

Growth Without Accountability: Measuring India's Institutional and Distributional Transformation, 2004– 2026

Varna Sri Raman · Independent researcher

2026-08-15 · Data vintage 2026-08-15 · CC BY 4.0 · Replication repository

Abstract

India's economy grew at an average of 6.1 per cent a year between 2014 and 2026. Over the same period the share of national income accruing to the bottom half of the population fell from 15.0 to 12.9 per cent, the share of workers holding jobs with any social-security cover fell from 13.0 to 10.9 per cent, and every available third-party measure of democratic quality deteriorated. This paper asks whether these facts are separable — whether India experienced a familiar and possibly tolerable trade of political openness for material progress — and concludes that it did not.

The difficulty in answering the question is that the instruments needed to answer it have themselves degraded. A decennial census was postponed indefinitely; a national consumption survey was conducted and then withheld; the national accounts were rebased in a manner that remains disputed and has never been independently reviewed; and the statistical system's institutional independence was reduced by merger and resignation. Any assessment of Indian economic performance after 2014 must therefore treat measurement capacity as an endogenous variable rather than as a fixed background condition.

We construct three indices to make that degradation legible and comparable across political eras. The Statistical Suppression Index (SSI, 0–10) aggregates graded severities across six documented streams of statistical interference, with a persistence term so that resolved episodes decay rather than vanish. The Fiscal Centralisation Index (FCI, 0–1) averages six min–max normalised components of the fiscal relationship between the union and the states, including the structural transfer of indirect taxing power effected by the Goods and Services Tax. The Democratic Quality Index (DQI, 0–1) is the geometric mean of four published third-party series — V-Dem's Liberal Democracy and Core Civil Society indices, Freedom House's aggregate score, and the Reporters Without Borders press-freedom rank.

All three are computed by a single deterministic script applied identically to 2004–2026, so that the comparison between the United Progressive Alliance decade (2004–2013) and the National Democratic Alliance period (2014–2026) rests on one consistent construction rather than on two differently-built measures.

The results are one-directional. SSI moves from 0.00 across the entire UPA decade to a peak of 9.00 in 2021–22 and remains at 6.40 in 2026. FCI rises from a UPA average of 0.09 to an NDA average of 0.57, peaking at 0.92 in 2020. DQI falls from a UPA average

of 0.57 to 0.29 by 2026. On eight of nine directly comparable cross-era measures India performs worse in the later period. The single exception is consumer price inflation, which is substantially lower — though we argue at length that this should not be read as a policy result: the comparison straddles the collapse of a global commodity cycle, and insofar as domestic policy contributed, it did so through a rule that constrains executive discretion rather than expands it.

We then subject the governance finding to external validation it did not generate. Grier and Grier (2026), working independently with a different method (synthetic control with randomisation inference), different data (V-Dem and the Penn World Table), and a different pre-treatment window (1984–2013), find that all ten V-Dem governance indicators they test deteriorate against a synthetic counterfactual at a joint p-value of 0.00. Their electoral-democracy gap of –60.2 per cent sits in the same territory as the –46.7 per cent raw decline in India's V-Dem Liberal Democracy Index recorded in our own data, despite the two quantities not being the same object. Convergence of this kind, across independent methods, is stronger evidence than either study provides alone.

We are explicit about what this paper does not establish. The DQI is built entirely from published third-party series and inherits their assumptions. The SSI severities are author-coded against a documented rubric and are the most contestable component here. The FCI is relative by construction and its values are not interpretable outside the 2004–2026 sample. None of the three indices identifies a causal effect. We do not claim that Indian incomes would have been higher under different governance, and we note where a companion study does make that claim on evidence we have not independently verified.

Keywords: India; political economy; democratic backsliding; statistical capacity; fiscal federalism; income inequality; composite indices; synthetic control

JEL codes: E01, H77, O53, P16

1. Introduction

1.1 The question

Between 2014 and 2026 India became the world's most populous country, the fourth-largest economy at market exchange rates, and — on the assessment of Varieties of Democracy, Freedom House, and Reporters Without Borders alike — substantially less democratic. Each of these facts is well documented on its own. The relationship between them is contested.

One reading treats the two developments as a bargain, and not an unusual one: a state that concentrates authority can act decisively, and the material returns compensate for the political costs. This argument has a long lineage in development economics and a specific Indian articulation in the case made for a strong executive capable of overriding the veto points that were held to have slowed Indian growth for decades. On this reading, the erosion of institutional constraint is a price, and the question is only whether the price was worth paying.

A second reading denies that the trade exists. On this account, accountability institutions are not a drag on growth but part of its machinery: they are how errors get detected, how rents get contested, and how the distribution of gains stays politically negotiable. Weaken them and the economy does not accelerate — it merely stops producing reliable information about itself, and its gains concentrate.

This paper adjudicates between these readings empirically, and finds for the second. It does so not by attempting to identify a causal effect of governance on growth — an identification problem we do not claim to solve — but by establishing what actually happened along both dimensions, on a consistent measurement basis, across two comparable political periods.

1.2 The obstacle

There is a difficulty that any such exercise must confront first, and that most treatments of Indian economic performance handle by ignoring.

The measurement apparatus deteriorated during the period under study. This is not a general complaint about data quality in developing countries. It refers to a specific and datable sequence of events:

- The decennial census due in 2021 was postponed and had not been conducted as of this writing, leaving welfare allocation, seat apportionment, and every population-denominated statistic running on 2011 figures for over a decade.
- The 2017–18 Consumer Expenditure Survey was completed and then withheld from publication after leaked findings indicated the first fall in real rural consumption in four decades. No comparable survey was published until HCES 2022–23, which changed methodology sufficiently that it does not cleanly splice onto the earlier series.
- The 2015 revision of the national accounts shifted the base year from 2004–05 to 2011–12 and, more consequentially, changed the measurement of manufacturing output from a production survey to the MCA21 corporate-filings database. The revision lowered measured UPA-era growth and raised measured NDA-era growth. Subramanian (2019) estimated that growth over 2011–17 was overstated by roughly 2.5 percentage points annually; the Economic Advisory Council to the Prime Minister issued a rebuttal; no independent review was ever conducted.
- The Periodic Labour Force Survey's 2017–18 results were withheld, prompting the resignation of two members of the National Statistical Commission, and the National Sample Survey Office was subsequently merged into the National Statistical Office, reducing its institutional separation from the ministry whose performance its data measures.
- Official COVID-19 mortality was recorded at approximately 0.5 million, against excess-mortality estimates in the range of 3–5 million from multiple independent teams.

Each of these is individually documented. Considered together, they constitute something more than a series of unfortunate administrative events: they describe a system in which the production of official statistics became less independent of the government those statistics evaluate. That is a fact about the Indian political economy, and it is measurable.

The methodological consequence is that treating measurement capacity as fixed background — the standard approach — will systematically understate the deterioration, because the very instruments that would register it are among the things that degraded. Our response is to make statistical capacity an explicit measured object rather than an assumption.

1.3 What this paper does

We construct three indices covering 2004–2026 and apply them identically across both political eras:

1. **The Statistical Suppression Index (SSI)**, on a 0–10 scale, measuring interference with the production of official statistics.
2. **The Fiscal Centralisation Index (FCI)**, on a 0–1 relative scale, measuring the shift of fiscal authority from states to the union.

3. **The Democratic Quality Index (DQI)**, on a 0–1 scale, aggregating four published third-party assessments of democratic quality.

Alongside these we assemble a conventional macroeconomic and distributional series — growth, unemployment, formal employment share, top 1 per cent and bottom 50 per cent income shares, consumer price inflation, fiscal deficit, and press-freedom rank — from official and established third-party sources.

We then compare the two eras, examine the internal structure of each index, and test the governance result against an entirely independent study using a different method.

1.4 Contribution and limits

The contribution is threefold. First, the SSI, to our knowledge, is the first quantitative index of statistical suppression constructed for India on a consistent multi-era basis; it makes an argument usually conducted anecdotally into something that can be plotted, disputed, and recomputed. Second, applying one deterministic construction across both eras removes a common source of spurious cross-era difference — that the periods were measured differently. Third, the entire dataset, the construction script, and the paper source are public, so every number here can be regenerated or contradicted.

The limits are equally important and we state them at the outset rather than burying them in a final section. This paper is descriptive. It does not identify causal effects. The SSI's severity codings are the author's, made against a published rubric, and a reasonable analyst applying the same rubric could produce somewhat different numbers — we quantify how much this matters in Section 8. The FCI is relative to its own sample and its absolute values carry no meaning outside it. The DQI inherits whatever biases exist in V-Dem, Freedom House, and RSF, including the well-known critique that expert-survey democracy measures may be sensitive to the international reputational climate surrounding a country. We return to each of these.

1.5 Structure

Section 2 sets out the measurement problem in more detail. Section 4 describes the data and its provenance. Section 5 specifies the three indices. Section 6 presents results. Section 7 reports the external validation exercise. Section 8 covers robustness. Section 9 states limitations and engages the strongest objections we are aware of. Section 10 discusses interpretation, and Section 11 concludes.

2. Framing: the alleged trade and how to test it

2.1 Two accounts of the relationship

The proposition that political constraint impedes economic development is old and not foolish. Its general form holds that democratic governance imposes costs on growth: coalition bargaining slows decisions, distributive politics diverts resources toward consumption, judicial review and federal veto points delay infrastructure, and electoral cycles shorten the horizon over which

governments will bear costs for later benefits. A government less encumbered by these can, on this account, do things that a more encumbered one cannot.

The Indian version of this argument has a specific history. The decades of slow growth before liberalisation were widely attributed to a state simultaneously over-controlling and under-capable, and the diagnosis of "policy paralysis" applied to the later UPA years located the problem in coalition politics and institutional friction rather than in policy content. The case for a decisive executive followed directly: what India lacked was not the right ideas but the capacity to act on them.

The competing account does not deny that constraint has costs. It denies that the costs are net, on the ground that accountability institutions perform functions internal to economic performance rather than external to it. They are the mechanism by which policy errors are detected before they compound, by which rent extraction is contested, by which the distribution of gains remains politically negotiable, and — critically for this paper — by which the state's own claims about its performance can be checked against evidence it does not control.

On this second account, weakening those institutions does not buy speed. It buys the *appearance* of speed, because the machinery that would have reported the shortfalls has been weakened at the same time.

2.2 What distinguishes the two empirically

These accounts make different predictions, and the difference is testable without solving an identification problem.

The bargain account predicts a trade: institutional quality falls, material outcomes improve. The second account predicts something more specific and more unusual — institutional quality falls, material outcomes do not improve, and the *measured* picture nonetheless looks better than the actual one, because measurement capacity is among the institutions that fell.

The second prediction has an observable signature. If it holds, we should find: no acceleration in growth; deterioration in distributional and employment measures that are hard to massage; and a pattern of statistical interference that is asymmetric with respect to the political convenience of the suppressed numbers. All three are things we can look for directly.

This is why the paper measures statistical capacity rather than assuming it, and why the SSI is not an ancillary index but the one that distinguishes the two hypotheses. A finding that institutional quality fell while growth did not accelerate is consistent with a failed bargain. A finding that institutional quality fell, growth did not accelerate, *and* the statistical apparatus degraded asymmetrically is evidence for the stronger claim that the bargain was never available.

2.3 Why not simply run a regression

The obvious alternative approach — regress growth on some measure of institutional quality across countries and years — is not available for the question as posed, for reasons worth stating.

Institutional quality and economic performance are jointly determined and share common causes, so cross-country panel estimates recover correlations whose causal interpretation depends on assumptions that the data cannot verify. The problem is compounded here because the dependent variable is measured by an apparatus whose independence is the treatment. And with a single country over a single episode, there is no useful within-sample variation to exploit.

Our approach is therefore descriptive by choice rather than by default: establish carefully what happened on both dimensions, on one consistent measurement basis, and let the pattern discipline the interpretation. Section 7 reports what a well-executed counterfactual design finds when applied to the governance half of the question, which is the closest thing to causal evidence available on this episode — and it was produced independently of us.

2.4 Where this sits in the literature, and what it adds

The account defended here is not free-standing. It modifies three existing bodies of work and introduces one concept we have not found named elsewhere.

Reframing modernisation theory. The classical expectation runs from growth to democracy: rising incomes produce an educated middle class with the means and the motive to demand accountable government. The Indian case suggests a refinement rather than a refutation. Growth can also *legitimise* erosion. Economic performance functions as political capital that purchases tolerance for institutional consolidation; a middle class that is benefiting can be co-opted rather than mobilised. What the data here describe is development coexisting with — and arguably underwriting — democratic decline, which the classical sequence does not anticipate.

We would stress the conditional. This is one country over twelve years, and a single case cannot overturn a cross-national regularity. The claim is that the mechanism is available, not that it is general.

Extending competitive authoritarianism. Levitsky and Way (2010) describe hybrid regimes in which meaningful electoral competition persists alongside an incumbent advantage substantial enough to make the playing field unfair. India after 2014 fits the category, but the mode of transition is distinctive in a way worth naming.

The erosion documented here is *incremental and legal*. No election was cancelled, no constitution suspended, no emergency declared — the contrast with 1975–77 is instructive precisely because the current V-Dem score is comparable to that period without any of its formal apparatus. Each individual step was procedurally lawful: a survey withheld on stated quality grounds, a census postponed on pandemic grounds, a statistical office reorganised by administrative order, a tax reform that happened to remove state fiscal discretion. What the indices capture is the *accumulation* of such steps, each defensible alone.

That suggests a refinement to the competitive-authoritarianism framework: the useful unit of analysis is not the discrete transgression but the graded, cumulative loss of accountability capacity. This is why the SSI is constructed from graded severities with a persistence term rather than from binary events — the measurement design follows from the theoretical claim about how this kind of transition proceeds.

A proposed concept: fiscal authoritarianism. The FCI measures something for which we lack a standard term: the use of federal fiscal architecture to centralise political power while leaving the constitutional structure formally intact.

The instruments are ordinary and individually unremarkable — cesses levied outside the divisible pool, conditional transfers that constrain state expenditure discretion, borrowing restrictions, and a consumption-tax regime that removes independent state rate-setting. None requires constitutional amendment. None reads, on its own, as a transfer of political authority. Together they shift the practical balance of the federation while every formal guarantee remains in place.

We offer *fiscal authoritarianism* as the name for this: revenue instruments functioning as control mechanisms, conditional transfers as political tools, and the subversion of fiscal federalism through means that never require the constitution to be touched. Whether

the term proves useful beyond India is an empirical question we cannot settle here — but the phenomenon it names is measurable, and Section 6.5 measures it.

Measurement as contribution. The final contribution is instrumental rather than theoretical. Arguments about statistical suppression are usually conducted anecdotally, one episode at a time, which makes them easy to answer one episode at a time. Rendering the phenomenon as a weighted, graded, reproducible index makes it something that can be plotted against other series, compared across political eras, and — most importantly — disputed on the record. A critic who thinks the SSI is wrong can say precisely where, and recompute.

3. The measurement problem

3.1 Why statistical capacity is not background

Economic analysis normally treats the statistical system as an instrument: imperfect, noisy, but exogenous to the phenomena it measures. Where the instrument is known to be biased, the bias is modelled and corrected.

That treatment fails when the instrument's accuracy is itself a function of the political variable under study. If a government's consolidation of authority extends to the agencies producing the statistics by which its performance is judged, then measured performance and actual performance diverge in a direction correlated with the independent variable. Standard corrections do not help, because the correction would need to be estimated from the same compromised source.

India after 2014 presents this problem in an unusually clean form, because the interference is datable. We are not inferring degradation from a discrepancy between Indian data and some prior; we are pointing at specific published events — a postponed census, a withheld survey, a disputed rebasing, a merger of a statistical office, two commissioner resignations — each with a date and a public record.

3.2 The direction of bias

It matters which way the resulting bias runs, and in this case the direction is not ambiguous.

Every one of the episodes listed above had the effect of removing from public view a number that was either known to be unfavourable or was widely expected to be. The withheld 2017–18 consumption survey had shown rural consumption falling. The postponed census would have updated a population denominator whose staleness inflates per-capita welfare coverage ratios. The national-accounts revision raised measured growth in the later period. Official COVID mortality was an order of magnitude below independent excess-mortality estimates.

We are aware that this observation could be made to sound like a conspiracy claim, and we do not intend it as one. There are ordinary administrative explanations available for several individual episodes: a pandemic genuinely does disrupt a census; methodology genuinely does need periodic revision; survey results are sometimes genuinely held for quality review. Our claim is narrower and, we think, harder to dismiss: the *joint distribution* of these events is asymmetric. Administrative friction that is uncorrelated with the political convenience of the resulting number should produce suppressions and delays in both directions.

What is observed is a sequence in which the withheld or delayed numbers are, with high consistency, the unflattering ones. That asymmetry is the object the SSI is designed to capture.

3.3 The consequence for cross-era comparison

If measurement quality differs systematically between the UPA and NDA periods, then any raw comparison of outcomes across the two is confounded. Specifically, if the later period's statistics are more favourably biased, then a finding that the later period performed *worse* on some measure is conservative: the true deterioration is at least as large as the measured one.

This is the sense in which our headline result is robust to the very problem it documents. We find that India performs worse on eight of nine comparable measures in the later period. If the later period's data are, as we argue, biased in the government's favour, then correcting for that bias would strengthen rather than weaken the finding.

The exception to note is that this logic does not apply symmetrically to the single measure on which the later period performs better — consumer price inflation. We return to that measure in Section 6.9, where we argue that its improvement is real as a measurement but is largely not a result of Indian policy.

4. Data

4.1 Coverage and structure

The dataset covers calendar years 2004 to 2026 inclusive. Observed series run through 2026 and reflect the latest official release available as of 15 August 2026. The 2004–2013 period corresponds to the two United Progressive Alliance governments; 2014–2026 spans the three National Democratic Alliance governments under the same Prime Minister.

The full machine-readable dataset is published as `data.json`, with a flat tabular export as `dataset.csv`. Every series carries a source string, and every constructed series carries a pointer to the script that generates it.

4.2 Series and provenance

We distinguish four provenance classes, and label every series accordingly, because conflating them is how composite indices acquire unearned authority.

Official or survey. Taken directly from a government or established statistical source with no transformation beyond unit conversion:

SERIES	SOURCE
GDP growth	MoSPI First Advance Estimates 2025–26; Economic Survey 2025–26 (2011–12 base)
Unemployment	PLFS Annual Report 2024–25; PLFS Monthly Bulletin (June 2026)
Youth unemployment	PLFS Annual; ILOSTAT youth (15–29)
Graduate unemployment	PLFS unit-level tabulation
Top 1% income share	World Inequality Database (WID.world) India series; WIR 2026
Bottom 50% income share	World Inequality Database (WID.world) India series; WIR 2026
Formal employment	PLFS workers with social-security cover; EPFO net additions cross-check
Press freedom rank	Reporters Without Borders, World Press Freedom Index
Fiscal deficit	Union Budget documents, central government, % of GDP
CPI inflation	MoSPI CPI (Combined)
Rupee / USD	RBI reference rates, annual average

Derived. Computed from an official series by an explicit rule — for example, informal employment share, which is 100 minus formal employment share by construction.

Constructed index. SSI, FCI, DQI. Specified in Section 5 and generated by `data/compute_indices.py`.

Compiled. Figures assembled from multiple public sources for the citizen-facing material, and for the external synthetic-control comparison in Section 6.

4.3 A basis break in the inflation series

One discontinuity requires explicit statement rather than silent splicing, and it is recent enough that most treatments will not yet have absorbed it.

From the January 2026 print, MoSPI replaced the CPI series on base 2012=100 with a new series on base 2024=100. The weights derive from the Household Consumption Expenditure Survey 2023–24, and the classification moves to COICOP 2018, replacing the previous six-group structure. The item basket comprises 358 weighted items priced across 1,395 urban markets in 434 towns, 1,465 rural markets, and 12 online markets.

The consequential change is the weight on food and beverages, which falls from 45.86 per cent under the old basket to 36.75 per cent under the new one. Housing carries 16.48 per cent, transport 9.43 per cent, and health 6.83 per cent.

This is a defensible and overdue revision: food's share of household expenditure genuinely does fall as incomes rise, and a basket fixed in 2012 had become unrepresentative. But it has a direct analytical consequence. Headline CPI computed on the new basket responds substantially less to a food-price shock than the old series did. A monsoon failure that would have moved the old index by a given amount moves the new one by roughly four-fifths as much, through the weighting change alone.

Our 2026 inflation figure (4.45 per cent, the July 2026 print) is therefore **not on the same basis** as the 2014–2025 values, and we flag it as such in the dataset rather than presenting a seamless series. Readers comparing 2026 inflation to earlier years should treat the comparison as approximate. We note without further comment that this is precisely the class of methodological change — legitimate in itself, consequential for comparison, and easy to miss — that the SSI is designed to make visible.

4.4 The pre-2014 reconstruction

The SSI, FCI and DQI values for 2004–2013 are reconstructed from the documented record rather than observed contemporaneously: survey timeliness and publication dates for SSI, Finance Commission devolution shares and Receipt Budget cess and centrally-sponsored-scheme shares for FCI, and the published V-Dem, Freedom House and RSF scores for DQI.

For the DQI this is not a reconstruction in any meaningful sense — all four inputs are third-party series that exist for the full period and require only assembly. For the FCI the official anchors exist and the interpolation between them is mechanical. For the SSI the reconstruction is more consequential, because it produces a value of exactly zero for all ten UPA years.

We address the obvious objection to that zero directly in Section 9.2, since a reader's assessment of this paper's central cross-era claim will turn substantially on whether they accept it.

4.5 Estimation principles for interpolated series

Several supporting series in the wider dataset — sanitation coverage, cess share, forest cover, out-of-pocket health expenditure, immunisation, tap-water coverage, state debt, coal share — exist only as periodic official anchors rather than annual observations. Where annual values are needed, they are produced by explicit interpolation under five principles, stated here because the difference between an observation and an estimate is exactly the kind of distinction this paper argues should not be blurred.

1. **Anchors supersede models.** No value is extrapolated beyond the range of verified anchors. Where the official record stops, the series stops.
2. **Transparency.** The interpolation rule and its parameters are recorded in the data file itself, not only in documentation.
3. **Consistency.** Every estimate is deterministic and replicable; identical inputs produce identical outputs.
4. **Neutrality.** Estimation is used to show continuity between known points, never to fill a gap in a direction that favours an interpretation.
5. **Auditability.** Every estimated series can be regenerated from the published note.

Uncertainty is represented explicitly: ± 5 per cent bounds on continuous variables, ± 3 per cent on fiscal ratios, and exact values for discrete counts, where an interpolated fractional count would be meaningless.

Two estimated series carry more modelling than the rest and are flagged accordingly. Sanitation coverage is fitted with a logistic curve between anchors of approximately 40 per cent (2014), 95 per cent (2019) and 98.5 per cent (2024), with parameters $L = 99.0$,

$k = 0.8$, $t_0 = 2018.0$ — chosen to reproduce the observed slow start, rapid middle and saturation. The cess share is piecewise linear between nine fiscal anchors running from 8.0 per cent (2004) to 29.0 per cent (2024).

None of the three headline indices depends on these interpolated supporting series. The SSI, FCI and DQI draw on the anchors and published third-party series directly.

4.6 How to read an official statistic

The analysis in this paper relies on a small number of interrogations applied to every figure it uses, and we set them out because they are also the means by which a reader can audit our own claims:

- Which definition is in use, and who set it?
- What does the number exclude?
- When was the underlying data actually collected, as opposed to published?
- Who benefits if the number looks favourable?
- Has the method changed within the comparison window?
- Was the inconvenient survey published?
- How does the figure compare with independent estimates of the same quantity?

Applied to the material here, these questions are what generate the paper's principal caveats. The unemployment rate looks low because of how "employed" is defined for self-employed and unpaid family workers. The 2026 inflation figure sits on a basket reweighted in 2026. The growth series depends on a revision never independently reviewed. The COVID mortality figure diverges from independent estimates by an order of magnitude. Each caveat is the output of one of the questions above, and each is stated in the relevant section rather than aggregated into a general disclaimer.

4.7 Reproducibility

Every figure in this paper derives from `data.json`, which is published under CC BY 4.0. The three constructed indices regenerate from `data/compute_indices.py`. Replication scripts are provided in Python and R. A continuous-integration guard verifies on every commit that the site's inline copy, the published `data.json`, and the output of the construction script agree to within 10^{-9} — a check introduced after the possibility of silent divergence between the three was identified.

5. Constructing the indices

5.1 Design principles

Three principles govern all three constructions.

One script, both eras. A recurring flaw in cross-era institutional comparison is that the two periods are measured by different instruments, so that some of the apparent difference is an artefact of construction. All three indices here are produced by a single deterministic script applied to the full 2004–2026 sample. Any difference between the eras is a difference in inputs, not in method.

Graded, not binary. Institutional degradation is continuous. A survey delayed by six months, a survey delayed by four years, and a survey discontinued are different events, and an index that codes all three as 1 discards most of the information. All severities are graded on [0,1] against a documented rubric.

Persistence. Institutional damage outlasts the event. When a suppressed statistic is eventually published, the informational gap it left does not close retrospectively, and the demonstrated willingness to suppress does not un-demonstrate itself. The SSI therefore includes a persistence term so that resolved episodes decay toward, but do not reach, zero.

5.2 The Statistical Suppression Index

The SSI is a weighted sum of graded severities across six streams:

$$\mathrm{SSI}_t = \sum_{i=1}^6 w_i \, s_{it}, \quad s_{it} \in [0,1], \quad \sum_i w_i = 10$$

Weights reflect each stream's consequence for the public capacity to evaluate government performance:

STREAM	WEIGHT	RATIONALE
Census delay	2.5	Population denominator for nearly every welfare and per-capita statistic; also determines seat apportionment
COVID mortality undercount	2.0	Largest single divergence between official and independent estimates in the period
Consumption survey suppression	1.5	Primary instrument for poverty and living-standards measurement
GDP back-series dispute	1.5	Determines the comparability of growth across the two eras
Institutional independence	1.5	NSC resignations; NSSO-NSO merger
Employment data delay	1.0	PLFS 2017–18 withholding

Severity coding follows a three-point rubric with interpolation: approximately 1.0 for full suppression or discontinuation, approximately 0.7 for a live methodology dispute, and approximately 0.3 for a delay or a post-resolution scar.

The index is anchored at zero for the UPA decade because no coded event occurred in that period.

5.2.1 The six streams, documented

Because the SSI is the least conventional construction here, and because its credibility rests entirely on whether the underlying events are real and correctly graded, we set out each stream's record.

Census delay (weight 2.5). India's decennial census has been conducted without interruption since 1881, through partition, war and famine. The 2021 census was postponed, initially on pandemic grounds, and had not been conducted as of this writing — a gap of over a decade against a ten-year cycle. Budgetary provision of ₹11,718 crore for a 2027 exercise has since been approved.

The consequences are not confined to demography. The census supplies the population denominator for essentially every per-capita and coverage statistic the state publishes, so welfare coverage ratios computed on 2011 denominators overstate reach by the extent of subsequent population growth. Estimates of the number of people excluded from subsidised food entitlements on this basis alone run to the tens of millions.

There is also a political consequence, discussed further in Section 10.3: Lok Sabha seat apportionment has been frozen on 1971 population data since 1976, and the freeze expires with the first census conducted after 2026. Postponing the census postpones a reallocation of parliamentary seats that projections suggest would transfer roughly eleven seats to Uttar Pradesh and ten to Bihar while removing approximately eight each from Tamil Nadu and Kerala. Average constituency size has risen from about 0.7 million people per member in 1951 to about 2.5 million today. A delayed census is therefore not only a statistical omission; it defers a redistribution of political representation.

Severity is coded at full weight from 2021 onward and retained thereafter, since the omission is not retrospectively repairable.

COVID mortality undercount (weight 2.0). Official Indian COVID-19 mortality was recorded at approximately 0.5 million. Independent excess-mortality estimates from multiple teams, using civil registration data, serological surveys and household surveys, cluster in the range of 3 to 5 million. A discrepancy of this magnitude in a headline vital statistic is the largest single divergence between official and independent estimates in the period.

Severity is coded at full weight for 2021–22 and decays thereafter, since the acute measurement episode passed while the unreconciled record persists.

Consumption survey suppression (weight 1.5). The 2017–18 Consumer Expenditure Survey was completed by the National Sample Survey Office. Following leaked findings indicating a fall in real rural consumption — which would have been the first such fall in over four decades — the survey was withheld from publication, with quality concerns cited. No comparable survey appeared until HCES 2022–23, producing an eleven-year gap in the primary instrument for measuring living standards and poverty. The 2022–23 survey employed a revised methodology that does not cleanly splice onto the earlier series, so the gap is not fully closed even retrospectively.

This is the clearest instance of the asymmetry described in Section 2.2: a survey was suppressed after its findings became known, and the findings were unfavourable.

GDP back-series dispute (weight 1.5). The 2015 revision shifted the national accounts base year from 2004–05 to 2011–12, moved manufacturing measurement from the Annual Survey of Industries production basis to the MCA21 corporate-filings

database, and valued output at market prices rather than factor cost. The resulting back-series lowered measured UPA-era growth and raised measured NDA-era growth.

Subramanian (2019) estimated that growth over 2011–17 was overstated by approximately 2.5 percentage points annually, attributing the discrepancy principally to the MCA21 shift and the treatment of dormant firms. The Economic Advisory Council to the Prime Minister issued a detailed rebuttal. The dispute was never referred to an independent review, and remains unresolved.

The stream is coded at the "live methodology dispute" level rather than at full suppression, since the revision was published and documented; what is absent is independent adjudication.

Institutional independence (weight 1.5). In January 2019 two members of the National Statistical Commission resigned, citing the government's failure to release the PLFS 2017–18 results and the sidelining of the Commission in the GDP back-series decision. The National Sample Survey Office was subsequently merged into the National Statistical Office, reducing the organisational separation between the body collecting the data and the ministry whose performance the data evaluates.

Unlike the other streams, this one does not decay: an organisational merger is a structural change that persists until reversed, and it has not been reversed. This is the principal reason the SSI does not return toward zero after 2022.

Employment data delay (weight 1.0). The PLFS 2017–18 results, which showed unemployment at a 45-year high, were withheld from publication and released only after the 2019 general election. The stream is coded at delay severity, with a residual scar thereafter.

5.2.2 What the weights encode

The weights are not estimated; they are asserted, and a reader is entitled to know on what basis. The ordering reflects each stream's consequence for the public capacity to evaluate government performance: the census ranks highest because it is an input to nearly every other statistic and to political representation itself; mortality ranks second because of the magnitude of the discrepancy; the employment delay ranks lowest because the data was ultimately published, if late.

A reader who disputes the ordering can recompute with alternative weights. Section 8.1 reports how much the cross-era result moves under such changes — the answer is that it survives the removal of any single stream, because no stream carries more than a quarter of the total.

5.3 The Fiscal Centralisation Index

The FCI is the unweighted mean of six components, each min–max normalised over the full 2004–2026 sample:

$$\mathrm{FCI}_t = \frac{1}{6} \sum_{j=1}^6 \tilde{C}_{jt}$$

with components:

- C_1 — cess and surcharge share of gross tax revenue. Cesses are not shared with states under the Finance Commission formula, so a rising cess share centralises revenue without any change in the formula itself.
- C_2 — one minus the devolution share to states.

- C_3 — one minus states' own-revenue share.
- C_4 — conditional transfer (centrally-sponsored scheme) share, which constrains state discretion over expenditure.
- C_5 — borrowing restriction, coded on {0, 0.5, 1}.
- C_6 — GST structural centralisation, on [0,1]: zero before the 2017 rollout, phasing to full weight once GST compensation to states ended in June 2022.

The inclusion of C_6 is a substantive judgement and we flag it as such. Cesses and conditional transfers are symptoms of centralisation that operate within the existing constitutional settlement. The Goods and Services Tax is a different kind of object: it removed states' capacity to set indirect tax rates independently, replacing it with a collective body in which the union holds an effective veto. Treating that as merely one more component alongside cess share would understate it; treating it as the regime shift it is, is why it enters with a phase-in tied to the end of the compensation guarantee.

Because normalisation spans both eras, **the FCI is relative**. A value of 0.00 identifies the least-centralised year in the sample (2004) and 0.92 the most (2020). The numbers carry no meaning outside this sample and should not be compared to a similarly-named index computed elsewhere.

5.4 The Democratic Quality Index

The DQI is the geometric mean of four published third-party measures, each mapped to [0,1]:

$$\mathrm{DQI}_t = \left(V_t \cdot \frac{F_t}{100} \cdot \frac{180 - R_t}{180} \cdot C_t \right)^{1/4}$$

where V is the V-Dem Liberal Democracy Index, F the Freedom House aggregate score, R the RSF press-freedom rank, and C the V-Dem Core Civil Society Index.

Two choices warrant defence.

Why a geometric mean. An arithmetic mean permits compensation: a collapse on one dimension can be offset by stability on another, yielding a moderate composite that describes no actual state of affairs. The geometric mean penalises imbalance, which is the correct behaviour for a concept like democratic quality where the dimensions are closer to jointly necessary than to substitutable. A country with a free press and no independent judiciary is not "averagely" democratic.

Why civil society is included. The RSF press-freedom rank barely separates the two eras — India ranked poorly throughout, at 140 in 2014 and 157 in 2026, having been as low as 161 in 2023 and as high as 105 during the UPA period. A composite built only on headline democracy scores and press rank would therefore understate the change. The V-Dem Core Civil Society Index, by contrast, falls from 0.669 to 0.302 over the period. That is the dimension along which the change is sharpest, and omitting it would flatten precisely the phenomenon under study.

We note the risk in this choice: including the component that moves most is a decision that increases the measured effect, and a sceptical reader is entitled to ask whether it was selected for that reason. Our answer is that the inclusion is theoretically motivated — associational freedom is a standard component of liberal-democratic quality, not an exotic addition — and that we report the DQI's behaviour without it in Section 8.2.

6. Results

6.1 Growth

Real GDP growth averaged 7.08 per cent across 2004–2013 and 6.08 per cent across 2014–2026. The later period contains the pandemic contraction of –7.3 per cent in 2020, which accounts for a substantial part of the gap; excluding 2020, the later-period average is materially higher, though still not above the earlier one.

The more important observation is that the growth figures for the two periods are not straightforwardly comparable, for the reason set out in Section 2: the 2015 rebasing altered the measurement of manufacturing output in a manner that lowered measured UPA-era growth and raised measured NDA-era growth, and the resulting back-series was never independently reviewed. Subramanian's (2019) estimate of a 2.5-percentage-point overstatement for 2011–17, if accepted, would place NDA-era growth materially below the UPA average. We do not adopt that correction — it is contested, and the EAC-PM rebuttal is on record — but we decline to present the two periods' growth rates as a clean comparison, because they are not one.

What can be said without relying on the disputed revision is that the later period did not deliver an *acceleration*. Whatever the trade of institutional constraint for economic dynamism was supposed to purchase, a higher growth rate is not visible in the data even before adjusting for a revision that runs in the government's favour.

6.2 Employment

The employment result is the clearest in the dataset, because it does not depend on the disputed national accounts at all.

The formal employment share — workers with any social-security cover — falls from 13.0 per cent in 2014 to 10.9 per cent in 2026. This is a decline in *share*, not merely in growth rate, and it is monotone or flat in every single year of the period. Informal employment correspondingly rises from 87.0 to 89.1 per cent.

Headline unemployment moves from 4.9 per cent in 2014 to 5.5 per cent in 2026, having recorded 3.1 per cent as recently as 2025 — a series whose volatility partly reflects definitional treatment of self-employment rather than labour-market change. Graduate unemployment stands at 26.5 per cent in 2026, against 18.5 per cent in 2014.

The graduate figure deserves emphasis. An unemployment rate above one in four among degree-holders, in an economy whose headline unemployment rate is 5.5 per cent, is not a story about aggregate labour demand. It describes an economy that has expanded low-productivity informal work while contracting the formal salaried positions that education is an investment in obtaining.

The conjunction of positive aggregate growth with a falling formal-employment share over twelve consecutive years is the single most consistent multi-year pattern in the dataset. It is difficult to reconcile with any account in which the period's institutional changes purchased broad-based material progress.

6.2.1 The capital intensity of the growth that occurred

Aggregate growth figures conceal the composition of what grew, and the composition explains the employment result.

Industrial policy in the period concentrated on capital-intensive sectors through Production-Linked Incentive schemes. The employment arithmetic of that choice is unfavourable by construction. Electronics manufacturing under PLI generates on the order of one job per ₹1.5 crore of investment; garment manufacturing, at comparable investment, generates employment approximately two orders of magnitude larger, and does so with a workforce roughly 70 per cent female — a composition consequential in an economy whose female labour-force participation rate stands near 21 per cent.

This is not an argument that electronics manufacturing is a mistake. It is an observation that an industrial strategy weighted toward capital intensity will raise output more than payroll, and that the observed pattern — output growing, formal employment share falling — is the predictable consequence rather than an anomaly requiring separate explanation.

The scale of unmet demand for formal employment is visible in a different series. Government job recruitment in the period drew approximately 22 crore applications against approximately 7.22 lakh positions filled — a success rate of roughly one in 305. Applications on that scale, for positions that are frequently modestly paid, are a revealed measure of how scarce secure formal employment has become relative to the number of people seeking it.

6.2.2 Structural transformation that did not occur

The standard development sequence moves labour out of agriculture into manufacturing, and subsequently into services, with sectoral employment shares converging toward sectoral output shares. India's current composition departs from this sequence sharply:

SECTOR	SHARE OF OUTPUT (GVA)	SHARE OF EMPLOYMENT
Agriculture	18%	42%
Manufacturing	17%	11%
Services	53%	33%
Other	12%	14%

Agriculture retains 42 per cent of employment while producing 18 per cent of output — a productivity gap that is the standard signature of surplus labour held in low-productivity work for want of alternatives. Manufacturing, the sector that in the canonical sequence absorbs that labour, holds only 11 per cent of employment against 17 per cent of output. Services produce over half of output with a third of employment.

The implication is that India's growth has been concentrated in sectors that do not absorb labour at scale, while the sector holding most of the labour has not been the locus of growth. This is a structural fact about the growth model, and it is the mechanism underlying the employment result in Section 6.2. It also bears on the distributional result: growth concentrated in high-productivity, low-employment sectors accrues disproportionately to capital and to a small skilled workforce, which is what the income-share movements in Section 6.3 record.

6.3 Distribution

The top 1 per cent income share rises from 21.3 per cent in 2014 to 23.0 per cent in 2026. The bottom 50 per cent share falls from 15.0 to 12.9 per cent over the same period.

Both movements are drawn from the World Inequality Database's India series. We flag a correction here that materially affects how this result should be read, and that we found in our own earlier public materials during preparation of this paper: three of the project's summary documents had stated the 2014 top 1 per cent share as 15 per cent, yielding an apparent rise of 7.6 percentage points. That figure was wrong — 15.0 per cent is the 2014 *bottom 50 per cent* share, and the two series had been crossed. The correct rise in the top 1 per cent share is 1.3 percentage points to 2023, or 1.7 to 2026.

We report this because the error inflated a headline claim in the project's favour, and because correcting it changes which fact carries the argument. The defensible statements are these. First, the *level* of the top 1 per cent share is the highest recorded in the WID series since 1922 — a period which includes the peak of the colonial extraction the comparison invokes. Second, the movement at the bottom is considerably larger in relative terms than the movement at the top: the bottom half lost roughly an eighth of its share of national income in twelve years. The distributional story is better told from the bottom of the distribution than the top, and the corrected figures tell it more accurately than the erroneous ones did.

6.4 Statistical suppression

The SSI is 0.00 for every year from 2004 to 2013.

It first moves in 2015, reaching 1.05, on the national-accounts rebasing. It rises through 2.05 (2017), 4.00 (2018) and 4.80 (2019) as the PLFS withholding and the consumption-survey suppression enter, reaches 7.00 in 2020, and peaks at **9.00 in 2021 and 2022**.

The peak is a stacking effect rather than a single event: the COVID mortality undercount at full severity, on top of an already-delayed census, an already-withheld consumption survey, an unresolved back-series dispute, and the by-then-permanent NSO merger.

The index then declines — 8.20 in 2023, 7.05 in 2024, 6.55 in 2025 — and settles at **6.40 in 2026**. It does not return to zero, and this is the persistence term working as designed. The census remains outstanding. The NSO merger is not reversed. The back-series was never independently reviewed. What decays is the acute phase; what remains is the structural change.

The cross-era contrast is the starkest in this paper: an era mean of 0.00 against 5.09.

6.5 Fiscal centralisation

The FCI rises from 0.00 in 2004 to 0.17 by 2013 — a modest drift within the UPA decade, driven mostly by a rising cess share.

Within the NDA period it moves from 0.19 (2014) to a peak of **0.92 in 2020**, then settles at 0.68 from 2023 onward. The era means are 0.09 against 0.57.

The trajectory has clear inflection points corresponding to identifiable policy events: the GST rollout in 2017 (FCI moves 0.25 → 0.41), the pandemic-period borrowing restrictions and cess expansion in 2020, and the end of GST compensation to states in mid-2022, after which C₆ carries full weight and the index stabilises at a structurally higher level than anything recorded in the earlier era.

The post-2020 decline from 0.92 to 0.68 should not be read as decentralisation. It reflects the unwinding of emergency borrowing restrictions, not any return of taxing authority. The floor is higher than the UPA peak by a factor of four.

6.6 Democratic quality

The DQI averages 0.57 across 2004–2013, peaking at 0.61 in 2005, 2006 and 2009, and standing at 0.51 in 2013.

Within the NDA period it falls from 0.49 (2014) to **0.29 (2026)**, with the steepest single-year declines after 2019 — the year V-Dem reclassified India from "electoral democracy" to "electoral autocracy". The era means are 0.57 against 0.38; the endpoint comparison is 0.49 against 0.29, a decline of 40.8 per cent within the later period alone.

The component behaviour is instructive. India's V-Dem Liberal Democracy Index falls from 0.488 to 0.260 — a 46.7 per cent decline. The press-freedom rank moves from 140 to 157, a deterioration but not a dramatic one against an already-poor base. Freedom House's aggregate score falls, with the reclassification from "Free" to "Partly Free" in 2021. The V-Dem Core Civil Society Index falls from 0.669 to 0.302, a fall of 54.9 per cent and the steepest of the four components.

In international context, India's V-Dem decline over 2014–2026 (0.488 → 0.260, a fall of 0.228) is the largest absolute fall in the comparison cohort, ahead of Hungary (0.521 → 0.315, −0.206) and Turkey (0.252 → 0.110, −0.142). In proportional terms Turkey falls further, but from a base already close to the floor. Brazil (0.783 → 0.704) and South Africa (0.641 → 0.633) show the comparatively mild movement one might expect of large democracies under ordinary political stress.

6.7 Macroeconomic and external position

Two further dimensions bear on the "bargain" question, since a defender of the period might argue that its achievement lies in stability rather than in growth or distribution.

Fiscal. The central fiscal deficit averaged 4.96 per cent of GDP across 2014–2026 against 4.64 per cent across 2004–2013. The later period contains a genuine consolidation from 4.1 per cent (2014) to 3.4 per cent (2018), an emergency expansion to 9.2 per cent in 2020, and a partial return to 4.4 per cent by 2025–26 — a level still above the pre-pandemic trough and above the earlier era's average. The consolidation before 2020 was real; it was not sustained through and beyond the shock.

External. The external position as of FY26 is mixed rather than fragile. Foreign-exchange reserves of approximately \$691 billion provide around eleven months of import cover, which is comfortable by any conventional standard, and the current-account deficit is contained at approximately 1.0 per cent of GDP.

The composition beneath that headline is less reassuring. The goods trade deficit runs at approximately 333 billion, offset by a services surplus of approximately 190 billion — meaning the external accounts depend heavily on services exports to cover a structural merchandise gap. Crude imports of approximately 123 billion against an import dependence of roughly 86 per cent leave the position exposed to oil-price movements, and gold imports of approximately 72 billion represent household savings flowing into an unproductive asset, itself a signal about the perceived returns available in formal financial instruments. Portfolio outflows of approximately \$16.5 billion were recorded in the period.

The rupee depreciated from 61 to the dollar in 2014 to a record low of approximately 95 in 2026 — a decline of roughly 36 per cent. Some depreciation is expected for an economy with an inflation differential against the United States, and a nominal rate is not

by itself a welfare measure. But it bears directly on the crude and gold import bills above, and it is not consistent with an account in which the period delivered exceptional macroeconomic management.

Neither dimension supports a "stability" defence of the period. Fiscal performance is slightly worse on average than the earlier era; the external position is adequately buffered but structurally dependent on services exports to offset a widening merchandise gap.

6.8 International comparison

A decline measured only against a country's own past invites the response that the instrument, or the era, moved rather than the country. The comparative picture addresses this.

Across 2014 to 2026, on the V-Dem Liberal Democracy Index and the RSF press-freedom rank:

COUNTRY	V-DEM 2014	V-DEM 2026	CHANG E	ABSOLUTE FALL	PRESS RANK 2014	PRESS RANK 2026
India	0.488	0.260	−46.7 %	−0.228	140	157
Turkey	0.252	0.110	−56.3 %	−0.142	154	159
Hungary	0.521	0.315	−39.5 %	−0.206	64	70
Brazil	0.783	0.704	−10.1 %	−0.079	99	80
South Africa	0.641	0.633	−1.2%	−0.008	42	38

All values are the V-Dem Liberal Democracy Index (`v2x_libdem`) as published; the 2026 column carries V-Dem's 2025 observation, the latest available. An earlier version of this table mixed two different V-Dem indices between the two columns; see Section 9.9.

Three observations follow.

First, India records the largest absolute decline in the cohort — larger than Hungary's and larger than Turkey's, the country most often treated as the reference case for backsliding in a large middle-income state. Turkey's proportional fall is steeper, but Turkey entered the period at 0.252, already near the bottom of the scale, and a country close to the floor has less left to lose. India entered at 0.488 and gave up 0.228 of it.

Second, the comparison discriminates. Brazil and South Africa, large democracies subject to serious political stress over the same years, register modest movement. Whatever is happening in the V-Dem instrument globally, it is not producing large declines indiscriminately.

Third, the press-freedom rank behaves differently from the composite in every case, and this is why Section 5.4 declines to rest the DQI on it. Brazil's press rank *improved* while its V-Dem score fell; India's press rank worsened by 17 places while its V-Dem score fell by almost half. A rank is a positional measure across 180 countries and is insensitive to large absolute changes when many countries move together. It is a poor instrument for exactly the question at issue here.

6.9 The cross-era scorecard

Comparing period means across the nine directly comparable measures:

MEASURE	UPA (2004-13)	NDA (2014-26)	DIRECTION
GDP growth (%)	7.08	6.08	Worse
CPI inflation (%)	8.05	4.98	Better
Fiscal deficit (% GDP)	4.64	4.96	Worse
Top 1% income share (%)	20.23	22.17	Worse
Bottom 50% income share (%)	15.67	13.79	Worse
Press freedom rank	120.7	145.0	Worse
Statistical Suppression Index	0.00	5.09	Worse
Fiscal Centralisation Index	0.09	0.57	Worse
Democratic Quality Index	0.57	0.38	Worse

Eight of nine measures deteriorate. One improves.

The exception is real as a measurement: average CPI inflation of 4.98 per cent against 8.05 per cent is a large difference in a variable that bears directly and regressively on household welfare. Section 6.10 takes it in full, because a paper that reports eight results in its own direction and disposes of the ninth in a sentence has not earned the eight.

6.10 The inflation exception, in full

The disinflation is the strongest available argument against this paper's reading of the period, and it deserves to be met rather than managed. We set out six factors, roughly in order of how much of the gap we think each explains, and then the strongest case we can construct against our own conclusion.

First, the comparison straddles the collapse of a global oil cycle, and this is the largest single factor. Brent crude fell from about 115 a barrel in June 2014 to below 30 by January 2016, a decline of roughly 75 per cent in nineteen months, driven by United States shale supply arriving at scale and OPEC's decision to defend market share rather than cut production. India imports approximately 86 per cent of the crude it consumes; energy enters the consumer price index directly through fuel and light and indirectly through the transport cost of everything else. The timing carries the argument: the shock does not merely fall inside

the NDA period, it arrives almost exactly at the boundary between the two periods being compared. Any price measure sensitive to world energy will break at 2014 whether or not anything changed in Delhi.

Second, and less often noted, the government absorbed much of that fall rather than passing it through. Central excise duty on petrol stood at ₹9.48 a litre and on diesel at ₹3.56 a litre in 2014. Between November 2014 and January 2016 it was raised on nine separate occasions — by ₹11.77 a litre on petrol and ₹13.47 on diesel — with the explicit purpose of transferring the gains from falling world prices to the exchequer rather than to the pump. Two implications follow. The measured disinflation is the residual after the tax take, so it understates the size of the imported shock. And to the extent that fiscal policy acted on fuel prices in this window, it acted to hold them *above* what world prices would have delivered. That is a defensible fiscal choice; it is not an anti-inflation achievement.

Third, the baseline is not neutral either. The UPA's second term coincided with the 2008 food and fuel spike, the commodity boom running through 2013, and the global stimulus that followed the financial crisis; Indian inflation printed 10.9 per cent in 2009, 12.0 in 2010 and 10.9 in 2013. Setting an era containing a boom against an era containing its collapse yields a difference that is substantially a statement about two adjacent decades of world prices.

Fourth, where domestic policy did contribute, the mechanism supports this paper's thesis rather than qualifying it. India adopted formal inflation targeting in this period, and it is a genuine institutional change: the Monetary Policy Framework Agreement was signed on 20 February 2015, the Finance Act 2016 amended the Reserve Bank of India Act 1934 to give it statutory backing and to create a Monetary Policy Committee constituted that September, and the target is 4 per cent with a two-point tolerance band. There is a reasonable case that this anchored expectations, and we do not dismiss it. But note what kind of instrument it is. Inflation targeting is a rule that *binds* the executive — it moves an instrument out of the discretionary reach of whichever government holds office, assigns it to a committee with a published mandate, and compels a written explanation to Parliament when the target is missed. If this is the channel, then the period's best macroeconomic outcome was produced by deliberately constraining executive discretion, which is the argument this paper makes throughout. It is an instance of the thesis, not an exception to it.

Fifth, the most recent figure is not on the same basis as the rest of the series. As set out in Section 4.3, from the January 2026 print the 2012=100 CPI was replaced by a series on 2024=100, weighted from HCES 2023-24 and classified on COICOP 2018. Food's weight falls from 45.86 to 36.75 per cent, so identical food-price movements now produce a smaller headline number. Reweighting was overdue and is not a manipulation, but the 2026 value and the 2014–2025 values are not measured on the same ruler.

Sixth, and stated as a limitation rather than a finding, a national average conceals distribution: a household spending most of its income on food faces something much nearer food inflation than the headline, and the reweighting above mechanically reduces food's presence in the headline for everyone. Establishing what inflation each income decile actually faced across 2004–2026 would require expenditure-weighted indices by decile. We have not constructed them and do not claim the result.

6.10.1 The strongest case against this section

The honest counter-argument is that oil fell for every importer, yet not every oil-importing economy brought inflation from double digits to the low single digits and then held it there for a decade. Something domestic plausibly did work, and there are candidates beyond monetary policy: more competent buffer-stock and food-price management, and fiscal consolidation between 2014 and 2018. That case has force and we have not refuted it.

What we claim is narrower. The raw cross-era gap cannot be read as a policy result, because a shock large enough to account for most of it arrived precisely at the era boundary; and the leading domestic candidate is a constraint on executive power rather than an exercise of it.

6.10.2 What would change our reading

A consistent measure of core inflation — excluding food and fuel — showing a fall of comparable magnitude across the two eras would be strong evidence of something beyond the commodity cycle, since it strips out the channel through which the oil shock runs. This paper does not carry a consistent core series back to 2004 and therefore cannot run that test; we will not imply that it has been run. Comparative evidence would move us similarly: if India's disinflation is materially larger than that of comparable oil-importing emerging economies over the same window, the residual requires a domestic explanation.

We therefore decline to read this row as an achievement of the period's governance, and we would resist anyone else doing so on our data. The deterioration on accountability, distribution and statistical integrity is uniform; the single measure moving the other way is one where the dominant driver was a world oil price, the fiscal response captured much of it, and the domestic institution with the best claim to have helped works by binding the executive's hands rather than freeing them.

6.11 What GDP does not count

Every growth figure in this paper is a GDP figure. GDP does not measure the work an economy does; it measures work inside a production boundary set by the System of National Accounts. That boundary is a convention, and the way to see it is to look at what it already crosses. Where a household owns the dwelling it occupies, the accounts estimate the rent it would otherwise have paid and add that sum to output, on the reasoning that a real service is being consumed. Cooking a meal in the same dwelling, minding a child, nursing a parent: the reasoning is not applied. Own-account household services fall outside the boundary and enter the accounts at zero. The framework prices the building and puts no price on the person working inside it.

The feasibility defence has run out in the Indian case, because the hours are counted. The National Statistical Office ran Time Use Surveys over January–December 2019 and January–December 2024, the later covering roughly 4.5 lakh people in 1.3 lakh households. Among those aged 15–59 who perform it at all, women spent 289 minutes a day on unpaid domestic services in 2024 and men 88. For a woman that is about 1,758 hours a year, appearing in no employment statistic and no output series. Care of household members sits above that: 41 per cent of women in the age band participated against 21.4 per cent of men, at 140 minutes against 74.

Two rounds allow the change to be measured. Male participation rose steeply, from 26.1 per cent in 2019 to 45.8 per cent in 2024. Minutes among participants fell for both sexes over the same period, men 97 to 88 and women 299 to 289. The participant gap therefore moved from 202 minutes to 201. One minute across five years, while female participation in employment rose from 21.8 to 25.0 per cent. More women entered paid work and the unpaid work stayed close to where it was.

Published Indian valuations of unpaid household work span 7.5 per cent of GDP (SBI Ecowrap 2023, minimum-wage valuation), a government working range of 15–17 per cent, 24.6 per cent on generalised opportunity cost and 32.4 per cent on replacement cost for 2019–20, and 26–36 per cent in *Economic and Political Weekly* (2024) for 2022–23. The spread is a factor of nearly five and it is not uncertainty about the hours, which are known. It is disagreement about the price to attach to them. We adjudicate between none of these and publish no preferred figure. Two things can be said without choosing a method: all of the estimates are large, the

lowest exceeding several sectors that are counted; and the question stays unsettled because no statistical agency has been asked to settle it.

This bears on a figure used elsewhere in the paper. Female labour force participation of roughly 21 per cent is customarily read as an absence of work. The Time Use Survey records 83.9 per cent of women aged 15–59 performing unpaid domestic work on the reference day. Both are correct. Together they indicate that the participation figure measures the absence of work the accounts recognise, and how much of India's participation gap is behavioural rather than definitional is an open question this paper has not answered.

We adjust nothing. Every growth series here is the official one on the official boundary. The claim concerns the boundary: it was drawn by people, India already collects the hours needed to price what falls outside it, the excluded quantity is plausibly between a thirteenth and a third of GDP, and a framework pricing a dwelling while entering zero for the care performed within it has settled a question about whose work constitutes production. That sits with the other measurement problems documented in Section 3. Nobody concealed it and it long predates the period under study. It is also the largest single omission from the headline figure. A household-production satellite account, which several statistical systems maintain and India's does not, would replace this section with a number.

7. External validation

7.1 The problem of a self-generated result

A composite index built by one researcher, from partly self-coded inputs, showing a result congenial to that researcher's stated thesis, is exactly the kind of finding that should attract scepticism. We share that scepticism, and the appropriate response is to ask whether anyone reaches a comparable conclusion without using our method, our data vintage, or our framing.

7.2 Grier and Grier (2026)

Kevin Grier and Robin Grier, of Texas Tech University, address the same period in "Promises, Promises: Modi's India in Comparative Perspective" using synthetic control with randomisation inference. Their design shares none of our machinery.

Where we build composite indices from observed series, they construct a counterfactual: a weighted blend of comparator countries fitted to India's pre-treatment path, against which post-2014 India is compared. Their pre-treatment window is 1984–2013, roughly thirty years, against our 2004–2013. Their governance data is V-Dem — an input we share — but their income data is the Penn World Table 11.0, which we do not use at all. Their democracy donor pool comprises twelve middle-income democracies: Brazil, Argentina, Chile, Colombia, Peru, Mexico, South Africa, Ghana, Indonesia, Malaysia, the Philippines, and Sri Lanka.

Their governance results, expressed as final-year gaps between India and Synthetic India as a percentage of India's own 2013 level:

INDICATOR	GAP	% OF 2013 LEVEL	JOINT P
Freedom of religion	−1.407	−168.7	0.00
Liberal democracy	−0.360	−69.1	0.00
Electoral democracy (polyarchy)	−0.391	−60.2	0.00
Freedom of expression	−0.488	−56.2	0.00
Political corruption	+0.272	+46.8 (worse)	0.00
Equality before the law	−0.221	−30.4	0.00
Freedom of association	−0.232	−29.6	0.00
Legislative constraints	−0.242	−28.8	0.00
Equal protection	−0.190	−26.1	0.00
Judicial constraints	−0.105	−12.7	0.00

All ten indicators move in the direction of worse governance, at a joint p-value of 0.00 in every case.

7.3 Convergence, and what it does and does not show

The correspondence with our results is close:

DIMENSION	GRIER & GRIER	THIS PAPER
Electoral democracy	−60.2% vs counterfactual	V-Dem LDI 0.488 → 0.260, −46.7%
Overall democratic quality	−69.1% liberal-democracy gap	DQI 0.49 → 0.29, −40.8%
Press and expression	Freedom of expression −56.2%	RSF rank 140 → 157

These are not the same quantity, and we want to be precise about that. Their figures are gaps against a counterfactual India that never elected the government it did; ours are observed changes over time. A gap of −60.2 per cent and a raw decline of −46.7 per cent describe different objects, and their rough proximity is not itself evidence of anything.

What the convergence does establish is more modest and more useful: that two research designs with almost no shared machinery, applied to the same country and period, agree on direction, on rough magnitude, and on which dimensions moved most. Our DQI result is not an artefact of our weighting scheme, our geometric mean, or our decision to include the civil-society index, because a study using none of those things finds the same thing.

7.4 Their robustness checks

The convergence argument in Section 7.3 is only as good as the study it leans on, so the robustness of their design matters to us as well as to them. Grier and Grier report several checks, and we summarise those relevant to the governance result we rely on.

Placebo in time. They re-run the entire procedure pretending treatment began in 2001 and ending the sample in 2010 — a period in which no treatment occurred. If the method manufactured gaps as an artefact, it would manufacture one here. It does not: no individual placebo effect is significant and the joint p-value is 0.5, placing the fake-treatment India squarely in the middle of the fake-effect distribution. This is the check that most directly addresses the concern that synthetic control finds effects where none exist.

Alternative donor pools. The governance results are re-estimated with a donor pool selected on economic rather than political criteria, and again with an alternative pool. The results survive both.

Leave-one-out. Removing individual donor countries does not overturn the governance findings.

Income robustness. For the income result, they report a specification excluding Ethiopia — the largest single donor weight, and the component we flagged in Section 7.5 as the most exposed to challenge. Excluding it produces a *larger* post-2013 gap, not a smaller one. This does not fully dispose of the donor-pool objection, since the concern is about the pool's overall composition rather than any single member, but it does address the most obvious form of the criticism.

We report these because a reader evaluating our convergence claim should be able to judge whether the study we are converging with is itself sound, rather than taking our word that it is.

7.5 Where the two studies diverge

Grier and Grier also report an income result: real per-capita income ending approximately 1,000 per person per year below their counterfactual, widening from approximately -400 in the first year, on a pre-treatment fit with an RMSPE of \$59, about 1 per cent of 2013 income. That is a very tight pre-treatment match.

We do not adopt this claim, and our own work does not test it. This paper documents the decoupling of growth from employment and the concentration of its gains; it makes no claim about what aggregate Indian income would have been under a different government. That is a counterfactual we have not constructed and cannot evaluate.

We would add one observation about their design, offered as a reader's caution rather than a rebuttal. Their income donor pool is weighted 38 per cent Ethiopia, 28 per cent China, 25 per cent Bangladesh, 7 per cent Pakistan, and 2 per cent Philippines. That combination is a defensible fit to India's pre-2014 income trajectory — and the RMSPE indicates it fits very well — but the choice of donor pool is the most consequential and most contested design decision in synthetic control, and a pool leaning this heavily on Ethiopia and Bangladesh is the part of their design most exposed to challenge. Readers should treat their governance finding as strongly corroborated by ours, and their income finding as a separate and independently contestable claim.

8. Robustness

8.1 Sensitivity of the SSI to severity coding

The SSI's severities are author-coded, making it the most subjective construction here. We assess sensitivity by re-computing under alternative codings.

Because the index is a weighted sum with weights summing to 10, a uniform proportional error in severity coding scales the index without changing its shape or the cross-era comparison: if every severity were overstated by 20 per cent, the NDA-era mean would fall from 5.09 to 4.07 and the UPA mean would remain 0.00. The cross-era conclusion is invariant to any uniform mis-calibration.

The conclusion is sensitive, in principle, to *differential* mis-coding — systematically overstating NDA-era severities relative to UPA-era ones. But the UPA-era value is zero because no coded event occurred, not because the events that occurred were coded leniently. Differential mis-coding would therefore have to take the form of omitting UPA-era events entirely, which is the objection we address in Section 9.2.

Dropping the single highest-weighted stream (census delay, weight 2.5) reduces the NDA-era mean from 5.09 to approximately 3.6 and lowers the 2021–22 peak from 9.00 to 6.50. The cross-era contrast survives the removal of any single stream, because no single stream accounts for more than a quarter of the total weight.

8.2 Sensitivity of the DQI to component selection

The DQI's most questionable choice is the inclusion of the V-Dem Core Civil Society Index, which is also the fastest-falling component.

Computed as a geometric mean of the remaining three components, the index still declines across the NDA period, though less steeply: the fall is driven substantially but not exclusively by civil society. The V-Dem Liberal Democracy Index alone falls 46.7 per cent over the period, so a DQI excluding the civil-society term still registers a large deterioration through that channel.

We report this rather than resolving it, because the honest position is that the four-component index shows a larger decline than a three-component one would, and readers should know that the choice affects the magnitude. It does not affect the sign, and it does not affect the cross-era ranking.

8.3 The relative construction of the FCI

Because the FCI's components are min–max normalised over the full sample, the index is guaranteed to have a minimum of approximately 0 and a maximum of approximately 1 somewhere in 2004–2026. It cannot, by construction, report that centralisation was uniformly low or uniformly high.

What it can report is *where* in the sample the extremes fall, and that is the informative content: the minimum falls in 2004 and the maximum in 2020, with a structural break at the 2017 GST rollout. A reader who rejects the normalisation can inspect the underlying components, all of which are published in absolute units.

8.4 The disputed national accounts

We do not adopt Subramanian's (2019) correction, and our growth comparison in Section 6.1 uses official figures throughout. Adopting the correction would strengthen our finding, since it would lower measured NDA-era growth further. Our decision not to adopt it means our growth result is conservative relative to the alternative in the literature.

9. Limitations and objections

9.1 This is descriptive, not causal

Nothing in this paper identifies a causal effect of institutional change on economic outcomes. We document co-movement across two periods. Both the institutional and the economic variables are plausibly driven by common factors — global conditions, demographic transition, technological change — and we make no attempt to separate these.

The specific claim we defend is narrower than a causal one and, we think, sufficient for the purpose: that the *bargain* interpretation of the period fails on its own terms. The bargain requires that constraint was traded for prosperity. What the data show is constraint reduced and, on eight of nine measures, prosperity not delivered. That conclusion does not require a causal identification; it requires only that both halves of the alleged trade be measured, which is what we have done.

9.2 The objection to SSI = 0 for the UPA decade

This is the strongest objection to the paper and we treat it at length.

An index that assigns exactly zero to one political era and a mean of 5.09 to the other invites the suspicion that the coder found what they were looking for. Indian statistical administration before 2014 was not beyond criticism: there were delays, methodological disputes, and instances of political discomfort with unwelcome findings.

Three responses, in ascending order of force.

First, the SSI codes a specific and narrow class of event: suppression, discontinuation, indefinite postponement, or a live and unresolved methodology dispute affecting a major official statistical product. It does not code ordinary administrative delay, quality problems, or resource constraint. On that definition, we are not aware of a UPA-era event that meets the threshold, and we have looked. The 2004–2013 decade ran its census on schedule, published its consumption surveys, and did not merge its statistical office into the ministry it reports on.

Second, the zero is a falsifiable claim, not an assumption. The rubric is published; the streams are enumerated; the weights are fixed. A critic who identifies a qualifying UPA-era event can code it and recompute, and the code to do so is public. We would regard such a contribution as an improvement to the index rather than a refutation of the paper.

Third, and most importantly, the cross-era conclusion does not depend on the zero. Suppose a critic assigns the UPA decade a mean of 2.0 — a substantial concession, implying suppression activity at 40 per cent of the NDA average. The comparison would still be 2.0 against 5.09. The result is a difference in kind that survives a generous re-coding.

We nonetheless record that the zero is the single most contestable number in this paper, and that a reader who rejects it is applying appropriate scrutiny.

9.3 The pandemic confound

The comparison window contains a global pandemic, and this is the most serious threat to the cross-era comparison after the SSI coding question.

COVID-19 affected both halves of the paper's argument simultaneously. On the economic side it produced a -7.3 per cent contraction in 2020, disrupted labour markets, and drove a fiscal deficit of 9.2 per cent of GDP. On the institutional side it supplied a legitimate justification for postponing a census, for emergency restrictions on assembly and movement, and for centralised control of borrowing — measures that a democratic government facing a genuine emergency might reasonably take, and that our indices register as deterioration.

We take three steps in response, and concede that none fully resolves it.

First, the arithmetic. Excluding 2020 and 2021 entirely from the NDA period raises the mean growth rate materially, but the direction of every other measure is unchanged: the formal employment share falls monotonically in the years on either side of the pandemic, the income-share movements are continuous through it, and the SSI, FCI and DQI are all substantially elevated in both 2018–19 and 2023–26 relative to any UPA year. The pandemic amplifies the pattern; it does not create it.

Second, the timing. The DQI's steepest decline begins in 2019, before the pandemic, and V-Dem's reclassification of India as an electoral autocracy is a 2019 event. The consumption-survey suppression is 2019, the PLFS withholding 2019, the NSC resignations January 2019, and the GDP revision 2015. Four of the six SSI streams pre-date the pandemic entirely.

Third, the persistence. Emergency measures justified by an emergency should lapse when it does. The census had not been conducted five years after the acute phase; the NSSO–NSO merger was not reversed; the FCI settles at 0.68 rather than returning to its 2014 level of 0.19. The signature that distinguishes an emergency response from a structural change is whether it unwinds, and on these measures it did not.

What we cannot do is separate the portion of the 2020–22 deterioration attributable to the pandemic from the portion attributable to political choice, and we do not claim to. A reader who attributes the entire 2020–22 spike to the emergency would still face a pre-2020 deterioration and a post-2022 floor well above anything in the earlier era.

9.4 The indices are not independent of one another

The SSI, FCI and DQI move together, and a reader may reasonably ask whether three indices that co-move are really three pieces of evidence or one piece counted three times.

They share no inputs. The DQI is built entirely from four published third-party series; the FCI from fiscal aggregates; the SSI from coded statistical events. No number enters more than one index. But they are plausibly measuring correlated aspects of a single underlying process, and if that process is "concentration of executive authority", then their co-movement is the expected result rather than independent confirmation.

We think the honest characterisation is that the three indices provide *coverage* rather than *replication*: they establish that the phenomenon appears across the statistical, fiscal and political-rights domains, which a single index could not show. They do not provide three independent tests of one hypothesis. This is a further reason the external validation in Section 7 matters: it is genuinely independent in a way our own three indices are not of each other.

9.5 Expert-survey measures and reputational feedback

The DQI rests substantially on expert-survey instruments. A known critique holds that such measures may respond partly to the international reputational climate around a country rather than to conditions on the ground: as a country acquires a reputation for backsliding, coders may mark it down in ways that partly reflect that reputation.

If this mechanism operates, it would inflate the measured DQI decline. Three considerations bound the concern. V-Dem's coding protocol uses multiple country experts with explicit uncertainty modelling designed to mitigate exactly this. The RSF and Freedom House instruments are constructed differently and move in the same direction. And the Grier and Grier synthetic-control results in Section 7, which use the same V-Dem source, at least establish that India's movement is large *relative to comparable countries assessed by the same instrument* — which is the relevant comparison if the concern is instrument-wide drift rather than India-specific reputational effects.

The concern is not eliminated. It is bounded.

9.6 Composite indices and unearned authority

Any composite index makes choices — which components, which weights, which aggregation — that are not themselves derivable from data, and then presents the result as a single authoritative number. This is a general problem and we do not claim to have solved it.

Our mitigations are procedural rather than technical: every weight and component is published; the construction script is public and deterministic; the components can be inspected separately; and we have reported above where a choice materially affects the magnitude of a result. Readers who reject a choice can recompute without it.

9.7 Scope

This paper is about aggregate national indicators. It does not address state-level variation, which is substantial and in several respects runs counter to the national picture. It does not address caste, gender, or religious disaggregation of the distributional results, all of which are material to any full account of who bore the costs described here. And its coverage of the informal economy is limited to what the PLFS measures, which is a real constraint given that the informal sector accounts for approximately 89 per cent of Indian employment.

9.8 An error corrected in preparation

As reported in Section 6.3, preparation of this paper surfaced an error in three of the project's own previously-published summary documents, which had stated the 2014 top 1 per cent income share as 15 per cent — the bottom 50 per cent figure — thereby overstating the rise in top-end concentration as 7.6 percentage points rather than 1.3. The documents have been corrected and a

continuous-integration check now compares every year-labelled figure in the project's published documents against the source dataset, failing the build on any contradiction.

We report this in the paper rather than correcting it silently, for two reasons. The error ran in the direction of the project's thesis, which is precisely the kind of error that deserves visible correction rather than a quiet edit. And a project whose central argument concerns the integrity of published statistics has an obligation to hold its own output to the standard it applies to others.

9.9 A second error corrected, in the DQI's own inputs

A review of the index inputs immediately before publication found a second error, in the DQI itself, and it is reported here on the same principle as Section 9.8.

Two of the DQI's four components are V-Dem series: the Liberal Democracy Index (`v2x_libdem`) and the Core Civil Society Index (`v2xcs_ccsi`). The table used to compute the index held approximated rather than published values for both, and the approximations were systematically too high in the UPA decade. For 2014 the table carried a Liberal Democracy value of 0.555 against a published 0.488, and a civil-society value of 0.87 against a published 0.669. Because the error inflated the earlier baseline while leaving the later years close to correct, it exaggerated the measured decline. It ran, in other words, in the direction of this paper's argument.

Both columns are now the published V-Dem series verbatim. The effect on the results is as follows. The UPA-decade DQI mean falls from 0.59 to 0.57; the 2014 value falls from 0.54 to 0.49; the 2026 value is unchanged at 0.29. The within-period decline is therefore 40.8 per cent rather than 46.3 per cent, and the cross-era gap narrows from 0.18 to 0.19 index points on means that are themselves lower. The direction of the result, its size relative to the other two indices, and the cross-era scorecard of eight-of-nine are all unaffected.

The same error appeared in a second place. The international comparison table had its 2014 column on V-Dem's Electoral Democracy Index and its 2026 column on the Liberal Democracy Index — two different instruments read as one series. India's decline appeared as $0.71 \rightarrow 0.26$ when the correct like-for-like comparison on the Liberal Democracy Index is $0.488 \rightarrow 0.260$. This too overstated the finding: the proportional decline falls from 63.4 per cent to 46.7 per cent, and India ceases to have the steepest proportional fall in the cohort, though it retains the largest absolute one. Section 6.8 now reports the corrected figures.

Neither error was caught by the consistency check introduced after Section 9.8, because that check compares published documents against the dataset and both errors were *in* the dataset. A second check now runs on the dataset itself: it fails the build unless the international comparison table and the DQI's component table report the same V-Dem figure for India, which they can only do if both are on the Liberal Democracy Index. A crossed-series error of this shape cannot pass silently again.

10. Discussion

10.1 The bargain fails on its own terms

The strongest case for the period under study is that a decisive executive, unencumbered by the veto points that had constrained earlier Indian governments, could deliver material progress that a more constrained government could not. This is a serious argument with real intellectual lineage, and it deserves to be tested rather than dismissed.

Tested, it fails — not because the institutional costs were too high relative to the gains, but because the gains are not visible. Growth did not accelerate. The formal employment share fell every year for twelve years. The bottom half's income share fell by roughly an eighth. The one measure moving the other way, inflation, is largely a world oil price: Brent fell roughly 75 per cent across 2014–16 in an economy importing 86 per cent of its crude. Whatever domestic policy contributed came from a rules-based framework that binds the executive rather than empowers it.

This is not a finding that the trade was unfavourable. It is a finding that there was no trade.

10.2 Why statistical capacity is the load-bearing variable

Of the three indices, the SSI is the most novel, the most contestable, and in our view the most important.

Democratic quality and fiscal centralisation are measurable degradations of institutions that constrain a government. Statistical capacity is different in kind: it is the institution that makes the other degradations *visible*. When the census is postponed, the consumption survey withheld, and the national accounts revised without independent review, the effect is not only that particular numbers go missing. It is that the evidentiary basis for contesting any claim about performance erodes.

This has a self-reinforcing character worth stating plainly. A government that suppresses unfavourable statistics faces a weaker evidentiary challenge at the next election, which lowers the cost of further suppression. The SSI's persistence term is a formal acknowledgement of this: the damage does not reverse when the acute episode ends, because what was demonstrated was a willingness, and that demonstration does not expire.

It is also why the 2026 value of 6.40, rather than the 9.00 peak, is the number we would emphasise. The acute phase passed. The structural change did not.

10.3 Counting as an allocation of power

The census delay appears in this paper as the highest-weighted component of the SSI, and it is worth drawing out why a missing enumeration belongs in an argument about political authority rather than only in one about data quality.

Lok Sabha seats have been apportioned on 1971 population figures since the freeze imposed in 1976, which was itself a federal bargain: states that succeeded in reducing fertility would not be penalised with reduced representation. That freeze expires with the first census conducted after 2026.

The arithmetic of unfreezing is severe. Average constituency size has risen from roughly 0.7 million people per member in 1951 to roughly 2.5 million today, and it has risen unevenly, because fertility fell earlier and faster in the southern states. Projections

indicate a transfer of approximately eleven seats to Uttar Pradesh and ten to Bihar, with Tamil Nadu and Kerala each losing roughly eight.

Two consequences follow for this paper's argument. First, the census is not merely an input to welfare statistics; it is the instrument that determines the distribution of parliamentary representation across a federation of considerable regional divergence. Postponing it postpones that reallocation while leaving representation frozen on data now more than half a century old. Second, this connects the SSI directly to the DQI rather than leaving them as parallel measures of unrelated institutional decay. A statistical omission with a determinate effect on the distribution of political power is not only a statistical omission.

We make no claim about intent here, and none is required. The point is structural: whichever government eventually conducts the census inherits a reapportionment problem that grows more difficult the longer the enumeration is deferred.

10.4 Federalism as the underrated channel

The FCI result receives less attention than the democracy result, and probably should receive more.

The Indian constitutional settlement distributes fiscal authority between the union and the states, and that distribution is a substantial part of what makes a country of India's scale and heterogeneity governable. The movement documented here — cess share rising, devolution falling, conditional transfers expanding, and the GST removing independent state indirect taxation altogether — represents a shift in that settlement accomplished largely without constitutional amendment and largely without public debate framed in those terms.

The GST is the clearest case. It was argued, defensibly, on efficiency grounds. Its federal consequence — that states no longer set indirect tax rates independently — was a structural change to the constitutional balance effected through a tax reform. Whatever one concludes about the efficiency case, the federal consequence is a fact about the distribution of power, and it is permanent in a way that a fiscal-deficit number is not.

10.5 What follows for policy

We are wary of appending policy prescriptions to a descriptive paper, since the analysis does not identify causal effects and therefore cannot establish that a given intervention would produce a given result. What the analysis does support is a narrower claim: that certain measured deteriorations are directly addressable, and that the first of them is a precondition for evaluating any of the others.

Statistical independence is the prior reform. Every other question raised here — whether growth accelerated, whether poverty fell, whether welfare reached its intended recipients — is answerable only against evidence produced by institutions with sufficient independence to publish unfavourable findings. Restoring that capacity means conducting the census, publishing consumption surveys on a fixed cycle regardless of findings, and re-establishing the organisational separation that the NSSO–NSO merger removed. An independent review of the disputed back-series would resolve the largest outstanding question about the comparability of the two eras' growth records.

This ordering is not rhetorical. A government that implemented every other reform below, while leaving the statistical system as it is, would leave the public unable to determine whether any of them had worked.

Formalisation is where the employment result points. The finding is not that too few jobs exist but that the jobs increasingly lack social-security cover, written contracts, and the protections that make employment a basis for planning. The instruments are conventional: extension of EPF/ESI coverage, enforcement of written contracts, and a social-security floor for platform and gig work, where the workforce has grown from roughly 0.8 million to approximately 8.8 million over the period with substantially no coverage.

Industrial policy weighting bears re-examination. Section 6.2.1 is not an argument against capital-intensive manufacturing, but the employment arithmetic is what it is. A strategy seeking employment absorption at scale would weight labour-intensive sectors — garments, food processing, construction, care — considerably more heavily than the current allocation does.

Fiscal federalism requires an explicit settlement. The FCI documents a shift accomplished largely through instruments — cesses outside the divisible pool, conditional transfers, the GST structure — that individually attracted little debate in federal terms. Rebalancing means constraining the cess share of gross tax revenue, restoring the share of transfers that are untied, and addressing the position states occupy in GST rate-setting.

Delimitation needs to be settled before it arrives, not after. The seat reapportionment deferred by the census delay is arithmetically severe and regionally asymmetric. It will be substantially harder to negotiate under the time pressure of an imminent census than in advance of one.

The full set, separated by how quickly each could move and stated as directions rather than drafted policy:

Immediate. Constitutional protection for statistical independence; restoration of the states' share in the divisible pool; electoral-finance reform following the Electoral Bonds judgment; anti-concentration rules for media ownership; and restoration of civil-society licensing under the foreign-contribution regime, this last bearing directly on the V-Dem civil-society index whose collapse drives much of the DQI result.

Structural. An urban employment guarantee on the MGNREGA model, addressing the informalisation documented in Section 6.2; wealth taxation, restored and strengthened, against the concentration in Section 6.3; a federal council institutionalising the GST Council's consultative form across a wider set of union–state questions; transparency in judicial appointments; and statutory limits on surveillance and data collection.

We do not model the effects of any of these, and we are conscious that listing them invites a demand for evidence we have not supplied. They are offered as the directions the measurements point in, not as costed proposals.

10.6 A research agenda

This paper measures a national aggregate over twelve years. That is a narrow slice of what the question deserves, and the omissions are as informative as the findings. We set out what we think should follow, partly because several of these would test the account here rather than merely extend it.

Empirical extensions. The most valuable is state-level variation: India's states diverge substantially in governance quality, and if some resisted the national pattern, the reasons why are the closest thing to a natural experiment this case offers. Beyond that: sectoral analysis of which industries gained and lost; disaggregation of the distributional results by caste, religion and gender, without which "the bottom 50 per cent" conceals more than it reveals; district-level construction of the indices, where the aggregate

may be hiding wide dispersion; and media-ownership concentration, which bears on the press-freedom and civil-society components directly.

Methodological advances. Machine-learning imputation for the survey gaps documented in Section 5.2.1; satellite data as an independent economic proxy, valuable precisely because it is outside the control of the statistical system whose independence is in question; network analysis of firm–state relationships; and text analysis of judgments and policy documents to measure institutional posture rather than infer it.

The satellite point generalises. The deepest methodological problem this paper raises is that its subject matter degrades its instruments. Measures that do not depend on the national statistical apparatus — remote sensing, private-sector data, third-party surveys — are therefore worth more here than their usual accuracy would justify.

Comparative work. Cross-national comparison with other large democracies under similar pressure; historical comparison with the 1975–77 Emergency, which is the obvious internal benchmark and would test the claim in Section 2.4 that the current episode achieves comparable outcomes through incremental legal means; and subnational comparison across state governance models.

Policy research. Institutional design for resilience; early-warning indicators; the mechanics of reversal, on which there is much less literature than on decline; and the role of international accountability where domestic accountability has weakened.

One question runs underneath all of these, and we would put it to other researchers rather than pretend to have answered it: **how do you study the health of a democracy when the data about that democracy is itself among the things being suppressed?** This paper's answer — measure the suppression explicitly, and treat statistical capacity as an endogenous variable — is one approach. It is unlikely to be the best one.

10.7 What would falsify this

We think it is worth stating what evidence would overturn the paper's conclusions, since a thesis that cannot be falsified is not doing empirical work.

The cross-era institutional finding would be substantially weakened by documented UPA-era events meeting the SSI's coding threshold in sufficient number to lift that era's mean materially above zero. It would be weakened by a demonstration that the V-Dem, Freedom House and RSF movements are artefacts of instrument drift rather than conditions — though this would need to account for the synthetic-control corroboration.

The distributional finding would be weakened by a revision to the WID India series, which is itself estimated from combined survey, fiscal and national-accounts sources and is not beyond revision.

The employment finding is the most robust, because it does not depend on the disputed national accounts, on any index we constructed, or on any expert survey: it is a direct survey measure of the share of workers with social-security cover, and it fell in every year of the period.

11. Conclusion

India between 2014 and 2026 did not exchange democratic quality for prosperity. It reduced democratic quality, centralised fiscal authority, degraded the statistical apparatus by which any of this might be independently assessed, and did not deliver superior growth, employment, or distribution. On eight of nine comparable cross-era measures the later period performs worse than the earlier one. The single exception, inflation, is best explained by a global commodity collapse rather than by Indian policy; insofar as policy mattered, it acted by binding the executive rather than empowering it.

The governance component of this finding is corroborated by an independent study using a different method, different data, and a different pre-treatment window, which finds all ten of the governance indicators it tests deteriorating against a synthetic counterfactual at a joint p-value of 0.00.

We have been explicit about what is not established: no causal claim is made; the SSI's UPA-era zero is the most contestable number here and we have set out the objection to it at length; the DQI inherits the assumptions of its third-party sources; the FCI is meaningful only within its own sample. We have also reported an error in the project's own prior materials that ran in the direction of its thesis, and corrected it.

The most consequential finding is the one that is least often quantified. A government's performance can be debated only against evidence, and the evidence-producing institutions were among the things that degraded during the period. That is not a secondary or procedural concern. It is the mechanism by which every other question in this paper becomes harder to answer next time — and the reason we regard the Statistical Suppression Index, whatever its imperfections, as the most important of the three constructions offered here.

Declarations

Funding

This research received no funding from any institution, government, political party, foundation, or commercial entity. It was conducted independently and at the author's own cost. The absence of funding is stated here not as a virtue but because, in a paper whose subject is the independence of research-producing institutions from the governments they assess, the author's own position is a material fact for the reader to weigh. Absence of funding is not absence of interest; the author's political affiliation is declared immediately below.

Competing interests

The author is a Manmohan Singh Fellow of the All India Professionals' Congress (AIPC), a department of the Indian National Congress. This paper compares the performance of a Congress-led government (UPA, 2004–14) against that of its successor (NDA, 2014–26) and reports against the successor on eight of nine comparable measures. The author therefore holds a political interest in the direction of the finding, and that interest is declared here in full rather than left to be discovered.

The author declares no competing financial interests. No party body commissioned, funded, reviewed, or approved this work, and no part of it was seen by any such body before publication.

The mitigation offered is not a claim of neutrality, which would be false, and which a paper arguing that institutional independence is a structural rather than a personal property has no business asserting about its own author. It is that every step between the raw series and the conclusion is inspectable. The complete dataset, the construction code, and the source of this paper are published; both eras are computed by the same script, so no part of the measured gap can originate in the two periods being handled differently; the index weightings are stated and their sensitivity reported in Section 8; the one measure that improved under the second period is reported as improved and is explicitly not attributed to the government's credit, and is argued out at length rather than disposed of (Section 6.10); an error that had run in the project's own favour is documented in Section 9.8; and the governance result is corroborated in Section 7 by an independent study using a method this paper does not use. A reader who distrusts the author's motives can recompute every figure under different assumptions and reach a different conclusion from the same published inputs. That is the only defence a partisan author can honestly offer, and it is offered here.

Data and code availability

All data, construction code, and the source of this paper are published under CC BY 4.0 at github.com/Varnasr/someperspective and someperspective.info. No data used in this paper is restricted, proprietary, or unavailable to a replicating researcher. Third-party series (V-Dem, Freedom House, RSF, WID, Penn World Table) are obtainable directly from their publishers under those publishers' own terms.

Use of generative AI

Drafting, code generation for the site's analysis tooling, and consistency checking of published figures against the source dataset were assisted by a large language model. All figures in this paper were verified programmatically against `data.json`; all substantive claims, index designs, weightings, and interpretations are the author's. The error described in Section 9.8 was identified during that consistency checking and is reported rather than silently corrected.

Version

Version 1.0, data vintage 15 August 2026. Subsequent revisions, including any corrections, are recorded at someperspective.info/updates.html and in the repository's changelog.

References

Bharti, N.K., Chancel, L., Piketty, T., and Somanchi, A. (2024). *Income and Wealth Inequality in India, 1922–2023: The Rise of the Billionaire Raj*. World Inequality Lab Working Paper 2024/09.

Coppedge, M., Gerring, J., Knutsen, C.H., Lindberg, S.I., Teorell, J., et al. *V-Dem Dataset v15*. Varieties of Democracy Institute, University of Gothenburg, 2026.

Economic Advisory Council to the Prime Minister (2019). *Rebuttal to "India's GDP Mis-estimation"*. Government of India, June 2019.

Freedom House. *Freedom in the World*. Annual editions, 2004–2026.

Government of India, Ministry of Finance. *Union Budget documents and Receipt Budget*. Annual editions.

Government of India, Ministry of Finance. *Economic Survey 2025–26*.

Levitsky, S., and Way, L.A. (2010). *Competitive Authoritarianism: Hybrid Regimes After the Cold War*. Cambridge University Press.

Grier, K., and Grier, R. (2026). *Promises, Promises: Modi's India in Comparative Perspective*. SSRN Working Paper 7248338. Replication materials: github.com/rgrier88/modi-promises-replication.

Ministry of Statistics and Programme Implementation (2026). *Consumer Price Index on Base 2024=100: First Press Release and Frequently Asked Questions*. Government of India, February 2026.

Ministry of Statistics and Programme Implementation. *Periodic Labour Force Survey, Annual Reports and Monthly Bulletins, 2017–18 to 2026*.

Ministry of Statistics and Programme Implementation. *Household Consumption Expenditure Survey 2022–23 and 2023–24*.

Ministry of Statistics and Programme Implementation. *First Advance Estimates of National Income 2025–26*.

National Health Accounts, Ministry of Health and Family Welfare. *Estimates for India, 2014–15 to 2021–22*.

Penn World Table version 11.0. Groningen Growth and Development Centre.

PRS Legislative Research. *Union Budget analyses and cess/surcharge share of gross tax revenue*. Annual editions.

Reporters Without Borders. *World Press Freedom Index*. Annual editions, 2004–2026.

Reserve Bank of India. *State Finances: A Study of Budgets*, 2021 and 2023 editions; reference exchange rates.

Subramanian, A. (2019). *India's GDP Mis-estimation: Likelihood, Magnitudes, Mechanisms, and Implications*. Harvard Center for International Development Faculty Working Paper No. 354.

World Inequality Database (WID.world). India series, accessed 2026. *World Inequality Report 2026*.

Raman, V.S. (2026). *Some Perspective: India's Economic Transformation*. Dataset and replication code. someperspective.info; github.com/Varnasr/someperspective. CC BY 4.0.

Appendix A: Index specifications

A.1 Statistical Suppression Index

$$\mathrm{SSI}_t = \sum_{i=1}^6 w_i \cdot s_{it}, \quad s_{it} \in [0,1], \quad \sum_i w_i = 10$$

Weights: census delay 2.5; COVID mortality undercount 2.0; consumption-survey suppression 1.5; GDP back-series dispute 1.5; institutional independence 1.5; employment-data delay 1.0.

Severity rubric: ≈ 1.0 full suppression or discontinuation; ≈ 0.7 live methodology dispute; ≈ 0.3 delay or post-resolution scar.

A.2 Fiscal Centralisation Index

$$\mathrm{FCI}_t = \frac{1}{6} \sum_{j=1}^6 \tilde{C}_{jt}, \quad \tilde{C}_{jt} = \frac{C_{jt} - \min_t C_{jt}}{\max_t C_{jt} - \min_t C_{jt}}$$

Components: C_1 cess share; $C_2 = 1 -$ devolution share; $C_3 = 1 -$ states' own-revenue share; C_4 conditional-transfer share; C_5 borrowing restriction $\in \{0, 0.5, 1\}$; C_6 GST structural centralisation $\in [0,1]$.

Normalisation spans 2004–2026. The index is relative to this sample.

A.3 Democratic Quality Index

$$\mathrm{DQI}_t = \left(V_t \cdot \frac{F_t}{100} \cdot \frac{180 - R_t}{180} \cdot C_t \right)^{1/4}$$

V = V-Dem Liberal Democracy Index; F = Freedom House aggregate score; R = RSF press-freedom rank; C = V-Dem Core Civil Society Index.

Appendix B: Principal series, 2004–2026

Full machine-readable data: `data.json` ; flat export: `downloads/dataset.csv` .

YEAR	GDP GROWTH	TOP 1%	BOTTOM 50%	PRESS RANK	SSI	FCI	DQI
2004	7.9	18.2	16.5	120	0.00	0.00	0.59
2005	9.3	18.8	16.3	106	0.00	0.03	0.61
2006	9.3	19.6	16.0	105	0.00	0.05	0.61
2007	7.7	20.1	15.8	120	0.00	0.07	0.58
2008	3.1	20.5	15.6	118	0.00	0.09	0.58
2009	7.9	20.6	15.5	105	0.00	0.11	0.61
2010	8.5	21.0	15.4	122	0.00	0.12	0.56
2011	5.2	21.3	15.3	131	0.00	0.14	0.54
2012	5.5	21.0	15.2	140	0.00	0.16	0.51
2013	6.4	21.2	15.1	140	0.00	0.17	0.51
2014	7.4	21.3	15.0	140	0.00	0.19	0.49
2015	8.0	21.7	14.7	136	1.05	0.21	0.46
2016	8.2	21.5	14.9	133	1.05	0.25	0.46
2017	7.2	21.5	14.9	136	2.05	0.41	0.44
2018	6.1	21.7	14.7	138	4.00	0.53	0.42
2019	4.2	22.1	13.5	140	4.80	0.65	0.39
2020	−7.3	22.3	13.3	142	7.00	0.92	0.36
2021	8.7	22.5	13.1	142	9.00	0.82	0.37
2022	7.2	22.6	13.1	150	9.00	0.70	0.34
2023	7.6	22.6	13.1	161	8.20	0.68	0.28
2024	7.8	22.6	13.1	159	7.05	0.68	0.28
2025	6.5	22.8	13.0	151	6.55	0.68	0.30
2026	7.4	23.0	12.9	157	6.40	0.68	0.29

Appendix C: Replication

The complete dataset, construction script, and this paper's source are published at github.com/Varnasr/someperspective under CC BY 4.0.

- `data.json` — source of record for every figure in this paper
- `data/compute_indices.py` — deterministic construction of SSI, FCI and DQI for both eras
- `data/METHODOLOGY.md` — full provenance for every series, including anchors and interpolation rules for estimated series
- `replication_code.py`, `replication_code.R` — index reconstruction
- `downloads/dataset.csv`, `downloads/dataset.xlsx` — flat exports
- `tools/check_data_parity.py` — verifies the published dataset, the site's inline copy, and the construction script agree
- `tools/check_docs_consistency.py` — verifies year-labelled figures in published documents against the dataset
- `paper/paper.md` — source of this document

To regenerate the indices:

```
python3 data/compute_indices.py
```

To verify that the published data matches the construction:

```
python3 tools/check_data_parity.py
```

Data vintage: 15 August 2026. This paper is generated from `paper/paper.md` ; the HTML and PDF editions are built by `tools/build_paper.py` and regenerated whenever the source or the dataset changes.