92 pp falls to around 2.08 pp, a reduction of roughly 30%. South Africa's 2.31 pp and Brazil's 2.91 pp each fall by 15–20%. China's volatility barely moves. Nigeria is the exception: its BLUE volatility under Model B (3.54 pp) exceeds the official 2.62 pp, reflecting oil-price-driven swings in $CO_2$ and electricity that do not appear in the reported growth series.

The two specifications illustrate the framework's flexibility. Specification 1, with only clean-error satellite signals, anchors the BLUE on official GDP at annual frequency but shifts decisively to urban land cover and nighttime lights at long horizons. Specification 2, by admitting the structural-error correlations between GDP and the energy signals, produces a meaningfully different growth series for China, India, and Brazil. The polynomial-sieve refinement of Specification 2-with-nonlinearity (Model B′) reinforces these corrections through the curvature of the nighttime light–output relationship.

### 6.5 External Validation: Model Disagreement and Official-Data Revisions

The country-level corrections invite an external check: are the countries whose estimates depart most from official GDP also those whose official statistics are independently regarded as least reliable? We compare the model's disagreement with official GDP to the revisions that the official growth figures themselves later received. For each country in year $t$, we take the revision of real GDP growth between its $t+1$ and $t+3$ vintages of the IMF World Economic Outlook (September/October version), $\text{rev}_{it}$, in percentage points, where a negative value means growth was later revised down; the revisions cover 2006–2022 for 180–193 countries per year.

If the gap $d_{it} = x_{1,it} - \hat{y}^*_{it}$ (official minus the Model B estimate) detected genuine over-reporting, over-reporting countries should subsequently be revised down, producing a negative association between $d_{it}$ and $\text{rev}_{it}$. We find no such relationship. Across the 163-country panel the cross-country correlation between the mean gap and the mean revision is weakly positive (Pearson +0.22, Spearman +0.14), and the moderately negative correlation visible among a few large emerging economies—including the six of Table 4—does not survive in the full sample. This is expected: vintage revisions reflect rebasing, methodological updates, and incoming source data rather than admissions of structural mismeasurement, so they do not identify the direction of bias—national accounts are seldom revised to concede inflated growth.

The picture changes once we compare magnitudes. Figure 3 plots, for each country, the average absolute gap $\frac{1}{T_i}\sum_t |x_{1,it} - \hat{y}^*_{it}|$ against the average absolute revision $\frac{1}{T_i}\sum_t |\text{rev}_{it}|$. The two are strongly related: Pearson $r = +0.57$ and Spearman $\rho = +0.45$ across 162 countries, with the rank correlation confirming that the pattern is not driven by a few outliers (scaling the revision by a country's average growth gives a similar +0.31). The economies where Model B departs most from official growth—Libya, Venezuela, Iraq, Equatorial Guinea, Timor-Leste, and Myanmar—are precisely those whose official figures are revised most, while the six stylized economies of Table 4 sit among the better-measured cases, with Nigeria the partial exception. The magnitude of the model's correction therefore tracks an independent proxy for the unreliability of a country's national accounts, even though the signed revision does not reveal which way that unreliability cuts.

Part of this magnitude relationship is a common-volatility effect: economies exposed to conflict, commodity-price swings, or hyperinflation exhibit both large model corrections and large data revisions. We therefore read Figure 3 as evidence that the framework flags the right countries as poorly measured—a credibility check on where the estimates move—rather than as a test of the level of bias, which the revisions cannot supply.

Figure 3: Model Disagreement vs. Official-Data Revisions, 2006–2022

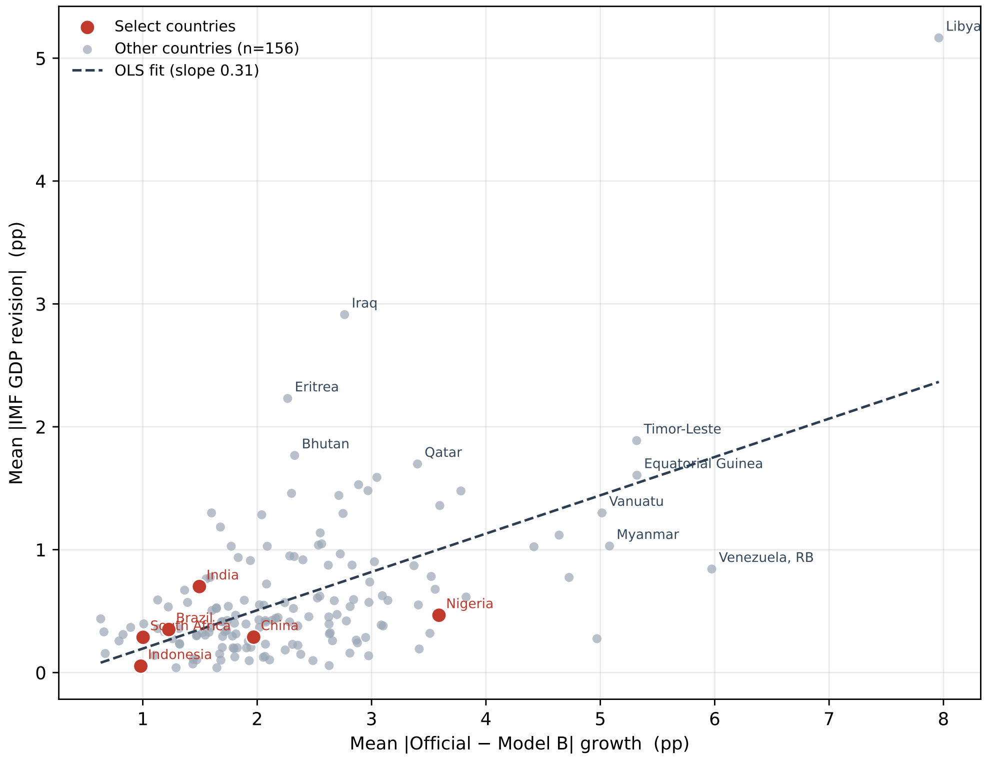


*Notes:* Each point is a country. The horizontal axis is the country's average absolute difference between official GDP log growth and the Model B (Specification 2) BLUE estimate, $\frac{1}{T_i}\sum_t |x_{1,it} - \hat{y}^*_{it}|$, in percentage points; the vertical axis is the average absolute IMF revision of official growth, $\frac{1}{T_i}\sum_t |\text{rev}_{it}|$, where $\text{rev}_{it}$ is the revision of year-$t$ growth between its $t+1$ and $t+3$ vintages. Red markers denote the six economies of Table 4; the dashed line is the OLS fit (slope 0.31). The sample is the 162 countries with at least five matched years. Pearson $r = +0.57$, Spearman $\rho = +0.45$.

# 7 Conclusion

This paper develops a general framework for approximating a latent variable using $K$ measurements, extending the two-signal model of Henderson, Storeygard, and Weil (2012) and the three-signal model of Civelli, Gaduh, and Yousuf (2024) in two important directions.

First, we generalize from three to $K$ signals. We derive two closed-form optimal linear combinations for $K$ signals with and without the unbiasedness constraint. We establish conditions under which the MSE converges to zero as $K$ grows: convergence obtains when the signal-to-noise ratios do not decay too fast, while the MSE remains bounded when they do.

Second, we incorporate nonlinear measurements. We show how nonlinear transformations of observables can serve as additional measurements, and prove that a polynomial sieve of generated measurements converges to the oracle bound. We establish that the total information content of generated measurements is bounded. This contrasts with independent signals, where the information sum can diverge and MSE can reach zero. The two types of measurements are complementary: independent signals drive the oracle bound toward zero, while generated measurements close the gap between the linear MSE and the oracle bound.

In our empirical application to measuring economic growth, we combine official GDP growth with a range of satellite signals and traditional energy statistics under two specifications—one pairing official GDP with satellite signals alone, the other adding $CO_2$ emissions and electricity consumption with correlated measurement errors—and extend each with a polynomial sieve. A key finding is that the informativeness of satellite data rises sharply with the time horizon. At annual frequency, official GDP is the least-noisy measure and dominates the combination, with nighttime lights the most useful of the satellite signals; over long differences, the satellite signals overtake official GDP, and urban land cover becomes the most informative signal, followed by nighttime lights.

Combining satellite signals with traditional energy statistics lowers the estimation error substantially and yields our best estimates, with the satellite signals providing identification and the energy statistics contributing most of the predictive content. The resulting BLUE estimates reveal substantial discrepancies between official and estimated growth: China and India's growth lies systematically below official statistics, Brazil's lies above, and Indonesia's official growth is excessively smooth relative to what the signals imply.

Our framework opens several avenues for future research. As new data sources—such as mobile phone activity, internet traffic, and high-frequency financial data—become increasingly available, the many-measurement approach developed here provides a principled method for combining them to better measure economic activity. The convergence results provide guidance on when additional signals are informative and when they are not. Extensions to panel data with time-varying parameters and to nonlinear aggregation methods are promising directions.

## References


Bai, J., and S. Ng (2002): "Determining the Number of Factors in Approximate Factor Models," *Econometrica*, 70(1), 191–221.

Beyer, R. C., Y. Hu, and J. Yao (2022): *Measuring quarterly economic growth from outer space*. International Monetary Fund.

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

Civelli, A., A. Gaduh, and H. Yousuf (2024): "Measuring Economic Growth With a Fully Identified Three-Signal Model," *Review of Economics and Statistics*, 106(4), 1126–1134.

Donaldson, D., and A. Storeygard (2016): "The View from Above: Applications of Satellite Data in Economics," *Journal of Economic Perspectives*, 30(4), 171–198.

Friedlingstein, P., M. Jones, M. O'Sullivan, R. Andrew, J. Hauck, G. Peters, W. Peters, J. Pongratz, S. Sitch, C. Le Quéré, et al. (2019): "Global Carbon Budget 2019," *Earth System Science Data*, 11(4), 1783–1838.

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

Hu, Y. (2008): "Identification and Estimation of Nonlinear Models with Misclassification Error Using Instrumental Variables: A General Solution," *Journal of Econometrics*, 144, 27–61.

——— (2017): "The Econometrics of Unobservables: Applications of Measurement Error Models in Empirical Industrial Organization and Labor Economics," *Journal of Econometrics*, 200, 154–168.

Hu, Y., and S. Schennach (2008): "Instrumental Variable Treatment of Nonclassical Measurement Error Models," *Econometrica*, 76, 195–216.

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

Kotlarski, I. (1965): "On Pairs of Independent Random Variables Whose Product Follows the Gamma Distribution," *Biometrika*, pp. 289–294.

Li, T. (2002): "Robust and Consistent Estimation of Nonlinear Errors-in-variables Models," *Journal of Econometrics*, 110(1), 1–26.

Martinez, L. R. (2022): "How much should we trust the dictator's GDP growth estimates?," *Journal of Political Economy*, 130(10), 2731–2769.

Nakamura, E., J. Steinsson, and M. Liu (2016): "Are Chinese growth and inflation too smooth? Evidence from Engel curves," *American Economic Journal: Macroeconomics*, 8(3), 113–144.

Rao, B. P. (1992): *Identifiability in Stochastic Models: Characterization of Probability Distributions*. Academic Press.

Schennach, S. M. (2004): "Estimation of Nonlinear Models with Measurement Error," *Econometrica*, 72(1), 33–75.

——— (2019): "Mismeasured and Unobserved Variables," *Handbook of Econometrics*.

Staiger, D., and J. H. Stock (1997): "Instrumental Variables Regression with Weak Instruments," *Econometrica*.

Stock, J., and M. Watson (2002): "Forecasting using Principal Components from a Large Number of Predictors," *Journal of the American Statistical Association*, 97(460), 1167–1179.

Stock, J. H., and J. H. Wright (2000): "GMM with Weak Identification," *Econometrica*, 68(5), 1055–1096.

# A Detailed Empirical Setup

This appendix collects the measurement equation, summary statistics, structural-zero patterns of the $V$ matrix, identification formulas, and the polynomial-sieve estimation table referenced in Section 6. Primary-source details for each signal are in Section 6.1.

## A.1 Measurement equation and signal indexing

Each signal admits the linear-projection representation

$$x_{k,it} \;=\; \gamma_k \, y^*_{it} + \varepsilon_{k,it}, \qquad \operatorname{Cov}(y^*_{it},\, \varepsilon_{k,it}) = 0, \qquad \gamma_1 = 1, \tag{13}$$

where $y^*_{it}$ is true latent GDP log growth in country $i$ at year $t$. Specification 1 indexes the six signals as $x_1 =$ GDP, $x_2 =$ NL, $x_3 = NO_2$, $x_4 =$ urban land cover, $x_5 =$ NDVI level change, $x_6 =$ NDBI level change. Specification 2 indexes the five signals as $x_1 =$ GDP, $x_2 =$ NL, $x_3 = NO_2$, $x_4 = CO_2$, $x_5 =$ electricity.

## A.2 Summary statistics

Table 5 reports the distribution of the eight candidate signals over the merged 2006–2022 panel, before any trimming.

Table 5: Summary Statistics of the Eight Candidate Signals

| | Signal | $N$ | Mean | SD | p10 | Median | p90 |
|---|---|---|---|---|---|---|---|
| $x_1$ | GDP log growth (%) | 3,540 | 2.79 | 6.15 | -2.70 | 3.28 | 7.73 |
| $x_2$ | Nighttime light log growth (%) | 3,144 | 7.34 | 34.05 | -13.70 | 4.80 | 31.97 |
| $x_3$ | $NO_2$ log growth (%) | 3,006 | -0.0403 | 11.73 | -11.25 | -0.3419 | 10.20 |
| $x_4$ | Urban land cover log growth (%) | 2,567 | 2.76 | 4.37 | 0.000000 | 1.59 | 6.40 |
| $x_5$ | NDVI level change | 3,094 | 0.000441 | 0.0162 | -0.0183 | 0.000121 | 0.0191 |
| $x_6$ | NDBI level change | 3,849 | -0.000198 | 0.0235 | -0.0246 | -0.000289 | 0.0237 |
| $x_7$ | $CO_2$ log growth (%) | 3,623 | 2.06 | 11.20 | -8.92 | 1.61 | 13.48 |
| $x_8$ | Electricity log growth (%) | 3,601 | 3.03 | 8.75 | -3.68 | 2.26 | 10.66 |

*Notes:* This table reports the summary statistics of the eight candidate signals, computed from the pooled distribution of each signal across the 2006–2022 country-year sample before trimming. GDP, NL, $NO_2$, urban land cover, $CO_2$, and electricity are reported in percent log growth, while NDVI and NDBI are unitless level changes. The cross-sectional standard deviations of the satellite-based growth signals – nighttime lights (34.0), $CO_2$ (11.2), $NO_2$ (11.7) – are substantially larger than that of official GDP log growth (6.2), reflecting both larger underlying $\gamma_k$ and noisier measurement processes. Urban land cover has a positive median (1.59) consistent with steady global urbanization over our sample, while NDVI and NDBI have median changes near zero, consistent with slow-moving stock measurements.

## A.3 Variance-covariance structure and identification

**Specification 1 (Model A).** Since $\varepsilon_k \perp \varepsilon_{\text{GDP}}$ for $k = 2, \ldots, 6$ and NDVI and NDBI share the MODIS spectroradiometer, the structural-zero pattern of $V$ is

$$V_{\text{Spec 1}} = \begin{pmatrix} \sigma_{11} & 0 & 0 & 0 & 0 & 0 \\ 0 & \sigma_{22} & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma_{55} & \sigma_{56} \\ 0 & 0 & 0 & 0 & \sigma_{65} & \sigma_{66} \end{pmatrix}.$$

With GDP as the fully-uncorrelated reference, identification follows the closed forms

$$\gamma_2 = \frac{\text{Cov}(x_2, x_3)}{\text{Cov}(x_1, x_3)}, \qquad \sigma^2_{y^*} = \frac{\text{Cov}(x_1, x_2)}{\gamma_2}, \tag{14}$$

$$\gamma_k = \frac{\text{Cov}(x_k, x_1)}{\sigma^2_{y^*}}, \qquad k = 3, 4, 5, 6. \tag{15}$$

**Specification 1 with nonlinearity (Model A′).** Dropping the uninformative NDVI and NDBI signals from Specification 1 and adding three squared terms $(x_2 - \bar{x}_2)^2$, $(x_3 - \bar{x}_3)^2$, $(x_4 - \bar{x}_4)^2$ and three pairwise cross terms of $(x_2, x_3, x_4) = (\text{NL}, \text{NO}_2, \text{Urban})$ produces a $K = 10$ system. Within the generated block, all pairwise off-diagonals are kept; within each base-signal set $\{x_k, (\text{base})^2, \text{base}\cdot\text{other}\}$, off-diagonals are open. Since every satellite error – and every product of satellite errors – is uncorrelated with the GDP error, GDP serves as the universal reference for $\gamma_k$ identification. Table 6 reports the full $K = 10$ result.

Table 6: Specification 1 with Polynomial Sieve (Model A′, $K = 10$)

| | Description | $\gamma_k$ | $\sigma^2_k$ | Signal SD | Noise SD | SNR | $\lambda_k$ | $\lambda_k\gamma_k$ |
|---|---|---|---|---|---|---|---|---|
| $x1$ | GDP | 1.000 | 12.474 | 4.125 | 3.532 | 1.364 | 0.872 | 0.872 |
| $x2$ | NL | 1.200 | 305.10 | 4.951 | 17.467 | 0.080 | 0.060 | 0.072 |
| $x3$ | $\text{NO}_2$ | 0.2308 | 61.915 | 0.9519 | 7.869 | 0.015 | 0.039 | 0.009 |
| $x4$ | Urban | 0.1070 | 9.750 | 0.4413 | 3.123 | 0.020 | 0.287 | 0.031 |
| $x5$ | $(\text{NL})^2$ | −13.562 | 406,606 | 55.937 | 637.66 | 0.008 | −0.001 | 0.014 |
| $x6$ | $(\text{NO}_2)^2$ | −0.4879 | 8242.49 | 2.012 | 90.788 | 0.000 | −0.000 | 0.000 |
| $x7$ | $(\text{Urban})^2$ | 0.1222 | 949.70 | 0.5039 | 30.817 | 0.000 | −0.022 | −0.003 |
| $x8$ | $\text{NL}\cdot\text{NO}_2$ | −1.235 | 22,935 | 5.094 | 151.44 | 0.001 | −0.001 | 0.001 |
| $x9$ | NL·Urban | 0.0100 | 3010.84 | 0.0413 | 54.871 | 0.000 | 0.001 | 0.000 |
| $x10$ | $\text{NO}_2\cdot\text{Urban}$ | −0.4319 | 505.46 | 1.781 | 22.483 | 0.006 | −0.009 | 0.004 |
| | | | | | Total $\lambda^T\gamma$: | | | 1.000 |

*Notes:* This table presents the Specification 1 estimation results with a polynomial sieve (Model A′, $K = 10$), using the 1%/99%-trimmed panel of 2,329 country-year observations on 148 countries. Identification uses GDP as the universal reference. The estimated latent variance is $\hat{\sigma}^2_{y^*} = 17.01$ and the BLUE MSE is 10.88, versus 11.51 for the linear $K = 4$ baseline on the same sample, a 5.4% reduction.

**Specification 2 (Model B).** $CO_2$ and electricity share data inputs with national-accounting systems, so $\sigma_{1,4}, \sigma_{1,5}, \sigma_{4,5} \neq 0$. Nighttime lights are uncorrelated with every other measurement error. The structural-zero pattern of $V$ is

$$V_{\text{Spec 2}} = \begin{pmatrix} \sigma_{11} & 0 & 0 & \sigma_{14} & \sigma_{15} \\ 0 & \sigma_{22} & 0 & 0 & 0 \\ 0 & 0 & \sigma_{33} & 0 & 0 \\ \sigma_{41} & 0 & 0 & \sigma_{44} & \sigma_{45} \\ \sigma_{51} & 0 & 0 & \sigma_{54} & \sigma_{55} \end{pmatrix}.$$

With NL ($x_2$) as the reference,

$$\gamma_2 = \frac{\mathrm{Cov}(x_2, x_3)}{\mathrm{Cov}(x_1, x_3)}, \qquad \sigma^2_{y^*} = \frac{\mathrm{Cov}(x_1, x_2)}{\gamma_2}, \tag{16}$$

$$\gamma_k = \frac{\mathrm{Cov}(x_k, x_2)}{\mathrm{Cov}(x_1, x_2)}, \qquad k = 3, 4, 5. \tag{17}$$

Once $\gamma$ and $\sigma^2_{y^*}$ are identified, the structural off-diagonals are recovered as $\sigma_{j,k} = \mathrm{Cov}(x_j, x_k) - \gamma_j \gamma_k \sigma^2_{y^*}$ at the four allowed positions.

**Specification 2 with nonlinearity (Model B′).** Adding four squared terms $(x_2 - \bar{x}_2)^2, \ldots, (x_5 - \bar{x}_5)^2$ and six pairwise cross terms produces a $K = 15$ system. Within the generated block, all pairwise off-diagonals are kept; $\sigma_{1,k} \neq 0$ for every base or generated term that involves $CO_2$ or electricity. We use a mixed-reference rule in which the reference signal for each $\gamma_k$ is chosen so that $\varepsilon_{\text{ref}}$ is uncorrelated with $\varepsilon_k$: nighttime lights are the default; for generated measurements that involve nighttime lights, we use $NO_2$; for the NL·$NO_2$ cross term, we use $CO_2$. Table 7 reports the full $K = 15$ result.

Table 7: Specification 2 with Polynomial Sieve (Model B$'$, $K = 15$)

| | Description | $\gamma_k$ | $\sigma_k^2$ | Signal SD | Noise SD | SNR | $\lambda_k$ | $\lambda_k\gamma_k$ |
|---|---|---|---|---|---|---|---|---|
| $x1$ | GDP | 1.000 | 13.499 | 2.898 | 3.674 | 0.622 | 0.266 | 0.266 |
| $x2$ | NL | 1.689 | 300.56 | 4.894 | 17.337 | 0.080 | 0.007 | 0.012 |
| $x3$ | $NO_2$ | 0.4244 | 59.015 | 1.230 | 7.682 | 0.026 | 0.022 | 0.009 |
| $x4$ | $CO_2$ | 1.995 | 46.599 | 5.782 | 6.826 | 0.717 | 0.144 | 0.287 |
| $x5$ | Electricity | 1.401 | 21.275 | 4.062 | 4.612 | 0.775 | 0.224 | 0.314 |
| $x6$ | $(\mathrm{NL})^2$ | 95.336 | 358,047 | 276.29 | 598.37 | 0.213 | 0.001 | 0.089 |
| $x7$ | $(NO_2)^2$ | 1.396 | 7772.90 | 4.045 | 88.164 | 0.002 | 0.001 | 0.001 |
| $x8$ | $(CO_2)^2$ | 2.544 | 19,592 | 7.371 | 139.97 | 0.003 | 0.001 | 0.002 |
| $x9$ | $(\mathrm{Elec})^2$ | 3.128 | 5827.73 | 9.066 | 76.340 | 0.014 | 0.001 | 0.003 |
| $x10$ | NL·$NO_2$ | −0.9028 | 22,331 | 2.617 | 149.44 | 0.000 | −0.001 | 0.001 |
| $x11$ | NL·$CO_2$ | −3.710 | 32,398 | 10.751 | 179.99 | 0.004 | −0.001 | 0.004 |
| $x12$ | NL·Elec | −3.461 | 16,024 | 10.030 | 126.59 | 0.006 | −0.002 | 0.006 |
| $x13$ | $NO_2$·$CO_2$ | −0.9323 | 5568.32 | 2.702 | 74.621 | 0.001 | 0.002 | −0.002 |
| $x14$ | $NO_2$·Elec | −0.8697 | 2635.63 | 2.521 | 51.338 | 0.002 | −0.002 | 0.001 |
| $x15$ | $CO_2$·Elec | −1.203 | 4740.44 | 3.488 | 68.851 | 0.003 | −0.005 | 0.006 |
| | | | | | | Total $\lambda^T\gamma$: | | 1.000 |

*Notes:* This table presents the Specification 2 estimation results with a polynomial sieve (Model B$'$, $K = 15$), using the 1%/99%-trimmed panel of 2,526 country-year observations on 163 countries. Identification uses the mixed-reference rule. The estimated latent variance is $\widehat{\sigma}^2_{y^*} = 8.40$ and the BLUE MSE is 2.85.