
Measuring Economic Growth with Many Signals

Yingyao Hu • Jiaxiong Yao

Johns Hopkins University • International Monetary Fund

October 2, 2026 University of Southern California

The views expressed are those of the authors and do not necessarily reflect those of the IMF, its Board of Directors, or the countries they represent.

The problem is no longer data scarcity

Signals of economic activity have multiplied

- ▶ Remote sensing: nighttime lights, NO₂, urban land cover, NDVI, NDBI
- ▶ Energy & emissions: CO₂, electricity
- ▶ Digital footprints: card transactions, search, mobile phones, ship tracking, news

Official GDP is itself measured with error

- ▶ Informal activity, statistical capacity, revisions, incentives to misreport

The question

How should many noisy, overlapping signals be combined — alongside official GDP — to recover true economic growth?

Two sub-questions

1. Does adding more signals always help?
2. What about nonlinear content in the signals?

Where the literature stands

	HSW (2012)	CGY (2024)	This paper
Signals	$K = 2$: GDP + lights	$K = 3$: + urban land	$K \geq 3$, arbitrary
Identification	Structural relationship under-identified; additional assumption on signal-to-noise ratio is needed	Exactly identified	Identified for any K
Error structure	Independent	Independent	General V (correlated errors)
Estimator	Weighted average	BLP weights	BLP and BLU
Large K	—	—	Convergence theory
Nonlinearity	—	—	Generated measurements

BLP: best linear predictor, which minimizes the MSE and trades some bias for lower variance.

BLU: best linear unbiased combination, which minimizes the MSE subject to unbiasedness.

Related: Hu & Yao (2022) imposes structure on the errors; Beyer et al. (2022) leverages heteroskedasticity. The measurement-error literature (Kotlarski; Hu & Schennach) identifies distributions — we target the value in each observation.

What we do

Theory

1. **Closed forms for K signals.** BLP and BLU combinations; CGY's three-signal formula is a special case.
2. **Convergence as $K \rightarrow \infty$.** $\text{MSE} \rightarrow 0$ if (1) signal strength accumulates across measurements, or (2) the errors in the signal-weighted portfolio become increasingly offsetting so that the aggregate noise decreases.
3. **Nonlinear generated measurements.** Transforms of existing observables are valid signals, but there is an information limit.
4. **Complementarity.** Generated and independent signals are complements.

Empirics: 8 signals, 146–163 countries, 2006–2022

- ▶ Satellite signals are far more informative at long horizons
- ▶ CO_2 + electricity cut the annual MSE by 84%
- ▶ Nonlinearity lowers the MSE further; only $(\text{NL})^2$ carries meaningful signal
- ▶ China and India lie below official growth, Brazil above, Indonesia is too smooth
- ▶ Where we disagree most is where the IMF later revises most

PART 1

A Many-Signal Model

A Many-Signal Model

Convergence and Estimation

Nonlinear Measurements

Simulations

Measuring Economic Growth

Setup and identification

Measurement equations

$$x_k = \gamma_k x^* + \varepsilon_k, \quad \text{COV}(x^*, \varepsilon_k) = 0,$$

$k = 1, \dots, K$ with $K \geq 3$; i.i.d. cross-section.

- ▶ These are linear projections — uncorrelatedness is not restrictive
- ▶ Errors may be correlated across signals and have nonzero means
- ▶ Classical measurement error is the special case $\gamma_k = 1, \varepsilon_k \perp x^*$

$$\gamma_2 = \frac{\text{COV}(x_2, x_3)}{\text{COV}(x_1, x_3)}, \quad \sigma_{x^*}^2 = \frac{\text{COV}(x_1, x_2) \text{COV}(x_1, x_3)}{\text{COV}(x_2, x_3)}, \quad \gamma_k = \frac{\text{COV}(x_k, x_1)}{\sigma_{x^*}^2}, \quad k \geq 3.$$

Assumption 1 (Three-signal core)

$\gamma_1, \gamma_2, \gamma_3 \neq 0$ with
 $\text{COV}(\varepsilon_1, \varepsilon_2) = \text{COV}(\varepsilon_1, \varepsilon_3) = \text{COV}(\varepsilon_2, \varepsilon_3) = 0$, and
 $\text{COV}(\varepsilon_1, \varepsilon_k) = 0$ for all $k \geq 4$.

Assumption 2 (Normalization & rank)

$\gamma_1 = 1, \mathbb{E}[\varepsilon_1] = 0$; $V(\varepsilon)$ has full rank.

To add a $(K+1)$ th signal, all we need is
 $\text{COV}(\varepsilon_1, \varepsilon_{K+1}) = 0$. It may be correlated with everything else.

Proposition 1: two optimal linear combinations

Consider $\hat{x}^*(\lambda) = \lambda_0 + \lambda_1^T x$, with λ_0 set so that $\mathbb{E}[\hat{x}^*] = \mathbb{E}[x_1]$.

BLP (bias-allowing)

minimizes $\mathbb{E}[\hat{x}^*(\lambda) - x^*]^2$ over all λ

$$\lambda_1^{\text{BLP}} = \frac{\sigma_{x^*}^2 V^{-1} \gamma}{1 + \sigma_{x^*}^2 \gamma^T V^{-1} \gamma}$$
$$\text{MSE} = \frac{1}{1/\sigma_{x^*}^2 + \gamma^T V^{-1} \gamma}$$

BLU (conditionally unbiased)

minimizes over $\Lambda : \lambda_1^T \gamma = 1$

$$\lambda_1^{\text{BLU}} = \frac{V^{-1} \gamma}{\gamma^T V^{-1} \gamma}$$
$$\text{MSE} = \frac{1}{\gamma^T V^{-1} \gamma}$$

- ▶ Both need only first and second moments; V is recovered structurally from $\sigma_{jk} = \text{cov}(x_j, x_k) - \gamma_j \gamma_k \sigma_{x^*}^2$
- ▶ BLP dominates BLU: it buys a small bias for a larger variance reduction; the gap is $1/\sigma_{x^*}^2$ in precision and vanishes as $\sigma_{x^*}^2 \rightarrow \infty$
- ▶ BLU remains useful when conditional unbiasedness is needed for interpretation or for downstream inference

Three remarks: special cases, SNR, and the choice of signals

Remark 1: CGY is a special case

Set $K = 3$ and $V = \text{diag}(\sigma_1^2, \sigma_2^2, \sigma_3^2)$. The weight on official GDP is

$$\lambda_{1,1}^{\text{BLP}} = \frac{1}{1 + \frac{\sigma_1^2}{\sigma_{x^*}^2} + \frac{\sigma_1^2}{\sigma_2^2} \gamma_2^2 + \frac{\sigma_1^2}{\sigma_3^2} \gamma_3^2}$$

which is exactly CGY's equation (11). The BLU counterpart simply drops $\sigma_1^2/\sigma_{x^*}^2$; the two coincide as $\sigma_{x^*}^2 \rightarrow \infty$. HSW's equation (8) is a special case when $\gamma_3 = 0$.

Remark 3: the three signals must be chosen a priori

Proposition 1 requires knowing which three signals have mutually uncorrelated errors, and which has an error uncorrelated with all others. This rests on how the signals are produced; if it fails, the estimates are inconsistent.

Remark 2: signal-to-noise ratio

$$\text{SNR}_k = \frac{\gamma_k^2 \sigma_{x^*}^2}{\sigma_k^2}$$

With diagonal V , $\lambda_{1,k}^{\text{BLU}} \gamma_k = \text{SNR}_k / \sum_j \text{SNR}_j$

- ▶ A signal's contribution is its SNR (not just γ_k)
- ▶ Pure noise ($\gamma_{K+1} = 0$) gets weight zero: adding useless signals is harmless
- ▶ Scale invariance: $x_k \mapsto c x_k$ leaves $\gamma^T V^{-1} \gamma$, the weights, and the prediction unchanged

PART 2

Convergence and Estimation

A Many-Signal Model

Convergence and Estimation

Nonlinear Measurements

Simulations

Measuring Economic Growth

When do more signals help?

Eigenvalues of V : $\varsigma_1^2 \geq \dots \geq \varsigma_K^2 > 0$. Noise in the signal direction: $\varsigma_\gamma^2 := \frac{\gamma^T V \gamma}{\gamma^T \gamma} \in [\varsigma_K^2, \varsigma_1^2]$.

Assumption 3 (directional noise)

$\varsigma_\gamma^2 = O(K^{-\beta})$, $\beta \geq 0$.

$\beta > 0$: errors in the signal-weighted portfolio increasingly offset.

Assumption 4 (signal strength)

$c_1 k^{-\alpha} \leq |\gamma_k| \leq c_2 k^{-\alpha}$ for large k , $\alpha \geq 0$.
 α governs how fast a new signal's information dries up.

Proposition 2 (Convergence)

$$\mathbb{E}[\hat{x}^*(\lambda^{\text{BLU}}) - x^*]^2 = O\left(\frac{K^{-\beta}}{\sum_{k=1}^K k^{-2\alpha}}\right)$$

$$\beta > 0 \quad \longrightarrow \quad 0$$

$$\beta = 0, \alpha < 1/2 \quad \longrightarrow \quad 0$$

$$\beta = 0, \alpha = 1/2 \quad \longrightarrow \quad 0, \text{ at rate } 1/\log K$$

$$\alpha > 1/2, \varsigma_K^2 \geq \varsigma^2 > 0 \quad \longrightarrow \quad \text{bounded away from 0}$$

- ▶ With $\beta = 0$, recovery needs $\sum_k \gamma_k^2 = \infty$, i.e., $\alpha \leq 1/2$; with $\beta > 0$ it can occur even when $\sum_k \gamma_k^2 < \infty$
- ▶ Same upper bound for the BLP

Estimation: how many signals can a sample support?

Plug-in estimator: $\hat{\lambda}_N = G(\hat{\mu})$, with $\hat{\mu}$ the vector of $K(K+3)/2$ sample moments. Screened version: use x_1 alone when the identifying covariances are near zero, $\hat{\sigma}_{x^*}^2$ leaves a fixed interval, or $\varsigma_K^2(\hat{V}) < \nu_K/2$, where $\nu_K := \min\{1, \varsigma_K^2\}$.

Proposition 3 (Asymptotics)

Under Assumptions 1–5, if $r_{K,N} := K^3 \log K / (N \nu_K^3) \rightarrow 0$, the screened estimator satisfies

$$\mathbb{E}|\hat{x}^*(\hat{\lambda}_N^{\text{BLU}}) - x^*|^2 = \underbrace{O(r_{K,N})}_{\text{estimation error}} + \underbrace{O\left(\frac{K^{-\beta}}{\sum_{k=1}^K k^{-2\alpha}}\right)}_{\text{approximation error}}$$

The approximation error behaves as in Proposition 2.

- ▶ K must grow more slowly than $N^{1/3}$ — much slower than the $O(N)$ allowed in the many-instruments literature — and more slowly still when V is near singular
- ▶ The screen is asymptotically inactive: with probability $\rightarrow 1$ the estimator is the ordinary plug-in

PART 3

Nonlinear Measurements

A Many-Signal Model

Convergence and Estimation

Nonlinear Measurements

Simulations

Measuring Economic Growth

Nonlinearity hides in the projection residual

If the true conditional mean is nonlinear,
 $\mathbb{E}[x_k|x^*] = m_k(x^*)$, the linear projection residual is

$$\varepsilon_k = \underbrace{(m_k(x^*) - \gamma_k x^*)}_{\text{nonlinear signal}} + \underbrace{u_k}_{\text{pure noise}}$$

- ▶ A linear combination throws this away
- ▶ But it is still in the observables — transform them to get it back

A concrete transformation

$$h(x_2, x_3) = \left(\frac{x_2}{\gamma_2} - \frac{x_3}{\gamma_3} \right)^\delta = \left(\frac{\varepsilon_2}{\gamma_2} - \frac{\varepsilon_3}{\gamma_3} \right)^\delta$$

a nonlinear function of exactly the residuals that carry the missing signal.

What makes $x_a \equiv h(x_2, x_3)$ admissible

Strengthen the reference-signal condition to conditional mean independence, $\mathbb{E}[\varepsilon_1 | x^*, \varepsilon_2, \varepsilon_3] = \mathbb{E}[\varepsilon_1]$. Since $x_a = \tilde{h}(x^*, \varepsilon_2, \varepsilon_3)$, this delivers $\text{cov}(\varepsilon_1, \varepsilon_a) = 0$ — exactly what Assumption 1 requires.

Corollary 1

Proposition 1 applies verbatim to (x_1, x_2, x_3, x_a) , for any h with $\mathbb{E}x_a^2 < \infty$, provided the augmented V has full rank.

ε_a is generically correlated with $\varepsilon_2, \varepsilon_3$ — which is fine: the general V of Proposition 1 absorbs it.

A polynomial sieve reaches the oracle bound

Sieve construction

Add bivariate monomials in (x_2, x_3) ,
beyond the constant and linear terms:

$$d=2 : x_2^2, x_2x_3, x_3^2$$

$$d=3 : x_2^3, x_2^2x_3, x_2x_3^2, x_3^3$$

$$d : x_2^j x_3^{d-j}, \quad j = 0, \dots, d$$

$\Rightarrow K_d = (d+1)(d+2)/2 - 3$ generated
signals.

The oracle bound

$O := \mathbb{E}[\text{Var}(x^* \mid x_1, x_2, x_3)]$ is the
minimum MSE from these observables by
any predictor.

Proposition 4 (Sieve convergence)

If polynomials are dense in $L^2(x_2, x_3)$ and
 $\mathbb{E}[x^* \mid x_1, x_2, x_3] = \beta x_1 + g(x_2, x_3)$ with $0 < O < \sigma_{x^*}^2$, then

$$\lim_{d \rightarrow \infty} \mathbb{E}[\hat{x}^*(\lambda_{K_d}^{\text{BLP}}) - x^*]^2 = O$$
$$\lim_{d \rightarrow \infty} \mathbb{E}[\hat{x}^*(\lambda_{K_d}^{\text{BLU}}) - x^*]^2 = \frac{\sigma_{x^*}^2 O}{\sigma_{x^*}^2 - O} > O$$

- ▶ The BLP attains the oracle bound: a linear combination of nonlinear transforms recovers the full nonlinear predictor
- ▶ The BLU cannot — unbiasedness ($\lambda_1^T \gamma = 1$) forbids the bias–variance trade, so it converges strictly above O
- ▶ Linear-only MSE $\sigma_{x^*}^2 / (1 + \sigma_{x^*}^2 \gamma^T V^{-1} \gamma)$ minus O is the nonlinear information

When does nonlinearity matter — and how much is there?

Two channels

If (x^*, x_2, x_3) are jointly Gaussian and the measurement equations are linear, the gap is zero and generated signals are worthless.

- (i) **Nonlinear measurement equations.** Light sensors saturate; emissions links are logarithmic
- (ii) **Non-Gaussian latent variable.** If growth is right-skewed, $\mathbb{E}[x^*|x_k]$ is nonlinear

But generated information is bounded

In an orthonormal basis $\{\phi_k\}$, Parseval and the law of total variance give

$$\sum_{k \geq 1} \gamma_{a,k}^2 = \frac{\text{Var}(\mathbb{E}[x^*|x_2, x_3])}{\sigma_{x^*}^4} \leq \frac{1}{\sigma_{x^*}^2} < \infty$$

- So $\gamma_{a,k} = o(k^{-1/2})$: generated signals sit in the $\alpha > 1/2$ regime
- They cannot drive MSE to zero — all are functions of the same observables

Complementarity

Independent signals push the oracle bound down; generated signals close the gap to it. Most effective approach: many independent signals augmented with nonlinear interactions.

PART 4

Simulations

A Many-Signal Model

Convergence and Estimation

Nonlinear Measurements

Simulations

Measuring Economic Growth

Simulation design

Each uses $N = 5,000$ observations and an 80/20 train–test split. Throughout, γ and V are estimated from cross-covariances of the observables alone (not from x^*).

1. Slow SNR decay

$x_k = \gamma_k x^* + \varepsilon_k$, with
 $x^* \sim N(0, 1)$ and $\varepsilon_k \stackrel{iid}{\sim} N(0, 1)$

$$\gamma_k = k^{-1/2}$$
$$\sum_{k \leq K} \frac{\gamma_k^2}{\sigma_k^2} = H_K \rightarrow \infty$$

2. Fast SNR decay

Same structure, but the loadings die off faster

$$\gamma_k = k^{-1}$$
$$\sum_{k \leq K} \frac{\gamma_k^2}{\sigma_k^2} \rightarrow \frac{\pi^2}{6}$$

3. Nonlinear base signals

$x^* \sim \text{LogNormal}(0, 1)$, so right-skewed




$$x_1 = x^* + \varepsilon_1, \quad \varepsilon_1 \sim N(0, 2.25)$$

$$x_2 = \frac{1}{3}\sqrt{x^*} + \varepsilon_2, \quad \varepsilon_2 \sim N(0, 0.0225)$$

$$x_3 = \frac{1}{2} \log x^* + \varepsilon_3, \quad \varepsilon_3 \sim N(0, 0.09)$$

then add monomials $x_2^j x_3^l$,
 $j+l = 2, 3, 4$

Three simulations confirm the theory



(a) Slow decay, $\gamma_k = k^{-1/2}$

$\sum_k \gamma_k^2 = H_K \rightarrow \infty$, so MSE
 $= 1/H_K \rightarrow 0$.

(b) Fast decay, $\gamma_k = k^{-1}$

$\sum_k \gamma_k^2 \rightarrow \pi^2/6$, so MSE
 $\rightarrow 0.608 > 0$: more does not help.

(c) Generated nonlinear signals

MSE 1.11 $\rightarrow$ 0.78 once x_2^2 enters,
 then flat above the oracle bound 0.63.

$N = 5,000$; 80/20 train-test split; (a)–(b) averaged over 50 Monte Carlo replications.

PART 5

Measuring Economic Growth

A Many-Signal Model


Convergence and Estimation

Nonlinear Measurements

Simulations

Measuring Economic Growth

What the signals look like



United States, 2021

East Coast and Midwest.

- ▶ Lights capture the extent and intensity of activity; urban land cover sees only the built-up footprint. Both visibly have a high SNR
- ▶ NO₂ is blocky whereas NDVI is close to uniform. NO₂ has a higher SNR, but both are less informative

Data: eight candidate signals, 2006–2022

	Signal	Source	Transformation
x_1	Official real GDP	World Bank WDI (PPP, 2017 int. \$)	log growth
x_2	Nighttime lights	SVNL harmonized DMSP–VIIRS, 500 m (Chen et al. 2024)	log growth
x_3	NO ₂ tropospheric column	NASA OMI/Aura OMNO2d, 0.25° daily grids	log growth
x_4	Urban land cover	Copernicus C3S land cover, 300 m	log change in share
x_5	NDVI	MODIS MOD13A1, 500 m	level change
x_6	NDBI	MODIS MOD09A1, 500 m	level change
x_7	CO ₂ emissions	Global Carbon Project / OWID	log growth
x_8	Electricity consumption	Ember + EIA / OWID	log growth

Specification 1 — Model A ($K = 6$)

GDP + five satellite signals, with $\sigma_{GDP,k} = 0$ and $\sigma_{25}, \sigma_{26}, \sigma_{56}$ free (NL, NDVI, NDBI share calibration errors).

Isolates what satellites know.

Specification 2 — Model B ($K = 5$)

GDP, NL, NO₂ + two traditional energy statistics, with $\sigma_{NL,k} = 0$ and $\sigma_{14}, \sigma_{15}, \sigma_{45}$ free; (NL is the reference.) How far correlated-error signals go.

Each is also extended with a polynomial sieve ($A', K=10$; $B', K=15$).

Specification 1: satellites at annual vs. long horizons

		Panel, 2006–2022			Long-difference		
	Signal	γ_k	SNR	$\lambda_k \gamma_k$	γ_k	SNR	$\lambda_k \gamma_k$
x_1	GDP	1.000	0.929	0.870	1.000	0.263	0.079
x_2	Nighttime lights	1.416	0.101	0.095	1.928	0.875	0.292
x_3	NO ₂	0.233	0.013	0.012	0.364	0.125	0.037
x_4	Urban land cover	0.121	0.022	0.021	1.356	1.739	0.519
x_5	NDVI	−0.0002	0.002	0.003	−0.0001	0.042	0.038
x_6	NDBI	−0.00004	0.000	0.000	−0.0002	0.040	0.034
$\widehat{\sigma}_{y^*}^2$ / Min. MSE / N		14.46	13.53	2,203	232.97	69.60	124

Annual panel: official GDP dominates

- ▶ GDP carries almost all of the effective weight
- ▶ Satellite SNRs are uniformly low; lights lead among them

15-year long difference: the balance flips

- ▶ The satellites now carry almost all of the weight
- ▶ Urban land cover becomes the single most informative signal

Why the ranking reverses

Noise washes out; the latent signal accumulates

Over 15 years measurement noise averages away while the cumulative latent signal grows with the horizon — the satellite SNRs rise by an order of magnitude.

And the ordering within the satellites flips

- ▶ Nighttime lights record both the extent of activity and its intensity, so they respond year to year
 - For each pixel, we know whether it is lit up and how much (intensity is a continuous number)
- ▶ Urban land cover records only spatial expansion — a slow structural process, nearly invisible annually but dominant over fifteen years
 - For each pixel, we only know whether it is urban or not (classification is binary)
- ▶ NDVI and NDBI stay weakest throughout: 500 m pixels average away the intensive margin

NO₂ predicts little — but it identifies

$\gamma_{NL} = \text{Cov}(NL, NO_2) / \text{Cov}(GDP, NO_2)$ needs a non-trivial $\text{Cov}(GDP, NO_2)$. Urban land cover is the most predictive signal at long horizons but the least reliable identifier there — NO₂ does the opposite job.

Specification 2: energy statistics reduce the MSE substantially

		Panel, 2006–2022			Long-difference		
	Signal	γ_k	SNR	$\lambda_k \gamma_k$	γ_k	SNR	$\lambda_k \gamma_k$
x_1	GDP	1.000	0.622	0.277	1.000	0.164	0.008
x_2	Nighttime lights	1.689	0.080	0.028	3.325	4.436	0.899
x_3	NO ₂	0.424	0.026	0.009	0.374	0.072	0.015
x_4	CO ₂	1.995	0.717	0.330	2.215	0.361	0.045
x_5	Electricity	1.401	0.775	0.357	1.914	0.323	0.033
$\hat{\sigma}_{y*}^2$ / Min. MSE / N		8.40	2.94	2,526	139.00	28.18	145

A clean division of labour

- ▶ On the common sample, panel MSE falls 10.25 $\rightarrow$ 1.60: an 84% reduction vs. Specification 1
- ▶ Energy carries prediction (SNR 0.78, 0.72 vs. 0.08 for lights)
- ▶ Satellites carry identification

The correlated errors are economically readable

- ▶ $\sigma_{14} = -2.11$ (GDP–CO₂), $\sigma_{15} = -1.76$ (GDP–electricity)
- ▶ When GDP overstates growth, energy looks low relative to the latent factor, and the BLU shifts weight away from GDP
- ▶ Without the satellites, none of this is identified

Does nonlinearity help in the data?

Model A' ($K = 10$): Spec. 1 + sieve

Drop NDVI/NDBI, then add squares and cross terms:

GDP, NL, NO₂, Urban, (NL)², (NO₂)²,
(Urban)², NL × NO₂, NL × Urban, NO₂ × Urban

BLU MSE: 11.51 → 10.88 (−5.4%)

Only (NL)² registers: $\lambda_5 \gamma_5 = 0.014$.

Model B' ($K = 15$): Spec. 2 + sieve


GDP, NL, NO₂, CO₂, Elec, (NL)², (NO₂)²,
(CO₂)², (Elec)², NL × NO₂, NL × CO₂,
NL × Elec, NO₂ × CO₂, NO₂ × Elec, CO₂ × Elec

BLU MSE: 2.94 → 2.51 (−14.5%)

Again only (NL)²: $\lambda_6 \gamma_6 = 0.079$, SNR 0.213.

- ▶ Only (NL)² adds signal — total generated information is finite
- ▶ Concentration in (NL)² is the light–output curvature of channel (i), showing up in the data
 - The simulated VIIRS night lights (SVNL) is not top-coded, so the curvature reflects genuine light–output relationship rather than a sensor artifact
- ▶ Squares of NO₂, CO₂, electricity each < 0.005; cross terms < 0.01
- ▶ In B', (NL)² alone gives about 5%; the other terms offset noise through their error covariances

Country-level estimates, 2006–2022



Model A tracks official GDP closely; Models B and B' move independently — energy signals shift the estimate whenever CO_2 and electricity diverge from reported output.

Common-sample country-years; each model uses the weights estimated on its own sample.

Systematic corrections in levels and in smoothness

Country	Mean growth (%)				Volatility, mean $ \Delta g $ (pp)			
	Official	A ($K=6$)	B ($K=5$)	B' ($K=15$)	Official	A	B	B'
China	7.69	7.49	5.78	5.28	1.60	2.25	1.71	1.45
India	5.91	5.95	5.06	4.70	2.92	3.04	2.08	1.94
Brazil	1.05	1.14	1.87	1.84	2.91	3.11	2.46	2.38
Indonesia	4.78	4.45	4.56	4.38	1.24	2.63	2.18	1.72
Nigeria	3.32	3.53	2.16	2.46	2.62	2.60	3.54	3.45
South Africa	2.27	2.32	1.80	1.78	1.58	1.88	1.41	1.37

Levels


- ▶ China 7.69 $\rightarrow$ 5.78/5.28; India 5.91 $\rightarrow$ 5.06/4.70 — 0.9–2.4 pp down, as in Nakamura et al. (2016), Hu & Yao (2022), Martinez (2022)
- ▶ Brazil moves the other way (1.05 $\rightarrow$ 1.87): its energy use outgrew reported GDP

Smoothness

- ▶ Indonesia's official series is the smoothest here (1.24 pp) — striking for a commodity exporter; the signals imply more than double that
- ▶ Elsewhere Model B compresses volatility (India –30%; Brazil –15%, South Africa –11%); Nigeria is the exception (3.54 vs. 2.62)

External validation: do we flag the right countries?

For each country-year, take the revision of real GDP growth between its $t+1$ and $t+3$ IMF World Economic Outlook vintages — an independent, model-free proxy for data reliability.



External validation: what the revisions show

Signed gaps: no relationship

- ▶ Mean gap vs. mean revision: Pearson $+0.22$, Spearman $+0.14$
- ▶ Expected — vintages reflect rebasing and new source data; accounts are seldom revised to concede inflated growth

Magnitudes: strongly related

- ▶ $|\text{gap}|$ vs. $|\text{revision}|$: Pearson $r = +0.57$, Spearman $\rho = +0.45$ across 162 countries
- ▶ Biggest disagreements — Libya, Venezuela, Iraq, Timor-Leste, Myanmar — are the most-revised series

Takeaways

Theory

- ▶ A general K -signal framework with closed-form BLP and BLU combinations and correlated errors; CGY is a special case, and HSW's under-identification is resolved
- ▶ More signals drive the MSE to zero if loadings decay no faster than $k^{-1/2}$ or noise cancels in the signal direction — not if loadings decay faster and errors stay well conditioned
- ▶ Nonlinear transforms of existing observables are legitimate new signals; a sieve attains the oracle bound
- ▶ Their information is bounded — generated and independent signals are complements

Measuring growth

- ▶ Satellite informativeness rises sharply with the horizon, and the ranking reverses: lights annually, urban land cover over 15 years
- ▶ CO₂ and electricity, admitted with their correlated errors, cut the annual MSE by 84%; satellites then serve identification
- ▶ China and India below official, Brazil above, Indonesia too smooth — and the countries we disagree with most are the ones later revised most
- ▶ **Next:** mobile-phone activity, internet traffic, high-frequency financial data; time-varying panels

Thank you

BACKUP

Additional Material

Backup: summary statistics of the eight 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

Pooled 2006–2022 country-year distribution, before trimming. Cross-sectional SDs of the growth signals — lights (34.0), NO₂ (11.7), CO₂ (11.2) — far exceed that of official GDP log growth (6.2), reflecting both larger γ_k and noisier measurement. Urban land cover has a positive median (1.59), consistent with steady urbanization; NDVI and NDBI have medians near zero, consistent with slow-moving stocks.

Backup: correlations of the satellite signals with GDP

Signal	Panel		Long-difference	
	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

1%/99%-trimmed sample; the panel is country-year over 2006–2022 and the long-difference is 2020–22 averages minus 2006–08 averages. NO₂'s correlation with GDP rises at long horizons, which is what anchors the identification of γ_{NL} . Urban land cover does the reverse: comparable in the panel, near zero in the long difference — so it is an unreliable identifying signal precisely where satellites matter most.

Backup: structural-zero patterns of V

Specification 1 (Model A)

$\varepsilon_k \perp \varepsilon_{\text{GDP}}$ for all k ; NL, NDVI, and NDBI errors may be correlated (shared near-infrared inputs).

$$V_{S1} = \begin{pmatrix} \sigma_{11} & 0 & 0 & 0 & 0 & 0 \\ 0 & \sigma_{22} & 0 & 0 & \sigma_{25} & \sigma_{26} \\ 0 & 0 & \sigma_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma_{44} & 0 & 0 \\ 0 & \sigma_{52} & 0 & 0 & \sigma_{55} & \sigma_{56} \\ 0 & \sigma_{62} & 0 & 0 & \sigma_{65} & \sigma_{66} \end{pmatrix}$$

GDP is the fully-uncorrelated reference:

$$\gamma_2 = \frac{\text{Cov}(x_2, x_3)}{\text{Cov}(x_1, x_3)}, \quad \sigma_{y^*}^2 = \frac{\text{Cov}(x_1, x_2)}{\gamma_2}, \\ \gamma_k = \frac{\text{Cov}(x_k, x_1)}{\sigma_{y^*}^2}.$$

For B' ($K = 15$) a mixed-reference rule is used: NL by default; NO₂ for generated terms involving NL; CO₂ for the NL·NO₂ cross term.

Specification 2 (Model B)

CO₂ and electricity share inputs with national accounts; NL is clean.

$$V_{S2} = \begin{pmatrix} \sigma_{11} & 0 & 0 & \sigma_{14} & \sigma_{15} \\ 0 & \sigma_{22} & 0 & 0 & 0 \\ 0 & 0 & \sigma_{33} & 0 & 0 \\ \sigma_{41} & 0 & 0 & \sigma_{44} & \sigma_{45} \\ \sigma_{51} & 0 & 0 & \sigma_{54} & \sigma_{55} \end{pmatrix}$$

NL (x_2) is the reference: $\gamma_k = \frac{\text{Cov}(x_k, x_2)}{\text{Cov}(x_1, x_2)}$ for $k = 3, 4, 5$; then $\sigma_{jk} = \text{Cov}(x_j, x_k) - \gamma_j \gamma_k \sigma_{y^*}^2$ at the three allowed positions.

Backup: Model A' polynomial sieve ($K = 10$)

	Description	γ_k	σ_k^2	Signal SD	Noise SD	SNR	λ_k	$\lambda_k \gamma_k$
x_1	GDP	1.000	12.474	4.125	3.532	1.364	0.872	0.872
x_2	NL	1.200	305.10	4.951	17.467	0.080	0.060	0.072
x_3	NO ₂	0.2308	61.915	0.9519	7.869	0.015	0.039	0.009
x_4	Urban	0.1070	9.750	0.4413	3.123	0.020	0.287	0.031
x_5	(NL) ²	-13.562	406,606	55.937	637.66	0.008	-0.001	0.014
x_6	(NO ₂) ²	-0.4879	8242.49	2.012	90.788	0.000	-0.000	0.000
x_7	(Urban) ²	0.1222	949.70	0.5039	30.817	0.000	-0.022	-0.003
x_8	NL·NO ₂	-1.235	22,935	5.094	151.44	0.001	-0.001	0.001
x_9	NL·Urban	0.0100	3010.84	0.0413	54.871	0.000	0.001	0.000
x_{10}	NO ₂ ·Urban	-0.4319	505.46	1.781	22.483	0.006	-0.009	0.004
Total $\lambda^T \gamma$:								1.000

1%/99%-trimmed panel of 2,329 country-year observations on 148 countries, GDP as the universal reference. $\hat{\sigma}_{y^*}^2 = 17.01$; BLU MSE = 10.88 versus 11.51 for the linear $K = 4$ baseline on the same sample (-5.4%).

Backup: Model B' polynomial sieve ($K = 15$)

	Description	γ_k	σ_k^2	Signal SD	Noise SD	SNR	λ_k	$\lambda_k \gamma_k$
x_1	GDP	1.000	13.499	2.898	3.674	0.622	0.242	0.242
x_2	NL	1.689	300.56	4.894	17.337	0.080	0.006	0.010
x_3	NO ₂	0.4244	59.015	1.230	7.682	0.026	0.019	0.008
x_4	CO ₂	1.995	46.599	5.782	6.826	0.717	0.154	0.307
x_5	Electricity	1.401	21.275	4.062	4.612	0.775	0.238	0.333
x_6	(NL) ²	95.336	358,047	276.29	598.37	0.213	0.001	0.079
x_7	(NO ₂) ²	1.396	7772.90	4.045	88.164	0.002	0.001	0.001
x_8	(CO ₂) ²	2.544	19,592	7.371	139.97	0.003	0.001	0.002
x_9	(Elec) ²	3.128	5827.73	9.066	76.340	0.014	0.000	0.001
x_{10}	NL·NO ₂	-0.9028	22,331	2.617	149.44	0.000	-0.001	0.001
x_{11}	NL·CO ₂	-3.710	32,398	10.751	179.99	0.004	-0.001	0.004
x_{12}	NL·Elec	-3.461	16,024	10.030	126.59	0.006	-0.002	0.006
x_{13}	NO ₂ ·CO ₂	-0.9323	5568.32	2.702	74.621	0.001	0.002	-0.002
x_{14}	NO ₂ ·Elec	-0.8697	2635.63	2.521	51.338	0.002	-0.002	0.001
x_{15}	CO ₂ ·Elec	-1.203	4740.44	3.488	68.851	0.003	-0.005	0.006
Total $\lambda^T \gamma$:								1.000

1%/99%-trimmed panel of 2,526 country-year observations on 163 countries, mixed-reference identification. $\hat{\sigma}_{y^*}^2 = 8.40$; BLU MSE = 2.51, 14.5% below Model B. Only (NL)² carries a meaningful generated contribution ($\lambda_6 \gamma_6 = 0.079$).

Backup: full Specification 1 table (Model A, $K = 6$)

Description	Panel (2006–2022)				Long-difference			
	γ_k	σ_k^2	SNR	$\lambda_k \gamma_k$	γ_k	σ_k^2	SNR	$\lambda_k \gamma_k$
x_1 GDP	1.000	15.556	0.929	0.870	1.000	886.16	0.263	0.079
x_2 NL	1.416	286.58	0.101	0.095	1.928	990.01	0.875	0.292
x_3 NO ₂	0.2331	61.139	0.013	0.012	0.3637	246.07	0.125	0.037
x_4 Urban land cover	0.1214	9.727	0.022	0.021	1.356	246.47	1.739	0.519
x_5 NDVI	−0.000156	0.000177	0.002	0.003	−0.000144	0.000113	0.042	0.038
x_6 NDBI	−0.000042	0.000318	0.000	0.000	−0.000178	0.000185	0.040	0.034
$\hat{\sigma}_{y^*}^2$		14.455				232.97		
Min. MSE		13.526				69.595		
N		2203				124		

Panel: 2,203 country-year observations on 146 countries, 2006–2022. Long-difference: a 124-country cross-section comparing 2020–22 averages with 2006–08 averages. GDP is the fully-uncorrelated reference; V is diagonal apart from $\sigma_{2,5}$, $\sigma_{2,6}$, and $\sigma_{5,6}$. The effective contributions $\lambda_k \gamma_k$ sum to one by unbiasedness.

Backup: full Specification 2 table (Model B, $K = 5$)

		Panel (2006–2022)				Long-difference			
	Description	γ_k	σ_k^2	SNR	$\lambda_k \gamma_k$	γ_k	σ_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
$\hat{\sigma}_{y^*}^2$		8.399				139.00			
Min. MSE		2.938				28.183			
N		2526				145			

Panel: 2,526 country-year observations on 163 countries, 2006–2022. Long-difference: a 145-country cross-section. Identification uses NL as the reference signal; V carries the structural off-diagonals $\sigma_{1,4}$, $\sigma_{1,5}$, and $\sigma_{4,5}$.

Backup: Assumption 5 and the rate in Proposition 3

Let $\hat{\mu}$ collect the $J_K = K(K+3)/2$ sample moments, with $\hat{\mu}_j - \mu_j = N^{-1} \sum_i Y_{j,i}$.

Assumption 5

There are constants $C_x, v, b < \infty$, $a > 0$, and $0 < s_- < s_+ < \infty$, independent of K and N , such that

- 1) $\max_{k \leq K} \mathbb{E}|x_k|^4 \leq C_x^2$; 2) $\mathbb{E}|Y_{j,i}|^q \leq \frac{q!}{2} v^2 b^{q-2}$ for every $q \geq 2$, uniformly in j, K, N ;
- 3) $\min\{|\text{cov}(x_1, x_2)|, |\text{cov}(x_1, x_3)|, |\text{cov}(x_2, x_3)|\} \geq a$ and $s_- \leq \sigma_{x^*}^2 \leq s_+$.

- ▶ Part 2 is a Bernstein moment condition: with a union bound over the J_K moments, it controls all sample moments at once
- ▶ Part 3 keeps the identifying ratios away from zero denominators; the screen enforces the same regularity in the sample
- ▶ The map from moments to weights contains V^{-1} , so $r_{K,N} = K^3 \log K / (N \nu_K^3)$ worsens as V approaches singularity
- ▶ Caveat: the rate is asymptotic and conservative, with unspecified constants — it is not a finite-sample cap on K

Backup: the BLP and BLU combinations

BLP = BLU shrunk toward the pooled mean

With $P = \gamma^T V^{-1} \gamma$,

$$\lambda_1^{\text{BLP}} = \theta \lambda_1^{\text{BLU}}, \quad \theta = \frac{\sigma_{y^*}^2 P}{1 + \sigma_{y^*}^2 P} = \frac{\sigma_{y^*}^2}{\sigma_{y^*}^2 + 1/P} \in (0, 1)$$

$$\hat{y}^*(\lambda^{\text{BLP}}) - \mathbb{E}[x_1] = \theta (\hat{y}^*(\lambda^{\text{BLU}}) - \mathbb{E}[x_1])$$

$$\text{Min. MSE: } \frac{1}{P} \rightarrow \left(\frac{1}{\sigma_{y^*}^2} + P \right)^{-1}$$


- ▶ $\mathbb{E}[x_1]$: pooled mean of official growth (3.24% in Model A's sample)
- ▶ θ : share of the BLU estimate's variance that is true signal
- ▶ Centered on $\mathbb{E}[x_1] + \theta(y^* - \mathbb{E}[x_1])$: understates growth above the pooled mean, overstates it below

	θ	Min. MSE	
		BLU	BLP
Model CGY ($K = 3$)	0.51	13.86	7.08
Model A ($K = 6$)	0.52	13.53	6.99
Model B ($K = 5$)	0.74	2.94	2.18

Model CGY: official GDP, nighttime lights, and NO₂ with diagonal V (CGY's three-signal structure, with NO₂ in place of urban land cover), estimated on Model A's sample.

The BLP's lower MSE is an average over countries and years. To measure each country's growth, including its departure from the typical country, we use the BLU.

Backup: BLP and BLU estimates for six countries



CGY (BLP) lies between CGY (BLU) and the pooled mean of 3.24%, with θ times its volatility.
For China it averages 5.29%, against 7.26% for CGY (BLU) and 7.69% official, from the same signals.
Common-sample country-years, as in Figure 2; Models A and B are the BLU estimates of the main text.