Does Rule of Law Cause Growth? Evidence from Indian Judicial Appointments

Jun Goto* W. Bentley MacLeod† Sultan Mehmood‡

March 25, 2026

Abstract

Independent courts are widely thought to promote investment and economic growth, yet causal evidence remains scarce: judicial institutions rarely change for reasons unrelated to broader political and economic change. We exploit a remarkable episode in India in which the Supreme Court reinterpreted the constitution to wrest appointment authority over High Court judges away from the executive, vesting it in a collegium of senior jurists, without legislative action or constitutional amendment. What makes this episode tractable for causal inference is that the reform's bite depended entirely on the age structure of sitting judges at the moment of the ruling. High Court judges retire at sixty-two, so a court whose bench happened to be older in 1993 cycled rapidly into collegium-appointed judges; a court with younger incumbents retained pre-reform appointees far longer. Which courts had older benches at the moment of the ruling was predetermined by the accumulated accidents of past appointment timing, not by any post-reform political or economic condition. This cross-court variation in retirement-driven turnover is the source of our identification. Using it, we document three results. First, faster bench turnover sharply reduces government win rates in property-rights disputes, by roughly one quarter relative to the sample mean. Second, judicial independence caused broad-based economic development: districts exposed to faster judicial turnover show significantly more growth in labor market formalization, the physical stock of built capital, and visible economic activity, with consistent results across the Economic Census, satellite imagery, and nightlights. Third, selection appears to be a mechanism: collegium-appointed judges are more likely to have judicial-service backgrounds and rule less deferentially toward the state, while incumbent judges without these credentials move in the opposite direction. Thus, limiting executive influence over judicial appointments can strengthen rule of law and, in turn, support economic development.

*National Graduate Institute for Policy Studies (GRIPS). Email: ju-goto@grips.ac.jp.
†Department of Economics, Yale University; Columbia University. Email: bentley.macleod@yale.edu.
‡New Economic School, Moscow. Email: smehmood@nes.ru.

1 Introduction

The rule of law is widely considered a precondition for economic development. Cross-country evidence consistently associates secure property rights, impartial courts, and constraints on executive power with higher investment, faster structural transformation, and greater long-run income (North, 1990; La Porta et al., 2002; Djankov et al., 2003b; Besley and Persson, 2009; Acemoglu et al., 2015; Dincecco and Katz, 2016). Yet establishing that courts cause development, rather than simply reflecting the deeper political equilibria that also produce development, has proved elusive. Courts are notoriously difficult to randomize. Reforms to judicial institutions are rarely exogenous: they tend to arrive in the wake of political upheaval, constitutional crises, or economic shocks that independently alter the trajectories of the outcomes researchers care about. The most widely used measures of judicial independence are either too coarse, too slow-moving, or too entangled with broader governance to support clean identification.

This paper identifies the impact of improvements to the rule of law by exploiting a sharp, administratively accidental shock to judicial independence in India. On October 6, 1993, the Supreme Court issued its landmark *Second Judges Case* ruling, reinterpreting the constitutional provisions governing judicial appointments and transferring de facto control over High Court nominations from the executive to a collegium of the Court's most senior jurists. The ruling left substantive law, court jurisdiction, and economic policy entirely unchanged. What it altered was the process by which judges entered the bench and, therefore, over time, who those judges were. The central empirical challenge is that the reform was national and simultaneous. To identify its effect, we exploit the fact that the age distribution of sitting judges differed substantially across courts in 1993, combined with the fact that High Court judges retire at the mandatory age of sixty-two.

Courts whose incumbents happened to cluster near retirement turned over into collegium-appointed benches within a few years; courts with younger incumbents retained pre-reform appointees far longer. Which courts happened to have older benches at the moment of the ruling was not something any political actor could easily have anticipated or engineered in advance, depending as it does on the accumulated accidents of past appointment timing across many separate courts over many years. We formalize this through a retirement-turnover ratio: cumulative post-October-1993 retirements in each High Court divided by its sanctioned bench strength. The ratio is identically zero before the reform and rises mechanically thereafter at a pace set entirely by the 1993 age structure of each bench. It is our measure of how far a court has moved from the pre-reform appointment regime.

For the judicial analysis, we construct a corpus of over 82,000 state-involving High Court judgments in land and property-rights disputes from 1988 to 2010, harvested from published Indian High Court records and linked to a biographical database of judges covering their educational backgrounds, prior career histories, and appointment dates. For the economic analysis, we draw on three independent district-level panels. The Economic Census provides measures of formal employment, establishment size, and labor formalization across four census rounds. The Global Human Settlement Layer, a satellite product of the European Commission's Joint Research Centre, provides built-up surface area, average building height, and built volume at 100-meter resolution across ten five-year epochs from 1975 to 2020. DMSP-OLS stable nightlights provide annual district-level luminosity from 1992 to 2013.

The satellite measures are valuable in this setting. Building height is derived from digital surface models that record rooftop elevations net of underlying terrain, and is therefore constructed entirely independently of administrative records: it captures where physical capital is actually being built, not what governments report. The convergence of results across these three data families — each measuring a distinct dimension of growth through an independent methodology — strengthens the causal interpretation.

Armed with this variation, we ask three linked questions. The first is whether faster replacement of pre-reform benches under the collegium regime changed how courts decide state-involving land and property-rights disputes, specifically whether the government wins less often and whether judicial reasoning becomes less deferential to the state. The second, and most consequential, is whether those changes in adjudication translated into measurable downstream economic development at the district level, in employment, physical capital accumulation, and visible economic activity.

The evidence supports a coherent chain from judicial appointments to adjudication and then to economic outcomes. First, courts that transition more rapidly into the collegium regime become less likely to deliver outcomes favorable to the government in state-involving land disputes with the citizenry. A one-standard-deviation increase in the retirement-turnover ratio is associated with a twelve-percentage-point decline in government win rates, a reduction of roughly one quarter relative to the sample mean of 49 percent. This decline in government wins is accompanied by more pro-individual-rights reasoning and more merits-oriented adjudication.

Second, bench turnover changes who sits on the court. Collegium-era appointees are more likely to have elite legal training, and the composition shift helps explain the aggregate change in court behavior.

Third, at the district level, faster retirement-driven transition under the parent High Court is associated with stronger economic responses in places with larger pre-reform backlog exposure. The mechanism connecting adjudication to investment is straightforward. When courts rule less favorably to the state in land and property disputes, private claimants face lower expropriation risk, and the expected return to investing in land and fixed capital rises (Besley, 1995; Acemoglu and Johnson, 2005).

Secure property rights encourage owners to invest, formalize economic activity, and redeploy land toward higher-value uses, effects documented empirically in a range of developing-country settings (Field, 2007; Galiani and Schargrodsky, 2010). In our context, where the pre-reform backlog of unresolved state-involving land disputes captures precisely where this expropriation risk was most salient, the economic predictions are precise. Consistent with this, formal employment rises significantly in more exposed districts, as do average establishment size and hired labor share, pointing toward expansion and formalization rather than a broad rise in informal activity.

The satellite-based outcomes reinforce this picture: built volume, built height, and total nightlight luminosity all increase significantly in more exposed districts. The building-height results are especially informative because height is estimated from satellite radar elevation data entirely independently of any administrative record, providing a clean and manipulation-resistant measure of vertical capital accumulation. In an alternative district-level design based on physical buildability and underbuilt slack, the strongest evidence remains concentrated in built outcomes and local visible activity. That the Economic Census,

the GHSL built-environment panel, and the nightlights series all point in the same direction, despite measuring development through entirely different methodologies and at different frequencies, substantially strengthens the causal interpretation. Our building-height measure uses satellite elevation data to capture vertical construction. Digital surface models record the reflected elevation of rooftops and other above-ground structures, while digital terrain models recover the underlying ground surface. The difference between the two gives building height. Because this measure is derived independently of administrative data, it provides a useful proxy for local development and a clean way to validate other built-environment measures, in the same spirit as work that uses remotely sensed data to measure economic activity (Henderson et al., 2012).

Our evidence points to selection as a mechanism. Under the collegium, peer-appointed judges are more likely to come from elite law schools and judicial-service backgrounds, and these judges rule less deferentially toward the state. By contrast, incumbent judges without these backgrounds move in the opposite direction after the reform. Our findings suggest that limiting executive influence over judicial appointments can strengthen rule of law in practice and, in turn, support economic development.

Contributions to the literature. This paper contributes to several strands of literature. A large body of literature in political economy and development economics associates the rule of law, secure property rights, and impartial courts with higher investment and greater long-run income (North, 1990; La Porta et al., 1998; Djankov et al., 2003b; La Porta et al., 2004, 2008; Besley and Persson, 2009; Acemoglu et al., 2015). Work in this tradition has shown that legal origin shapes investor protection and financial development (La Porta et al., 1998, 2008), that the procedural formalism of courts predicts worse economic outcomes across countries (Djankov et al., 2003b), and that the form of legal institutions reflects how societies resolve the fundamental tradeoff between the costs of disorder and the costs of state power (Djankov et al., 2003a). This paper contributes to that tradition by providing causal evidence: exploiting a discrete and plausibly exogenous shift in appointment authority, we show that reducing executive influence over judicial selection strengthens judicial independence in practice and produces measurable downstream economic development. Closer to our setting, Chemin (2009) and Visaria (2009) show that judicial speed and legal reform affect credit markets and firm contracting in India, and Ponticelli and Alencar (2016) shows that court enforcement shapes firm investment in Brazil; we complement this work by studying how the independence rather than the efficiency of courts affects adjudication, and by tracing development effects specifically on the appointments channel.

Second, the paper contributes to the literature on state capacity and public sector selection. Besley and Persson (2009) show that legal and fiscal capacity are complementary investments that constrain what states can do, and that building legal capacity — the ability to enforce property rights and support markets — is central to long-run development. A related literature emphasizes that the quality of state institutions depends not only on formal rules but on who fills them: the characteristics, incentives, and selection of public officials shape what institutions actually deliver (Finan et al., 2017; Dal Bó et al., 2013; Choi et al., 2010; Ash and MacLeod, 2021; Moreira and Pérez, 2024). Courts occupy a particularly consequential position in this framework. They are the institution through which legal capacity

is exercised in practice: a state may formally protect property rights, but whether those rights constrain executive action depends on whether the judges hearing disputes between citizens and the state are willing and able to rule on the merits. We show that shifting the appointment authority from the executive to the judiciary changes who enters the bench in systematic ways, and that this compositional change is the primary channel through which the reform altered court behavior. The finding that selection rather than incentives drives the aggregate result speaks to broader debates about how institutional reforms propagate through organizations whose personnel turn over gradually, and connects the apparatus of state capacity to the specific question of how judicial independence is built in practice.

Third, the paper contributes to the literature on property rights and investment by showing that judicial independence matters most precisely where it should: in disputes and districts where state-citizen land conflict was most concentrated before the reform. A large microeconomic literature documents that tenure insecurity and expropriation risk suppress investment across a wide range of settings: from agricultural plots in Ghana where insecure land rights reduce soil fertility investment and output (Besley, 1995; Goldstein and Udry, 2008), to urban slums in Peru where titling raises labor supply and formalization (Field, 2007), to post-communist transition economies where weak property rights discourage firms from reinvesting profits even when credit is available (Johnson et al., 2002), to the American frontier where the technology of fencing enabled property enforcement and drove agricultural development (Hornbeck, 2010). At a more aggregate level, Acemoglu and Johnson (2005) show that it is specifically the property rights dimension of institutions, rather than contracting institutions, that drives long-run income differences across countries, and Galiani and Schargrodsky (2010) provide experimental evidence that land titling raises investment and welfare among the poor. Our contribution to this literature is to show that the judicial enforcement of property rights against the state itself is the binding constraint in settings where government action is the primary source of land insecurity. When courts rule less favorably to the state in land and property disputes, private claimants face lower expropriation risk and the expected return to fixed investment rises. That the economic effects are largest in formal employment, built capital, and vertical construction in the most backlog-exposed districts is consistent with this mechanism and difficult to reconcile with alternative accounts.

The remainder of the paper is organized as follows. Section 2 describes the institutional setting and the 1993 appointment reform. Section 3 describes the data. Section 4 studies judicial consequences of the reform. Section 5 studies economic consequences under the main backlog-based design and the alternative buildability-based design. Section 6 concludes.

2 Institutional Background

2.1 Higher Judiciary and the Appointment System

India's higher judiciary has a two-tier structure. At the apex stands the Supreme Court of India, which hears constitutional questions and appeals from the High Courts. Below it are the State High Courts, which serve as courts of record with original and appellate jurisdiction over civil and criminal matters within each state. High Courts also exercise writ

jurisdiction under Article 226 of the Constitution, which allows citizens and firms to directly challenge actions of the state, including administrative orders and land-acquisition actions. Because many economically salient disputes involve government action, High Courts often act as the first judicial forum in state-citizen conflicts over land and property rights.

Judges of the Supreme Court and the High Courts are formally appointed by the President of India under Articles 124 and 217 of the Constitution. In practice, however, the locus of appointment power has shifted over time. Prior to the early 1990s, appointments were widely understood to be dominated by the executive branch, with substantial influence exercised by the central government and state political elites over both nominations and promotions (Chandra et al., 2023). The higher judiciary retained security of tenure and constitutionally protected salaries, but the process by which judges were selected and elevated was heavily shaped by political actors.

A distinctive feature of the Indian system is the presence of mandatory retirement ages, set at 65 for Supreme Court judges and 62 for High Court judges. These fixed retirement ages imply predictable judicial exits and generate a steady flow of vacancies across courts. After the appointment regime changed in 1993, these age-based exits caused bench composition to shift gradually over time as pre-reform judges retired and were replaced by new appointees. Because the age structure of sitting judges varied widely across High Courts at the time of the reform, the speed of transition into the post-1993 appointment regime differed substantially across courts.

2.2 The 1993 Reform and the Collegium System

The central institutional turning point for this paper is the Supreme Court's 1993 decision in the *Second Judges Case* (*Supreme Court Advocates-on-Record Association v. Union of India*). In this landmark reinterpretation, the Court transformed the meaning of "consultation" in judicial appointments and held that primacy in appointments lay with the judiciary rather than the executive (Chandra et al., 2023). In operational terms, the ruling created the collegium system, under which the Chief Justice of India, together with a small group of the Court's most senior judges, collectively recommends candidates for appointment and elevation. Under this system, the executive formally retains the power to appoint, but its discretion is sharply constrained: recommendations can be returned for reconsideration once, but if reiterated they are conventionally binding.

In practice, the collegium's influence is exercised through a sequence of internal consultations and written recommendations that the executive processes through the Ministry of Law and Justice before the President issues the formal warrant of appointment. For High Court appointments, proposals originate from each High Court and the respective state government, while the collegium retains the final say on which candidates are forwarded for appointment. The office of the Chief Justice of India is filled through the same constitutional channel, but by a strong convention of seniority, which sharply limits discretionary political selection at the apex.

Reducing executive discretion over appointments changes the incentives faced by judges. When political actors have greater influence over appointments, transfers, and promotions, judges may have stronger career-related reasons to avoid decisions that impose costs on the government, especially in cases where the state is itself a litigant. By shifting de facto

appointment power toward senior judges, the collegium system plausibly increases judicial independence by weakening these career concerns. In our setting, where land acquisition, administrative action, and state involvement in property disputes are common, greater independence should be reflected in lower deference to the state and greater emphasis on merits and evidence in written opinions.

The lower judiciary is governed differently. District judges are appointed by the Governor in consultation with the High Court, and High Courts exercise administrative control over the subordinate judiciary. As a result, district judges are not directly appointed through the collegium system, but High Court leadership still plays a central role in the governance of the lower courts. The lower judiciary therefore occupies a different institutional channel; the empirical analysis here focuses on High Court decisions and district outcomes linked to High Court jurisdiction.

2.3 Land and Property Disputes in the Judicial Hierarchy

Many land and property disputes begin in the district or subordinate civil courts. However, when a dispute involves government action, such as a challenge to land acquisition, compensation, administrative title decisions, or other state action affecting land use, litigants can file directly in the High Court as a writ petition under Article 226. In such cases, the High Court acts as the first judicial forum.

High Court matters can be heard by a single judge or by a division bench, depending on subject-matter allocation and internal rules. The Chief Justice of each High Court allocates cases through the roster. In practice, many writ petitions and routine land and property matters are decided by single judges, while intra-court appeals and certain categories of cases are heard by division benches.

High Court decisions reached in writ jurisdiction may be subject to intra-court appeals from a single judge to a division bench where such appeals are maintainable. Separately, cases can reach the High Court through statutory appeals or revisions from subordinate courts and specialized authorities. In land-acquisition compensation disputes, for example, statutory provisions provide High Court appellate review. Our litigation dataset consists of High Court judgments, so it includes both original writ decisions and appellate or revisional decisions to the extent that they appear as High Court judgments in the corpus. Each observation in the judicial analysis is therefore a High Court judgment rather than a full dispute-level case history.

3 Data

3.1 Judicial Outcomes Sample

The judicial analysis uses a case-level sample of state-involving land and property-rights judgments from the High Courts of India over 1988–2010. The unit of observation is a High Court judgment. The corpus is harvested from the published High Court judgments available through Indian Kanoon and then merged with the judge-biography and court-treatment files used elsewhere in the paper. The sample restriction is mechanical. First, petitioner and

respondent names are parsed from the case title, and a case is classified as state-involving when one side matches a pre-specified state-entity dictionary (including *State of ...*, *Union of India*, ministries, departments, development authorities, municipal corporations, boards, commissions, and other public bodies). Second, the case is retained only if the judgment contains pre-specified land/property-rights conflict keywords or cites one of the pre-specified property-rights statutes. This is the domain in which judicial independence should be especially consequential for property-rights credibility.

The main court-level outcome is *Government Wins*, a binary indicator equal to one when the government fully or partly wins the case and zero when the government loses. The analysis also studies text-based and case-quality outcomes, including measures of pro-rights reasoning, merits-oriented adjudication, and procedural dispositions. Appendix Table 2A and Appendix A.1 define these variables in detail.

3.2 District-Level Economic Panels

The economic analysis combines three independent district-level outcome panels. First, the Economic Census, harmonized through SHRUG, provides district-level measures of formal employment, total employment, informal employment, average establishment size, hired labor share, and formal establishment share over four rounds: 1990, 1998, 2005, and 2013. Second, the Global Human Settlement Layer (GHSL R2023A) provides district-level built-environment measures over ten five-year epochs from 1975 to 2020. We use three GHSL quantities: built surface, defined as the aggregate number of square meters classified as built-up within each district; built height, defined as average gross building height; and built volume, defined mechanically as built surface times height. These measures are first estimated on a global 100-meter grid and then aggregated to district boundaries. Third, DMSP-OLS stable nightlights provide annual district-level measures of total and mean luminosity over 1992–2013. Total light captures overall visible economic activity, while mean light captures average light intensity and is therefore less mechanically tied to district size. Appendix Table 5A reports descriptive statistics for these outcome families, while Appendix Table 6A explains why the number of observations, districts, and High Courts differs across the Economic Census, GHSL, and nightlights panels.

These outcome families play different roles in the empirical interpretation. The Economic Census outcomes speak to employment composition and formalization. GHSL measures the stock of physical structures and therefore speaks more directly to construction, land redevelopment, and the spatial form of capital accumulation. Built surface captures horizontal spread, built height captures vertical intensity, and built volume combines the two into a single measure of total built mass. Nightlights are conceptually different again: they proxy visible local economic activity rather than the stock of structures itself. The evidence is therefore strongest when these three outcome families converge. This distinction is especially important because some districts can add built stock mainly through low-rise outward expansion, while others can add it through vertical intensification. Likewise, visible economic activity can rise even when the physical built footprint changes only gradually.

All district-level outcomes are harmonized to a common district geography and linked to the parent High Court that governs each district. Because the judicial treatment is defined at the High Court level, the district panels inherit treatment exposure through the

court that has jurisdiction over the district. Appendix B.1 reports the outcome distributions and sample-accounting facts for these panels. Appendix B.2 then provides a self-contained description of the GHSL and nightlights measures, including the satellite inputs, the construction of built surface and height, the aggregation from gridded products to district-level outcomes, and the interpretation of these measures in the paper's setting. That appendix also reproduces the technical process figures and district-level examples used to validate the meaning of the GHSL built-surface, built-height, and built-volume variables.

4 Effect of Collegium System on Rule of Law

4.1 Treatment Measures and Descriptive Evidence

The main treatment variable is the *retirement-turnover ratio*. For each High Court and year, this variable equals the cumulative number of post-October-1993 retirements divided by a fixed court-specific sanctioned bench strength. It measures how far a court has moved from the pre-reform appointment regime toward a bench composed of judges appointed under the collegium system. Because the age structure of sitting judges differed sharply across High Courts in 1993, the retirement-turnover ratio rises faster in some courts than in others even though the institutional reform was national. Pre-1993 retirements are not counted in the numerator because the treatment is not generic judicial turnover; it is turnover into the post-1993 judge-dominant appointments regime. Values above one therefore indicate that a court has replaced more than one full sanctioned bench since the reform, and the baseline specification treats this cumulative replacement intensity linearly.

We also estimate specifications under three alternative treatment definitions. The first is an *Above-Median versus Below-Median retirement-exposure* binary that classifies a court as Above Median if it reaches one full bench turnover by 2000. The second is a *threshold-crossing* binary that turns on in the first year in which the retirement-turnover ratio reaches one. The third is an *accumulated new appointments ratio*, which counts realized post-reform appointments relative to sanctioned strength. The appointments-based measure is informative but not preferred because realized appointments can also reflect delays in final executive action after collegium nomination. The paper treats the 1993 *Second Judges Case* as the institutional break because that is when executive primacy in appointments was displaced; the 1998 *Third Judges Case* refined collegium procedures, but it did not restore the pre-1993 appointments regime or create a second reset in court exposure.

Figure 2 visualizes the two main court-level treatment measures. Panel A shows the retirement-turnover ratio across High Courts and years, revealing substantial cross-court variation in the pace of institutional replacement. Panel B shows the corresponding threshold-crossing binary treatment, which turns on when a court first reaches one full bench turnover.

Table 1 reports descriptive statistics for the judicial estimation sample. The government wins approximately 49% of state-involving cases under the main coding, and the median retirement-turnover ratio is 1.86. The sample also displays substantial dispersion in judgment length and citations, which underscores the importance of controlling for observable case characteristics in the judicial-outcome regressions.

4.2 Empirical Strategy

The court-level analysis estimates whether faster transition into the post-1993 appointment regime changes adjudication in state-involving High Court cases. The main specification is

$$Y_{ict} = \beta \,\text{Retirement-turnover ratio}_{ct} + \gamma_c + \delta_t + X'_{ict}\theta + \varepsilon_{ict},$$

where Y_{ict} is an outcome for judgment i in High Court c and year t, γ_c are High Court fixed effects, δ_t are year fixed effects, and X_{ict} includes the state-petitioner indicator and log judgment length. Standard errors are clustered at the High Court×year level.

The identifying idea is straightforward. Before the reform, the retirement-turnover ratio is zero because no collegium-era replacement has yet occurred. After the reform, the ratio rises mechanically as judges appointed under the old regime retire. Conditional on court and year fixed effects, the coefficient on the retirement-turnover ratio therefore captures whether courts that replace their pre-reform benches more rapidly become less government-favorable in state-involving disputes. Judge transfers, delayed replacements, and vacancies can blur the mapping from retirement-driven court exposure to the composition of the bench actually hearing a given case, but those forces add noise to the court-year treatment rather than creating the retirement patterns that identify the specification.

4.2.1 Alternative Treatment Designs

The paper uses three alternative treatment definitions to probe robustness and timing.

First, the Above-Median versus Below-Median retirement-exposure design compares courts that complete one full bench turnover by 2000 with those that do not:

$$\text{Above-Median retirement exposure}_c = 1\left\{\min\{t : \text{Retirement-turnover ratio}_{ct} \geq 1\} \leq 2000\right\}.$$

The corresponding specification interacts this indicator with the post-1993 period. This binary partition serves two roles: it defines the treatment groups in the event-study and binary-difference specifications, and it provides the partition for pre-reform balance checks. The main judicial regressions still use the continuous retirement-turnover ratio; the Above-Median versus Below-Median split is used only where a coarse treatment partition is easier to interpret.

Second, the threshold-crossing binary design defines treatment as turning on once the retirement-turnover ratio reaches one. A court is untreated before it reaches one full bench turnover and treated in every year afterward. The design is staggered because courts cross that threshold in different years. For example, one High Court might first reach a full bench turnover in 1998, another in 2001, and another in 2004. Treatment therefore begins in different calendar years across courts, which is why the Sun-Abraham estimator is appropriate in this design. Formally,

$$\text{Treatment}_{ct} = 1\{t \geq T_c^*\}, \qquad T_c^* = \min\{t : \text{Retirement-turnover ratio}_{ct} \geq 1\}.$$

This design supports staggered-adoption event-study estimates using the Sun-Abraham interaction-weighted estimator.

Third, the accumulated new appointments ratio captures realized institutional replacement, but because the timing of actual appointments can reflect administrative delay after

collegium recommendation, it provides a less clean measure of exposure than the retirement-based ratio.

4.2.2 Balance Tests

For this interpretation to be credible, faster eventual turnover cannot simply stand in for pre-existing differences across courts that would have predicted post-1993 adjudication even without the appointments reform. Two concerns are especially important: pre-reform case composition and pre-reform bench composition. The balance exercises therefore use pre-reform data only and are constructed specifically for those threats. The case-composition exercise aggregates the judicial sample to a High Court×year panel and asks whether Above-Median and Below-Median courts were already hearing systematically different kinds of state-involving land and property disputes before October 1993. The bench-composition exercise uses a separate High Court×month panel of sitting-judge characteristics and asks whether the observable composition of the bench already differed before the reform. Appendix Figure 1A reports these two exercises, and Appendix Table 4A reports the numeric judge-composition coefficients. No displayed judicial-balance coefficient remains significant at the 10 percent level after Holm correction within the relevant family of tests.

4.3 Government Wins

We begin with the most direct question: whether faster replacement of pre-reform benches under the collegium regime changes who wins state-involving land and property-rights disputes. Table 2 reports the core estimates. The coefficient on the retirement-turnover ratio is negative, stable across specifications, and precisely estimated. In the preferred specification, the estimate is -0.0636 with a standard error of 0.0187 ($p = 0.0007$; $N = 82{,}334$), implying that High Courts that transition more rapidly into the post-1993 appointment regime become less likely to deliver outcomes favorable to the government. Using the standard deviation from Table 1, a one-standard-deviation increase in the retirement-turnover ratio corresponds to roughly a 0.12 decline in the government-win rate, or about 12 percentage points relative to a sample mean just below 0.49.

The specification ladder shows that this pattern does not depend on a narrow modeling choice. The estimate remains negative as the design moves from year fixed effects only, to the addition of High Court fixed effects, and then to the full specification with case controls. This stability is important because it suggests that the result is not being driven by cross-court differences in baseline win rates or by simple changes in observable case composition. Instead, the decline in government wins appears tightly connected to the pace of institutional replacement itself.

The same conclusion emerges under the alternative treatment designs. The Above-Median versus Below-Median retirement-exposure specification yields a negative and statistically significant estimate of -0.0763, and the staggered-adoption Sun-Abraham estimator yields a negative ATT of -0.0417. Because these binary treatments do not trace the full path of institutional replacement, they carry less identifying power, but they point in the same direction. The accumulated new appointments ratio also yields negative estimates, but

with less precision, which is consistent with the view that realized appointments are noisier than retirements as a measure of exposure to the new regime.

Dynamic evidence also supports the main interpretation. Figure 3 reports event-study estimates for the Above-Median versus Below-Median retirement-exposure design. The pre-period coefficients are close to zero on average, while the post-period coefficients shift downward, consistent with the main specification. The retirement-turnover ratio remains the preferred treatment because it traces the full path of institutional replacement, but the event-study timing pattern shows that the decline in government wins is not being generated by a pre-existing trend.

Taken together, these results provide the paper's first key fact: judicial appointments matter for adjudication in a substantively meaningful way. High Courts that move more quickly into the collegium regime become less favorable to the government in the litigation domain where judicial independence should matter most. Appendix C reports additional robustness for the supplementary binary designs.

4.4 Judicial Reasoning and Case Quality

Changes in win rates are important, but they are not sufficient on their own to show that judicial behavior becomes more independent in a substantively meaningful sense. We therefore examine whether faster institutional replacement also changes how judges reason and how they dispose of state-involving cases. Table 3 groups the most informative outcomes into two families: pro-individual-rights reasoning and case-merit or procedural outcomes.

Panel A focuses on textual indicators of more rights-protective judicial reasoning. The first outcome, *scrutiny versus deference*, measures whether the opinion uses language associated with searching judicial review of state action rather than language associated with restraint and deference. The second, *individual versus systemic*, measures whether the opinion emphasizes the concrete claimant and individualized relief rather than generalized systemic remedies. The third is a narrower concluding-section contrast score that captures whether the operative portion of the judgment places greater weight on human rights, individual rights, scrutiny, and individualized relief than on deference, order, and systemic remedies. The final column averages the standardized versions of these measures into a single pro-individual-rights reasoning index. Appendix A.1 explains how these variables are constructed. The first two are court-relative phrase-axis measures built from fixed dictionaries and normalized to each court's pre-1988 reasoning baseline; the third is a conclusion-focused contrast score constructed from support and counter phrase sets in the operative portion of the judgment.

The broad pattern in Panel A is consistent with a court system that becomes less deferential to the state as bench replacement progresses. The scrutiny-versus-deference outcome rises significantly, and the summary index is also positive and precisely estimated. The individualized-relief measure and the exploratory concluding-section contrast score are less precise on their own, but they move in the same qualitative direction. The picture is therefore not one of a single textual measure changing in isolation, but of a broader shift in reasoning away from administrative deference and toward closer attention to the individual claimant.

Panel B turns from reasoning to adjudicatory style. The first outcome counts language that explicitly weighs the evidentiary basis of the case. The second counts language rejecting unsupported government claims. The third captures procedural closure—phrases indicating

disposal on procedural grounds rather than on the merits. The final column combines these pieces into a merits-oriented adjudication index. Here the clearest evidence is that faster turnover increases the rejection of unsupported government claims and reduces procedural closure. The broader merits-oriented index is positive but less precisely estimated, which suggests that the reform changes some concrete dimensions of adjudication more clearly than it changes the full bundle of quality-related outcomes all at once.

These results strengthen the interpretation of the government-win findings. Courts that turn over more rapidly do not merely produce a different mix of winners and losers. They also write opinions that are less deferential, more willing to scrutinize unsupported government claims, and less likely to terminate disputes on purely procedural grounds. That pattern is exactly what one would expect if the collegium-era appointment regime increased the judiciary's willingness to constrain the state on the merits.

4.5 Mechanisms: Selection Around the Reform

The court-level results show that faster institutional replacement reduces government wins and shifts reasoning away from deference. The most direct mechanism evidence here concerns *selection*: whether the collegium changed who entered the High Courts.

Specification. The unit of observation is a judge, anchored at the first High Court appointment. For each characteristic outcome Y_{ic} for judge i appointed to High Court c, we estimate:

$$Y_{ic} = \alpha + \beta \cdot \text{CollegiumAppointed}_i + \gamma_c + \varepsilon_{ic},$$

where $\text{CollegiumAppointed}_i$ equals one if the judge's first High Court appointment falls on or after October 6, 1993, and γ_c are High Court fixed effects. Standard errors are clustered at the High Court level. Figure 4 reports the dynamic version of this comparison using four-year appointment-date bins, which also allows assessment of pre-reform trends.

Results. Figure 4 focuses on two attributes that bear most directly on the judicial-independence mechanism: elite law school attendance (Panel A) and prior judicial-service experience (Panel B). Each panel plots binned coefficients with High Court fixed effects and 95 percent confidence intervals, with dashed lines marking the pre-reform and post-reform unconditional means.

In both panels, the pre-reform bins are close to zero and show no discernible trend, while the post-reform bins shift upward. The pattern is clearest for judicial-service background (Panel B), where the post-reform mean rises by roughly seven percentage points relative to the pre-reform mean on a base of approximately 35 percent. Elite law school attendance (Panel A) shows a similar directional shift, with the post-reform mean approximately tripling relative to a low pre-reform base, though individual bins are less precisely estimated. These within-High-Court comparisons indicate that the collegium appointed systematically different judges, not merely that appointments happened to coincide with a secular trend in the legal profession.

These two attributes matter for the mechanism documented in Sections 4.3 and 4.4. Judges with judicial-service backgrounds enter the High Court with direct experience adjudicating disputes between citizens and the state in the subordinate courts, which plausibly

reduces deference to the government as a litigant. Elite legal training may similarly foster confidence in independent legal reasoning. That the collegium drew more heavily from these backgrounds is therefore consistent with the decline in government win rates and the shift toward more rights-protective reasoning as bench replacement progresses. A fuller sorting-versus-incentive decomposition would require modelling judicial career concerns and is beyond the scope of this paper.

5 Economic Consequences of Judicial Independence

5.1 District Exposure Measures

The preferred district-level heterogeneity measure is *Backlog Exposure*. For each district, the paper counts the number of state-involving property-rights disputes from that district that were pending in the parent High Court immediately before October 6, 1993. This count is then standardized within each parent High Court to form a within-court z-score. Intuitively, this measure captures where unresolved state-citizen property conflict was already concentrated at the reform date. Districts with larger pre-reform backlogs are the places where subsequent judicial independence should matter most if the reform changes how High Courts handle disputes involving the state. Appendix Table 6A reports the distribution of the raw backlog count, the within-parent-High-Court backlog share, and the standardized exposure used in the main economic design.

The main economic design is supplemented by an alternative district-level exposure based on physical development potential rather than litigation history. This alternative exposure is designed to ask whether the same retirement-driven judicial shift is associated with stronger built-environment responses in districts that were physically easier to develop and still underbuilt at baseline.

The alternative measure, *Buildability Exposure*, combines two predetermined components. The first is *physical buildability*, a geophysical index that averages standardized terrain and hydrology constraints so that higher values denote flatter, less rugged, less water-constrained districts. In practice, the index uses slope, low-slope share, ruggedness, elevation dispersion, elevation range, and surface-water constraints. The second is *underbuilt slack*, which captures whether a district had less built stock in 1990 than would be expected given that physical feasibility. In the preferred version, underbuilt slack is the negative standardized residual from a cross-sectional regression of 1990 built volume on the physical buildability index alone. The main alternative exposure is the standardized product of these two objects. A stricter appendix variant replaces the slack term with the negative standardized residual from a 1990 built-volume regression that conditions on both physical buildability and accessibility. The alternative design is therefore meant to capture places that were physically feasible but not yet fully built out at the start of the reform era.

5.2 Empirical Strategy

The preferred district-level economic specification interacts the court-level retirement-turnover ratio with pre-reform backlog exposure:

$$Y_{dt} = \delta\Big(\text{Retirement-turnover ratio}_{h(d)t} \times \text{Backlog exposure}_d\Big) + \alpha_d + \lambda_t + \text{constituent terms} + \varepsilon_{dt},$$

where Y_{dt} is a district-level outcome in district d and period t, $h(d)$ denotes the parent High Court, α_d are district fixed effects, and λ_t are period fixed effects at the relevant frequency. The specification also includes the court-level treatment, the district exposure term, and the constituent two-way interaction terms needed for a saturated triple-difference design. The coefficient of interest is therefore identified from differential post-reform changes associated with the interaction between court-level institutional replacement and predetermined district exposure.

Because the retirement-turnover ratio is identically zero before October 1993, the court-level treatment term in this equation is numerically equivalent to a post-reform treatment interaction. In the implementation, the lower-order terms are the court-level treatment itself and the district-exposure-by-period term that accompanies the main interaction. The coefficient can therefore be read as the additional district response associated with moving one within-parent-High-Court standard deviation higher in pre-reform backlog exposure for a given increment in retirement-driven court turnover.

This design does not attempt to identify the effect of any single judge assignment on any single district. Instead, it measures whether the reform had larger downstream economic effects in districts where unresolved state-involving property-rights disputes were already more concentrated when the collegium regime began. The interpretation is therefore one of differential district exposure to a court-level institutional change.

Backlog Exposure is a predetermined heterogeneity term, not a claim of district-level random assignment. The main economic design therefore asks whether the same court-level institutional change had larger post-1993 effects in districts that entered the reform era with more unresolved state-citizen property conflict. The buildability design then asks whether a structurally different, non-litigation-based district exposure term points in the same direction.

5.2.1 Alternative Exposure Design

The paper also reports an alternative district-level economic design based on predetermined physical development potential. This design combines the two underlying ingredients into a single exposure term:

$$\text{Buildability exposure}_d = \text{PhysicalBuildability}_d \times \text{UnderbuiltSlack}_d.$$

The associated estimating equation is

$$Y_{dt} = \kappa\Big(\text{Retirement-turnover ratio}_{h(d)t} \times \text{Buildability exposure}_d\Big) + \alpha_d + \lambda_t + \text{constituent terms} + \varepsilon_{dt}.$$

The buildability-based exposure is geographic and predetermined rather than litigation-based, so its identification assumptions are distinct. It is also narrower: the strongest effects appear in built-environment and visible-activity outcomes rather than across the full set of economic measures.

5.2.2 Balance Tests

A threat to internal validity is that High Courts with faster eventual turnover might systematically serve jurisdictions that were already larger, more urban, more developed, or more conflict-prone before the reform. If so, the downstream economic heterogeneity could partly reflect baseline regional differences rather than the interaction between court-level institutional change and pre-reform district exposure. The paper therefore reports a separate pre-reform regional and economic balance exercise. Unlike the main economic regressions, which use district-level outcome panels and the continuous retirement-turnover ratio, this balance exercise uses a court-jurisdiction cross-section because the object of interest is whether High Courts above and below the median of eventual retirement exposure already differed in baseline regional conditions before the reform. Appendix Figure 2A compares Above-Median and Below-Median retirement-exposure courts on court-jurisdiction averages of log population, capped log urban population, log rural population, log total employment, backlog share, district mention share, and log district case count. No displayed coefficient remains significant at the 10 percent level after Holm correction within this family of tests.

Having established that the reform changed adjudication, we now test whether those changes produced downstream economic effects in the districts served by different High Courts. The preferred district-level design uses pre-reform backlog exposure, so the interpretation is explicitly tied to places where unresolved state-involving property-rights disputes were already concentrated before the reform.

5.3 Descriptive Evidence

Figure 5 provides a visual summary of the district-level variation used in the main economic design. Panel (a) maps treatment intensity, defined as the product of the parent High Court's retirement-turnover ratio and the district's pre-reform backlog exposure. Panel (b) maps built-height growth from 1990 to 2020. The raw correlation is consistent with the main hypothesis: districts with higher treatment intensity tend to exhibit above-average growth in built height, while the least-exposed districts show below-average physical-capital accumulation.

The economic-outcome measures deserve careful explanation because they capture different margins of adjustment. Built surface is the district-wide built footprint aggregated from GHSL's 100-meter built-up grid. Built height is average gross building height, estimated by GHSL from satellite-derived elevation data and then aggregated to the district. Built volume combines the two and therefore captures the total mass of built structures rather than only land coverage. This distinction matters in our setting because a district can add physical capital either by expanding outward in built surface or by intensifying vertically in height. The nightlights measures are conceptually different: they proxy visible economic activity rather than the physical stock of structures. Total light captures overall luminous activity,

while mean light captures average light intensity and is therefore less mechanically tied to district area. Taken together, the built-environment outcomes and nightlights allow us to distinguish changes in the stock of physical capital from changes in broader local economic intensity.

The timing structure also differs across data sources. The Economic Census is observed in four discrete rounds. GHSL is observed in five-year epochs from 1975 to 2020 and is therefore best interpreted as a long-run built-stock panel rather than as a direct construction-flow dataset. DMSP-OLS is observed annually from 1992 to 2013. The GHSL stock outcomes — built surface, built height, and built volume — are the primary built-environment measures. The epoch-to-epoch differences ("new built surface" and "new built volume") are flow proxies derived from those stocks, not administrative construction records, and should be interpreted accordingly. All three outcome families illuminate different dimensions of economic adjustment to the same court-level institutional change. Appendix B.2 provides a full description of the satellite measurement process, shows the underlying GHSL construction pipeline, and explains why built height is noisier than built surface and why the paper emphasizes the joint pattern across built surface, built height, built volume, and nightlights rather than any single satellite measure in isolation.

Appendix Table 6A also makes clear why the district counts differ across economic tables. The raw Economic Census panel covers 640 districts, whereas the GHSL and nightlights panels cover 641. The backlog design is defined only where districts can be linked to the pre-reform backlog inventory. The buildability designs impose different support requirements because they additionally require geophysical inputs, 1990 built-volume support, and, in the stricter appendix design, non-missing accessibility. The number of High Courts falls further in the buildability tables because those specifications remap later court splits to their pre-reform parent High Courts.

5.4 Main Estimates

Table 4 reports the main triple-difference estimates. The coefficient of interest is the interaction between the retirement-turnover ratio and district backlog exposure. Across outcome families, the estimates imply that judicial independence matters most in districts where unresolved state-involving property-rights conflict was most concentrated at baseline.

The results are strongest in two parts of the economy. First, the Economic Census outcomes show significant increases in formal employment, total employment, average establishment size, hired labor share, and formal establishment share. This pattern points toward expansion and formalization rather than a broad rise in informal activity. Second, the satellite-based outcomes show significant increases in built volume, built height, and total nightlight luminosity, while the built-surface coefficient is positive but imprecisely estimated. These results are consistent with a district-level environment in which improved judicial independence facilitates physical investment, construction, and visible economic activity in precisely the places where pre-reform state-related property disputes were most congested. Read in conjunction with the outcome definitions above, the positive built-volume and built-height coefficients imply that the reform is associated not merely with spatial expansion of development, but with denser and taller built capital in the districts most exposed to unresolved property-rights backlog.

5.5 Alternative Economic Identification

The main economic design is intentionally tied to litigation history through backlog exposure. A distinct question is whether a physically grounded district exposure, which does not rely on the pre-reform litigation distribution, reaches similar conclusions. The alternative design answers that question by replacing backlog exposure with *Buildability Exposure*, a measure that combines physical buildability with underbuilt slack at baseline.

This design answers a narrower question: whether judicial independence matters more in districts that were physically feasible for development but still underbuilt at baseline. Because it is based on geography rather than litigation history, it is more clearly predetermined. Because it is tied specifically to the location of construction and visible activity, it is also narrower in scope.

Table 5 reports the alternative estimates only for the outcome families that are conceptually closest to this design: built outcomes and nightlights. The preferred buildability-exposure measure uses physical buildability multiplied by conditional-volume underbuilt slack. Under this design, built volume, built height, built surface, and mean nightlights all remain positive, and the built-stock outcomes are precisely estimated. Built volume rises by 0.0236, built height by 0.0213, and built surface by 0.0223. Mean light also rises by 0.0185. Appendix Table 7A shows that the same qualitative pattern survives when the exposure term is rebuilt using the stricter accessibility-adjusted residual-slack measure.

The interpretation is therefore complementary to the main backlog-based results. The preferred economic design shows that judicial independence matters most where unresolved property conflict was already concentrated. The alternative design shows that a related positive pattern also appears in places that were physically developable but still underbuilt. Taken together, the two designs suggest that the reform affected both dispute resolution in conflict-intensive districts and the broader physical environment for investment where land could plausibly be redeployed.

6 Conclusion

This paper asks whether reducing executive influence over judicial appointments strengthens the rule of law in practice and, if so, whether that change translates into economic development. The 1993 *Second Judges Case* created a discrete institutional shock to appointment authority in India, but the practical pace of transition differed sharply across High Courts because judges retire at fixed ages. Exploiting this retirement-driven heterogeneity, the paper traces a chain from judicial selection to judicial behavior and then to district-level economic outcomes.

The first set of findings concerns adjudication. High Courts that transition more rapidly into the collegium regime become less likely to deliver outcomes favorable to the government in state-involving land and property-rights disputes, with a one-standard-deviation increase in bench replacement associated with a twelve-percentage-point decline in government win rates relative to a sample mean of 49 percent. These same courts also exhibit more rights-protective and less deferential reasoning, together with a greater willingness to reject unsupported government claims and a reduced reliance on procedural closure. The reform

therefore changes not only who wins, but how courts justify their decisions.

The second set of findings concerns downstream economic consequences. When courts constrain the state in land and property disputes, private claimants face lower expropriation risk and the expected return to fixed investment rises. Consistent with this, faster retirement-driven court transition is associated with stronger formal-sector expansion, larger built-environment gains, and higher nightlight growth in districts with larger pre-reform backlogs of unresolved state-involving property-rights disputes. An alternative design based on physical buildability and underbuilt slack reaches a narrower but related conclusion. Together, the two economic designs support the view that judicial independence mattered most where property conflict and land reallocation were economically salient, and that the rule of law has real consequences for development precisely in the settings where the state is the primary source of insecurity.

The third set of findings concerns mechanism. The reform changed the composition of appointees, particularly by increasing the share of judges with judicial-service backgrounds and elite legal training. The mechanism evidence therefore points to selection: judicial independence, in this setting, is built through who enters the bench under the new appointments regime.

More broadly, the paper shows that judicial independence can be studied through appointment rules rather than through broad constitutional indices alone, that who enters the bench shapes what courts deliver, and that reforms to judicial selection can have measurable downstream consequences for economic development. The causal chain runs from appointment rules, through the composition and behavior of courts, through adjudication of state-involving land disputes, to investment and visible economic activity on the ground. Strengthening the rule of law, this evidence suggests, need not require sweeping constitutional change: it can begin with who is allowed to choose the judges.

References

Acemoglu, D., García-Jimeno, C., and Robinson, J. A. (2015). State capacity and economic development: A network approach. *American Economic Review*, 105(8):2364–2409.

Acemoglu, D. and Johnson, S. (2005). Unbundling institutions. *Journal of Political Economy*, 113(5):949–995.

Ash, E. and MacLeod, W. B. (2021). Reducing partisanship in judicial elections can improve judge quality: Evidence from u.s. state supreme courts. *Journal of Public Economics*, 201:104478.

Besley, T. (1995). Property rights and investment incentives: Theory and evidence from Ghana. *Journal of Political Economy*, 103(5):903–937.

Besley, T. and Persson, T. (2009). The origins of state capacity: Property rights, taxation, and politics. *American Economic Review*, 99(4):1218–1244.

Chandra, A., Kalantry, S., and Hubbard, W. H. J. (2023). *Court on Trial: A Data-Driven Account of the Supreme Court of India*. Penguin Random House India.

Chemin, M. (2009). Do judiciaries matter for development? Evidence from India. *Journal of Comparative Economics*, 37(2):230–250.

Choi, S. J., Gulati, G. M., and Posner, E. A. (2010). Professionals or politicians: The uncertain empirical case for an elected rather than appointed judiciary. *Journal of Law, Economics, and Organization*, 26(2):290–336.

Dal Bó, E., Finan, F., and Rossi, M. A. (2013). Strengthening state capabilities: The role of financial incentives in the call to public service. *Quarterly Journal of Economics*, 128(3):1169–1218.

Dincecco, M. and Katz, G. (2016). State capacity and long-run economic performance. *Economic Journal*, 126(590):189–218.

Djankov, S., Glaeser, E., La Porta, R., Lopez-de Silanes, F., and Shleifer, A. (2003a). The new comparative economics. *Journal of Comparative Economics*, 31(4):595–619.

Djankov, S., La Porta, R., Lopez-de Silanes, F., and Shleifer, A. (2003b). Courts. *Quarterly Journal of Economics*, 118(2):453–517.

Field, E. (2007). Entitled to work: Urban property rights and labor supply in Peru. *Quarterly Journal of Economics*, 122(4):1561–1602.

Finan, F., Olken, B. A., and Pande, R. (2017). The personnel economics of the developing state. In Banerjee, A. V. and Duflo, E., editors, *Handbook of Economic Field Experiments*, volume 2, pages 467–514. North-Holland.

Galiani, S. and Schargrodsky, E. (2010). Property rights for the poor: Effects of land titling. *Journal of Public Economics*, 94(9–10):700–729.

Goldstein, M. and Udry, C. (2008). The profits of power: Land rights and agricultural investment in Ghana. *Journal of Political Economy*, 116(6):981–1022.

Henderson, J. V., Storeygard, A., and Weil, D. N. (2012). Measuring economic growth from outer space. *American Economic Review*, 102(2):994–1028.

Hornbeck, R. (2010). Barbed wire: Property rights and agricultural development. *Quarterly Journal of Economics*, 125(2):767–810.

Johnson, S., McMillan, J., and Woodruff, C. (2002). Property rights and finance. *American Economic Review*, 92(5):1335–1356.

La Porta, R., Lopez-de Silanes, F., Pop-Eleches, C., and Shleifer, A. (2004). Judicial checks and balances. *Journal of Political Economy*, 112(2):445–470.

La Porta, R., Lopez-de Silanes, F., and Shleifer, A. (2008). The economic consequences of legal origins. *Journal of Economic Literature*, 46(2):285–332.

La Porta, R., Lopez-de Silanes, F., Shleifer, A., and Vishny, R. W. (1998). Law and finance. *Journal of Political Economy*, 106(6):1113–1155.

La Porta, R., Lopez-de Silanes, F., Shleifer, A., and Vishny, R. W. (2002). Investor protection and corporate valuation. *Journal of Finance*, 57(3):1147–1170.

Moreira, D. and Pérez, S. (2024). Civil service exams and organizational performance: Evidence from the pendleton act. *American Economic Journal: Applied Economics*, 16(3):250–291.

North, D. C. (1990). *Institutions, Institutional Change, and Economic Performance*. Cambridge University Press, New York.

Ponticelli, J. and Alencar, L. S. (2016). Court enforcement, bank loans, and firm investment: Evidence from a bankruptcy reform in Brazil. *Quarterly Journal of Economics*, 131(3):1365–1413.

Visaria, S. (2009). Legal reform and loan repayment: The microeconomic impact of debt recovery tribunals in India. *American Economic Journal: Applied Economics*, 1(3):59–81.

Figures

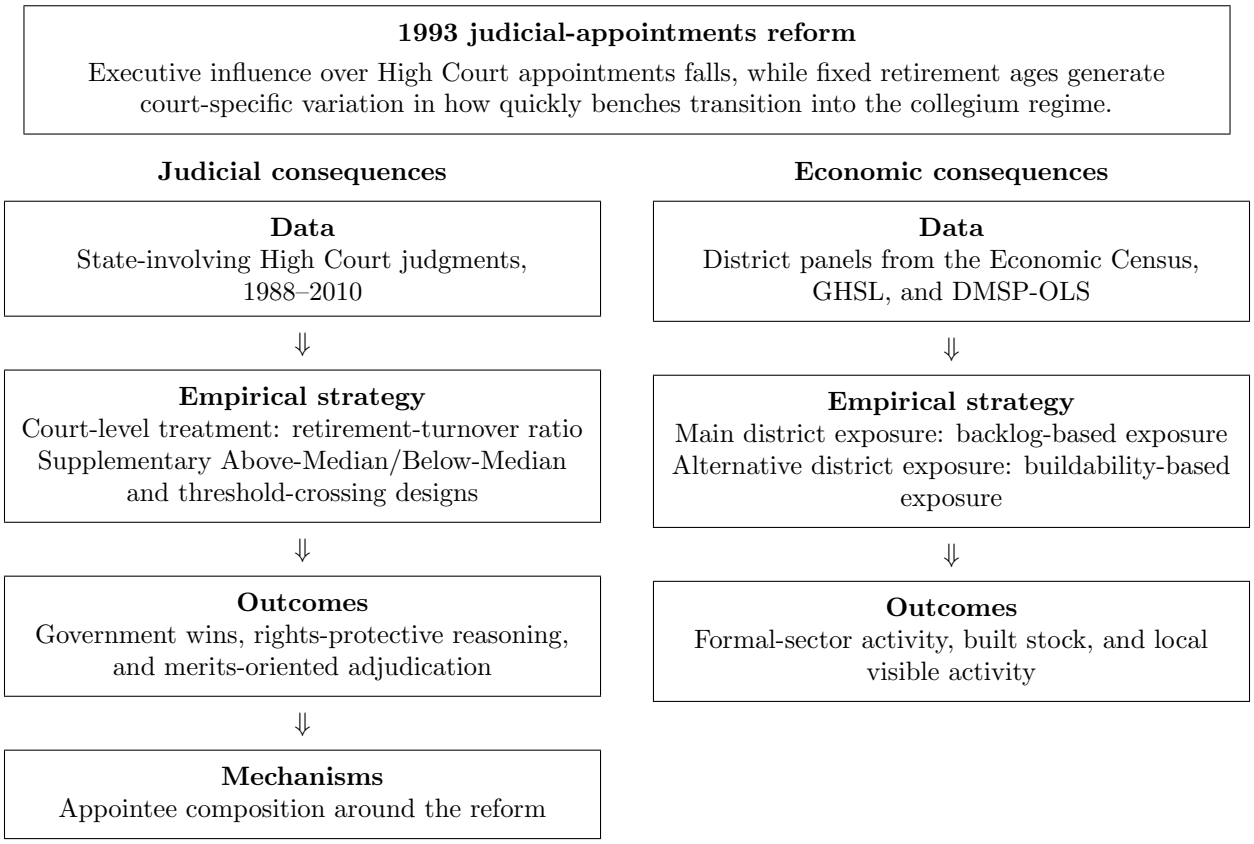


Figure 1: Roadmap of the Empirical Analysis

Notes: The figure summarizes the paper's architecture. The judicial branch uses High Court judgments to study how retirement-driven exposure to the post-1993 collegium regime changes adjudication and through which mechanisms. The economic branch uses district panels to study whether the same court-level institutional change produces larger downstream effects where pre-reform dispute backlog or baseline physical development potential implies greater exposure.

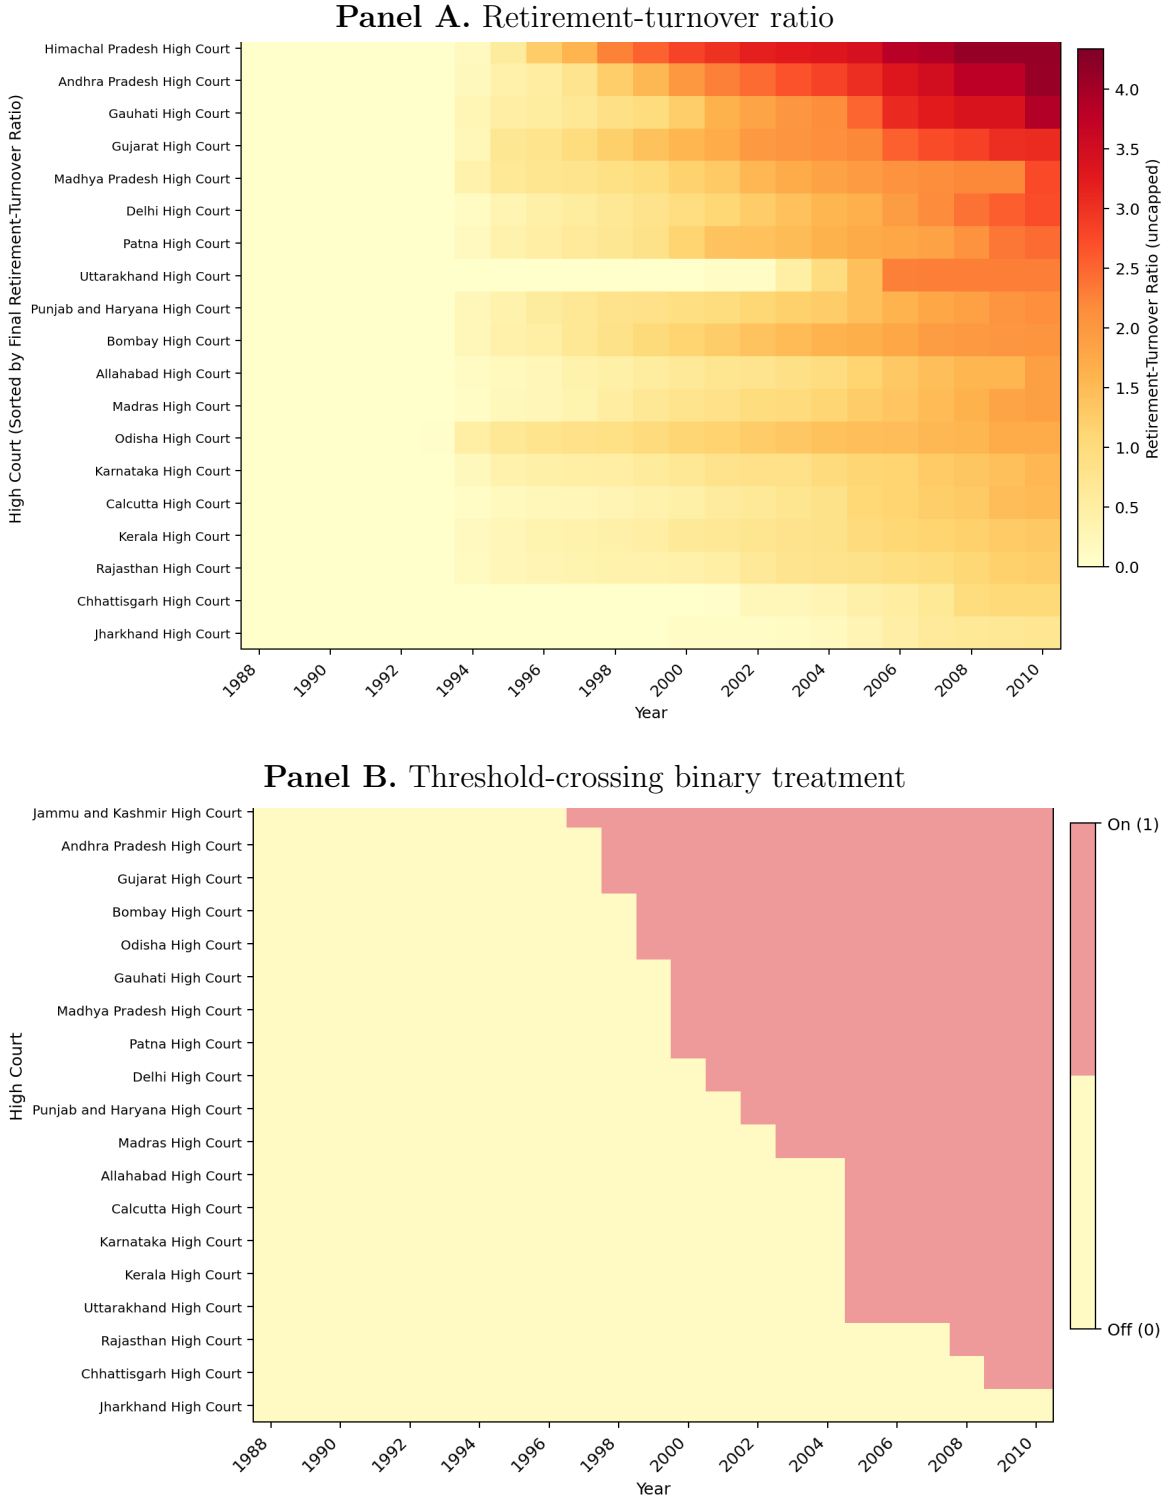


Figure 2: Court-Year Treatment Heatmaps

Notes: Panel A plots the retirement-turnover ratio by High Court and year, where the retirement-turnover ratio is cumulative post-1993 retirements divided by sanctioned bench strength. Panel B plots the corresponding threshold-crossing binary treatment, which equals one once a High Court first reaches one full bench turnover. In both panels, the unit of observation is a High Court-year, High Courts are sorted by their final observed retirement-turnover ratio, and darker shading indicates higher treatment intensity or treated status.

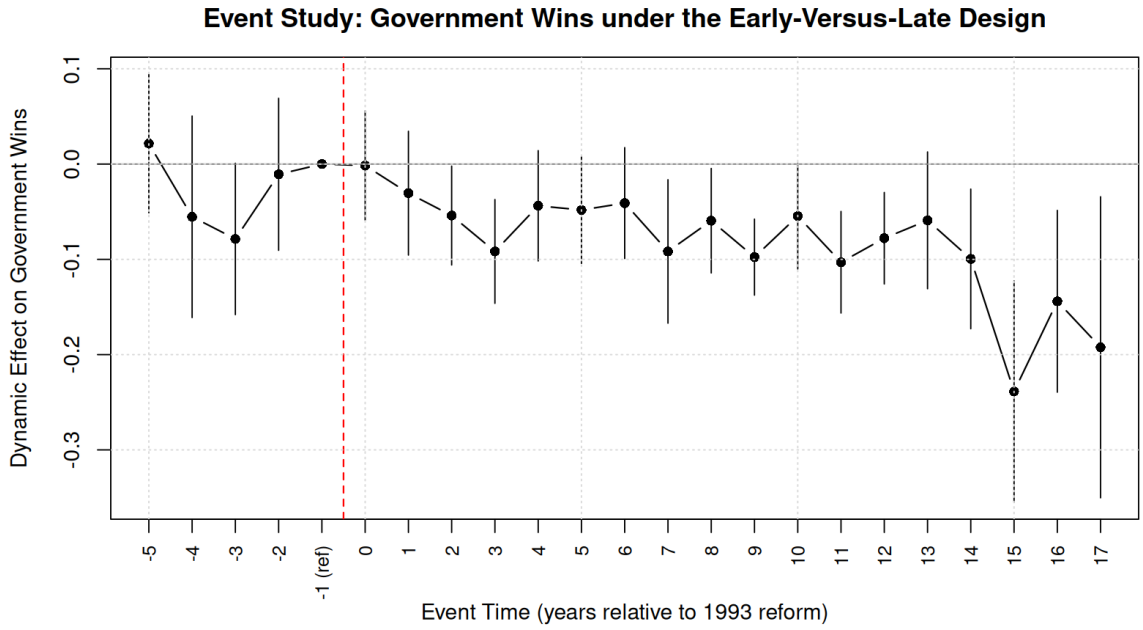


Figure 3: Dynamic Treatment Effects under the Above-Median versus Below-Median Retirement-Exposure Design

Notes: The figure reports event-time coefficients and 95% confidence intervals from the supplementary Above-Median versus Below-Median retirement-exposure design, where the treated group consists of High Courts that complete one full bench turnover by 2000. This split is equivalent to the court-level median partition in eventual retirement exposure. The unit of observation is a state-involving High Court judgment over 1988–2010. Event time is measured in years relative to the 1993 reform, and the omitted category is the year immediately before the reform (event time -1). Negative coefficients indicate lower government-win rates in Above-Median courts relative to Below-Median courts.

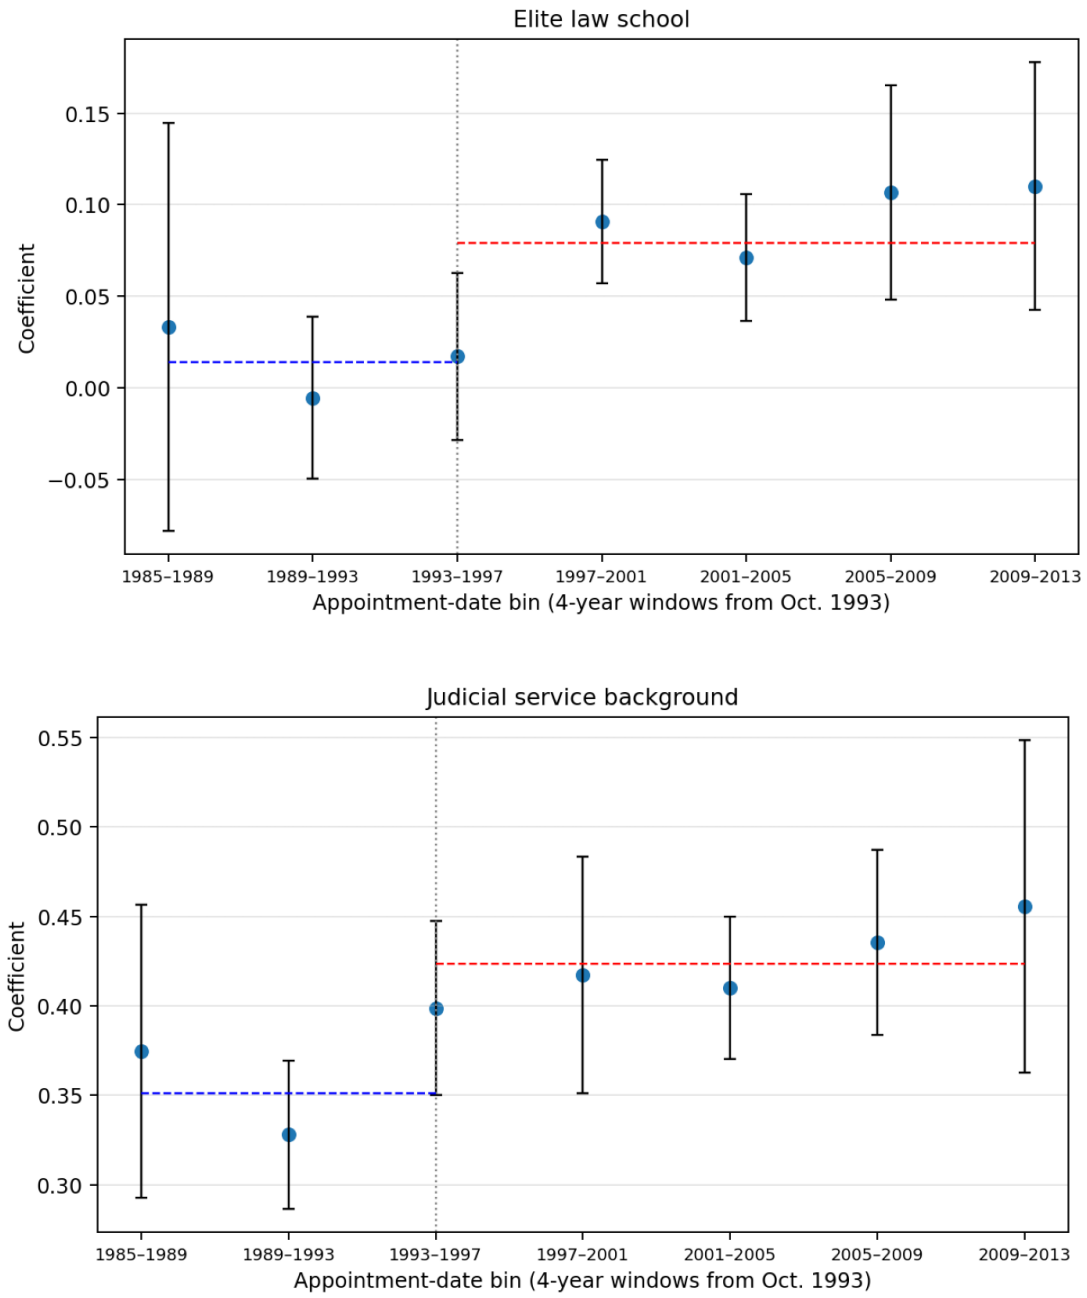


Figure 4: Binned Coefficient Estimates around the 1993 Reform

Notes: Panel A reports four-year appointment-bin coefficients for elite law school attendance. Panel B reports the corresponding coefficients for prior judicial-service experience. The unit of observation is a judge, observed once at the date of first High Court appointment over 1988–2010. Each panel comes from a separate regression with High Court fixed effects, and the vertical bars denote 95% confidence intervals based on standard errors clustered at the High Court level. The blue dashed line is the pre-reform mean, the red dashed line is the post-reform mean, and the vertical dotted line marks the first appointment bin after the October 1993 reform.

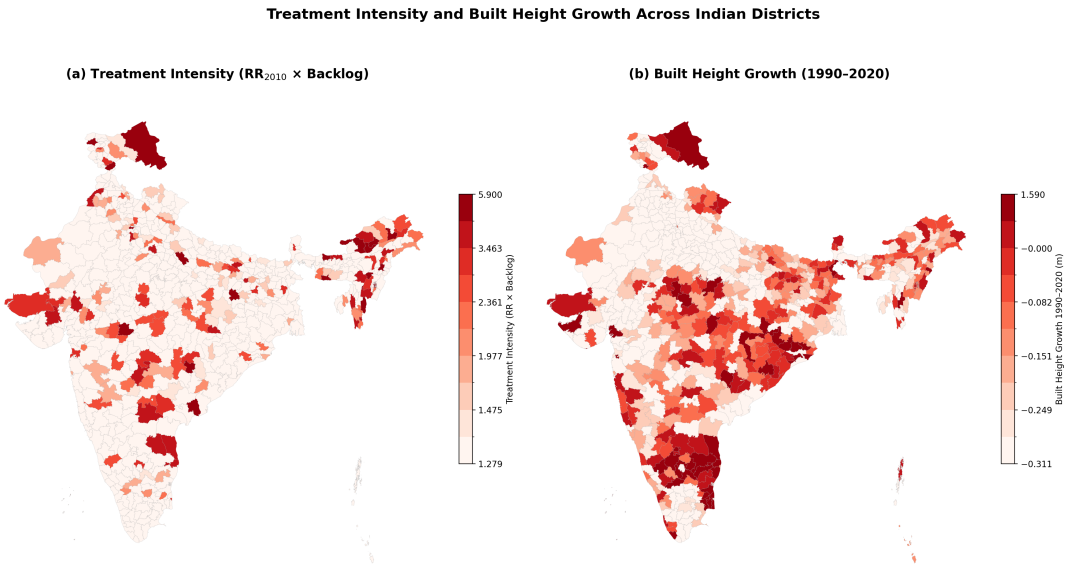


Figure 5: Treatment Intensity and Built Height Growth Across Indian Districts

Notes: Both panels use a single red color scale with decile breaks, where darker shading indicates higher values. Panel (a) shows treatment intensity at the district level, defined as the product of district d's parent High Court retirement-turnover ratio in 2010 and district d's within-High-Court standardized pre-reform backlog of pending state-involving property-rights cases. This mirrors the interaction term in the triple-difference specification and varies at the district level because the standardized backlog captures within-High-Court variation. Panel (b) shows built height growth from 1990 to 2020, measured as the change in GHSL built volume per built surface area. Raw height growth is negative in most districts (mean $= -0.30$m) due to suburban expansion of low-rise built surface; the text reports standardized values (mean zero, unit standard deviation) for interpretability.

Tables

Table 1: Descriptive Statistics

	N	Mean	SD	Min	P25	Median	Max
Outcome variables							
Government Wins	82,334	0.489	0.500	0.000	0.000	0.000	1.0
Treatment variable							
Retirement-Turnover Ratio	82,334	2.121	1.878	0.000	0.236	1.864	23.0
Case characteristics							
State is petitioner (=1)	82,334	0.226	0.418	0.000	0.000	0.000	1.0
Log judgment length	82,334	9.024	0.963	5.366	8.372	9.052	13.4
Judgment length (chars)	82,334	13,042	15,511	214	4,322	8,532	674,221
Number of citations	82,334	6.654	9.179	0.000	1.000	4.000	292
Number of times cited	82,334	2.321	25.625	0.000	0.000	0.000	5,560
Has citations (=1)	82,334	0.839	0.368	0.000	1.000	1.000	1.0
Has been cited (=1)	82,334	0.366	0.482	0.000	0.000	0.000	1.0

Notes: The sample consists of state-involving High Court judgments (1988–2010) with non-missing Government Wins outcome under the main coding. Government Wins equals one when the government fully or partly wins and zero when the government loses. The Retirement-Turnover Ratio is cumulative post-1993 retirements divided by sanctioned bench strength.

Table 2: Effect of Collegium Appointments on Government Wins involving Land cases

Dependent variable: Government Wins					
	(1)	(2)	(3)	(4)	(5)
Retirement-turnover ratio	-0.0712***	-0.0676***	-0.0636***		
	(0.0179)	(0.0187)	(0.0187)		
Post × Above-Median retirement exposure				-0.0763***	
				(0.0207)	
Staggered-adoption average treatment effect					-0.0417***
					(0.0132)
Year fixed effects	Yes	Yes	Yes	Yes	Yes
High Court fixed effects		Yes	Yes	Yes	Yes
Case controls			Yes	Yes	Yes
Observations	82,334	82,334	82,334	82,334	82,334
R-squared	0.020	0.024	0.026	0.026	0.026

Notes: The dependent variable is Government Wins, a judgment-level indicator equal to one when the government fully wins or partly wins the case and equal to zero when the government loses. The unit of observation is a state-involving High Court judgment over 1988–2010 in which the government appears on one side and a private party appears on the other. Columns (1)–(3) report a specification ladder for the main continuous treatment, the retirement-turnover ratio. The retirement-turnover ratio equals cumulative post-October-1993 retirements in High Court c divided by that court's sanctioned bench strength; it is identically zero before the reform and rises afterward as pre-reform judges retire. Column (1) includes year fixed effects, column (2) adds High Court fixed effects, and column (3) adds case controls (log judgment length and a state-petitioner indicator). Column (4) uses the Above-Median versus Below-Median retirement-exposure binary, where the Above-Median indicator equals one for High Courts that reach one full bench turnover by 2000; this split coincides with the court-level median partition in eventual retirement exposure. Column (5) reports the Sun and Abraham (2021) staggered-adoption average treatment effect under a design in which each High Court's treatment onset is the year its retirement-turnover ratio first reaches 1.0; this estimator accounts for treatment-effect heterogeneity across adoption cohorts. Standard errors in column (5) are clustered at the High Court level. Columns (4)–(5) use the full specification with all fixed effects and case controls. Standard errors in columns (1)–(4) are shown in parentheses and are clustered at the High Court×year level. Appendix Table 3A reports estimates using the accumulated new appointments ratio, which is less precisely identified because realized appointments can reflect administrative delay. Alternative outcome definitions are reported in Appendix Table 2A. * p<0.10, ** p<0.05, *** p<0.01.

Table 3: Judicial Reasoning and Case-Merit Outcomes

Panel A. Pro-Individual-Rights Judicial Reasoning Outcomes

Dependent variables: the reasoning outcomes named in columns (1)–(4)

	Scrutiny vs. deference (1)	Individual vs. systemic (2)	Human-rights and individual-rights contrast (3)	Pro-individual-rights index (4)
Retirement-turnover ratio	0.0153***	0.0250	0.0045	0.0306***
	(0.0050)	(0.0172)	(0.0074)	(0.0097)
Case controls	Yes	Yes	Yes	Yes
High Court fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	82,334	82,334	82,334	82,334
R-squared	0.005	0.045	0.016	0.024

Panel B. Case-Merit and Procedural Outcomes

Dependent variables: the case-merit and procedural outcomes named in columns (1)–(4)

	Explicit evidentiary weighing (1)	Unsupported government-claim rejection (2)	Procedural closure (3)	Merits-oriented adjudication index (4)
Retirement-turnover ratio	-0.0074	0.0093**	-0.0024**	0.0193
	(0.0110)	(0.0047)	(0.0011)	(0.0131)
Case controls	Yes	Yes	Yes	Yes
High Court fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	82,334	82,334	82,334	82,334
R-squared	0.120	0.035	0.007	0.091

Notes: The unit of observation is a state-involving High Court judgment over 1988–2010 in which the government appears on one side and a private party appears on the other. Every column uses the same full judicial-outcome specification with High Court fixed effects, year fixed effects, and case controls, log judgment length and a state-petitioner indicator. The treatment is the retirement-turnover ratio, defined as cumulative post-October-1993 retirements in High Court c divided by that court's sanctioned bench strength; because the ratio is zero before the reform, the reported coefficient is the post-reform interaction with this retirement-driven exposure. In Panel A, column (1) measures whether the judgment language reflects more scrutiny and less deference toward state action; column (2) measures whether the judgment emphasizes individualized relief rather than generalized systemic remedies; column (3) is a narrower concluding-section contrast that is higher when the operative portion of the judgment places greater weight on human rights, individual rights, scrutiny, and claimant-specific relief than on deference, order, and systemic remedies; and column (4) is the standardized average of columns (1) through (3) within the common estimation sample. In Panel B, column (1) counts explicit evidentiary-weighing language; column (2) counts language rejecting unsupported government claims; column (3) counts procedural-closure language in the conclusion or operative portion of the judgment; and column (4) is a standardized merits-oriented adjudication index equal to the average of the standardized versions of columns (1) and (2) minus the standardized version of column (3). Standard errors are shown in parentheses and are clustered at the High Court×year level. The human-rights and individual-rights contrast score is retained as moderate evidence because text audit of the highest-scoring cases shows compensation, personal liberty, illegal detention, and petitioner-specific relief language, even though the measure remains exploratory. * p<0.10, ** p<0.05, *** p<0.01.

Table 4: Economic Effects of Judicial Independence

Panel A: Economic Census (SHRUG), District × EC Round						
	(1) Formal Emp (log)	(2) Total Emp (log)	(3) Informal Emp (log)	(4) Avg Estab Size	(5) Hired Labor Share	(6) Formal Estab Share
Retirement-turnover ratio × backlog exposure	0.0691**	0.0264*	0.0171	0.0489**	0.0065***	0.0004*
	(0.0268)	(0.0152)	(0.0113)	(0.0219)	(0.0021)	(0.0002)
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Round FE	Yes	Yes	Yes	Yes	Yes	Yes
Constituent interaction terms	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,349	2,349	2,349	2,349	2,349	2,349
R^2	0.854	0.910	0.914	0.670	0.909	0.789
Panel B: Built Environment (GHSL R2023A), District × 5-Year Epoch						
	(7) Built Volume (log m^3)	(8) Built Height (m)	(9) New Built Surface (km^2)	(10) New Built Volume (M m^3)	(11) Built Surface (log km^2)	
Retirement-turnover ratio × backlog exposure	0.0117***	0.0110***	0.1808**	1.741***	0.0032	
	(0.0041)	(0.0037)	(0.0766)	(0.533)	(0.0055)	
District FE	Yes	Yes	Yes	Yes	Yes	
Epoch FE	Yes	Yes	Yes	Yes	Yes	
Constituent interaction terms	Yes	Yes	Yes	Yes	Yes	
Observations	6,009	6,009	5,428	5,428	6,009	
R^2	0.986	0.984	0.654	0.757	0.979	
Panel C: Nightlights (DMSP-OLS), District × Year						
	(12) Total Light (log)					
Retirement-turnover ratio × backlog exposure	0.0388***					
	(0.0084)					
District FE	Yes					
Year FE	Yes					
Constituent interaction terms	Yes					
Observations	13,338					
R^2	0.953					

Notes: The dependent variables are listed at the top of each column. Each column reports $\hat{\delta}$ from a specification that interacts the parent High Court's retirement-turnover ratio with district-level backlog exposure. The retirement-turnover ratio equals cumulative post-October-1993 retirements in the parent High Court divided by that court's sanctioned bench strength, so it is identically zero before the reform and rises mechanically as pre-reform judges retire. Backlog exposure is the within-High-Court z-score of pre-October-1993 pending state-involving property-rights cases from district d.

Panel A uses the Economic Census across four rounds (1990, 1998, 2005, 2013). Panel B uses GHSL R2023A built surface, built volume, and built height across ten five-year epochs (1975–2020); "M m^3" denotes million cubic meters. Panel C uses DMSP-OLS stable nightlights (1992–2013).

All columns include district fixed effects, period fixed effects, the court-level treatment, the district exposure term, and the constituent two-way interaction terms implied by the triple-difference design. Standard errors clustered at the High Court level (21 clusters) are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 5: Alternative Economic Effects Under the Buildability-Exposure Identification

Panel A: Built Environment (GHSL R2023A), District × 5-Year Epoch

	(1) Built Volume (log m^3)	(2) Built Height (m)	(3) Built Surface (log km^2)
Retirement-turnover ratio × buildability exposure	0.0236***	0.0213**	0.0223***
	(0.0059)	(0.0095)	(0.0043)
District FE	Yes	Yes	Yes
Epoch FE	Yes	Yes	Yes
Constituent design terms	Yes	Yes	Yes
Observations	6,410	6,410	6,410
Districts	641	641	641
High Courts	18	18	18
R^2	0.987	0.983	0.980

Panel B: Nightlights (DMSP-OLS), District × Year

	(4) Total Light (log)	(5) Mean Light (log)
Retirement-turnover ratio × buildability exposure	0.0276	0.0185**
	(0.0198)	(0.0081)
District FE	Yes	Yes
Year FE	Yes	Yes
Constituent design terms	Yes	Yes
Observations	14,102	14,102
Districts	641	641
High Courts	18	18
R^2	0.952	0.971

Notes: The dependent variables are listed at the top of each column. Each column reports the coefficient on the interaction between the parent High Court's retirement-turnover ratio and district buildability exposure. The retirement-turnover ratio equals cumulative post-October-1993 retirements in the parent High Court divided by that court's sanctioned bench strength, so it is identically zero before the reform and rises mechanically as pre-reform judges retire. Buildability exposure equals physical buildability multiplied by underbuilt slack, where underbuilt slack is defined as low 1990 built volume conditional on physical buildability.

Panel A uses GHSL built outcomes across ten five-year epochs (1975–2020). Panel B uses DMSP-OLS nightlights (1992–2013). All columns include district fixed effects, period fixed effects, the court-level treatment, the district exposure term, and the constituent two-way interaction terms implied by the design. Standard errors clustered at the High Court level are in parentheses. Appendix Table 7A reports the stricter accessibility-adjusted residual-slack version.

* p<0.10, ** p<0.05, *** p<0.01.

Appendix

Appendix A collects judicial-outcome definitions and judicial-design robustness checks. Appendix B collects district-level economic-outcome definitions, exposure construction, pre-reform economic balance checks, sample accounting, and the robustness version of the alternative economic identification strategy. Appendix C reports the additional binary-design evidence referenced in the judicial-results section.

- **Appendix A. Judicial-outcome definitions and judicial-design robustness**
  - A.1 Judicial data construction and machine-assisted coding
  - A.2 Outcome-definition robustness
  - A.3 Alternative continuous treatment definition
  - A.4 Pre-reform judicial balance checks
- **Appendix B. Economic outcomes and alternative economic identification**
  - B.1 District-level economic panels, exposure measures, and sample support
  - B.2 Satellite built-environment and nightlights methodology
  - B.3 Alternative economic identification robustness
- **Appendix C. Supplementary binary-design robustness**

Appendix A. Judicial-Outcome Definitions and Judicial-Design Robustness

Appendix A.1 explains how the judicial sample is defined and how the core outcomes and judge-level variables are constructed. Appendix A.2 shows that the main judicial result does not depend on one particular coding of mixed case outcomes. Appendix A.3 reports a parallel specification that replaces the retirement-based treatment with an appointments-based treatment. Appendix A.4 verifies that courts above and below the median of eventual retirement exposure do not appear systematically different in observable pre-reform case composition or judge composition before the 1993 reform.

A.1 Judicial Data Construction and Machine-Assisted Coding

This subsection defines the judicial outcomes and judge-level selection variables.

Judicial sample and unit of observation. The main judicial sample consists of state-involving High Court judgments in land and property-rights disputes over 1988–2010. The unit of observation is a High Court judgment. The underlying judgment text is harvested from Indian Kanoon's published High Court corpus. Government involvement is coded from the parsed petitioner and respondent names in the case title. A case is classified as state-involving when one side matches a pre-specified state-entity dictionary, which includes *State of ...*, *Union of India*, ministries, departments, municipal corporations, development authorities, boards, commissions, police authorities, revenue authorities, and other public bodies. The property-rights restriction uses a combined screen that equals one when the case either triggers the pre-specified land/property-rights conflict keyword screen or cites one of the pre-specified property-rights statutes. The judicial-outcome regressions in the main text and in Tables 2 and 3 are estimated on this judgment-level sample.

Appendix Table 1A reports how the 82,334-judgment sample is distributed across High Courts. Coverage is not uniform: Madras and Kerala each contribute more than 10,000 judgments, while several smaller or later-created courts contribute far fewer. This heterogeneity matters for description, but the identifying variation in the judicial regressions comes from within-court changes over time, with High Court fixed effects absorbing persistent differences in publication volume and court size.

Table 1A: Judicial Sample Coverage by High Court

High Court	Judgments	Share of sample (%)	First year	Last year
Madras High Court	11,060	13.43	1988	2010
Kerala High Court	10,735	13.04	1988	2010
Delhi High Court	8,829	10.72	1988	2010
Punjab and Haryana High Court	7,080	8.60	1988	2010
Karnataka High Court	7,072	8.59	1988	2010
Bombay High Court	5,705	6.93	1988	2010
Allahabad High Court	5,008	6.08	1988	2010
Rajasthan High Court	4,687	5.69	1988	2010
Gujarat High Court	4,535	5.51	1988	2010
Andhra Pradesh High Court	4,485	5.45	1988	2010
Calcutta High Court	3,205	3.89	1988	2010
Madhya Pradesh High Court	2,849	3.46	1988	2010
Patna High Court	2,730	3.32	1988	2010
Odisha High Court	1,686	2.05	1988	2010
Gauhati High Court	1,067	1.30	1988	2010
Jharkhand High Court	602	0.73	1996	2010
Jammu and Kashmir High Court	452	0.55	1988	2010
Himachal Pradesh High Court	354	0.43	1988	2008
Chhattisgarh High Court	106	0.13	2000	2010
Uttarakhand High Court	70	0.09	2001	2007
Sikkim High Court	17	0.02	1990	2006

Notes: The table reports the distribution of the 82,334 state-involving land and property-rights judgments used in the main judicial analysis. Each row sums the number of sample judgments observed in the published High Court corpus for that court over 1988–2010. The table is meant to document coverage heterogeneity, not to define treatment intensity. Treatment is identified from within-court changes in retirement-driven bench turnover over time, with court fixed effects absorbing persistent differences in court size and publication volume.

Case form and bench structure. The judgment-level sample includes both original writ matters and appellate or revisional High Court decisions whenever those judgments satisfy the paper's state-involving property-rights restriction. This is why the unit of observation is a High Court judgment rather than a dispute-level case history. The paper does not require every observation to be a writ petition, because economically relevant state-citizen property conflict reaches the High Courts through both original and appellate channels.

The selection analysis is correspondingly anchored to the identified authoring or presiding judge rather than to the full bench composition of every case. Where a published judgment clearly identifies a single author or presiding judge, that judge is the natural unit for linking case outcomes to judge characteristics. Where bench metadata are incomplete or multi-judge

decisions do not expose a clean author field, the case remains in the main judicial sample but may drop from the judge-identified subsample used for the appointment-composition evidence. Figure 4 should therefore be read as selection evidence for the judge-identified subset of the judicial sample, not as a complete decomposition of every bench member's vote in every High Court judgment.

Government Wins. The main judicial outcome, *Government Wins*, is constructed in two steps. First, each judgment is assigned a dispositive label based on the concluding portion of the opinion, where Indian High Court judgments typically state whether the petition or appeal is allowed, partly allowed, dismissed, or otherwise disposed of. We classify each judgment into one of five mutually exclusive categories: `ALLOWED`, `PARTLY_ALLOWED`, `DISMISSED`, `DISPOSED`, and `UNCLEAR`. The coding task is run under a fixed instruction set so that each judgment receives a reproducible raw disposition label. Second, we convert that label into a petitioner-win score and then into a government-win score using the identity of the government litigant.

Let PetitionerWin_i equal 1 for `ALLOWED`, 0.5 for `PARTLY_ALLOWED`, 0 for `DISMISSED`, and missing for `DISPOSED` or `UNCLEAR`. Let StatePetitioner_i equal one when the government is the petitioner and zero when the government is the respondent. The government-win score is then

$$\begin{aligned}\text{GovernmentWin}_i = \text{StatePetitioner}_i \cdot \text{PetitionerWin}_i \\ + (1 - \text{StatePetitioner}_i) \cdot (1 - \text{PetitionerWin}_i)\,.\end{aligned}$$

The paper reports three codings of this object. The main text uses a binary outcome that counts partly allowed cases as government wins because that coding is the closest to the question of whether the state prevailed in the dispute. Appendix Table 2A then shows that the main result is robust when partly allowed cases are instead coded as 0.5 or removed from the sample.

Two checks validate this construction. First, the disposition coding is checked against manually inspected judgments to verify that the label corresponds to the actual dispositive paragraph. Second, the coding is compared across alternative text windows to confirm that reading the ending portion of the judgment is sufficient to recover the operative outcome. These checks matter because the judicial results would be difficult to interpret if the state-win variable were merely a noisy text label.

Judge-level characteristics used in the selection analysis. The selection discussion links the timing of appointment to the characteristics of the incoming judge. The unit of observation for these variables is the judge, anchored at the date of first High Court appointment. The core judge attributes are those emphasized in Figure 4: elite law school attendance and prior judicial-service experience before High Court appointment. The underlying biography file also records age at appointment, gender, advocate background, executive-service background, and name-based religion proxies. These variables are assembled from structured judicial biographies and standardized profile fields, with manual review whenever multiple public sources disagree on dates, schooling, or prior service. The judge-identified subsample is smaller than the full judicial sample because each observation must be linked to a uniquely

identified presiding judge in the biography file. In this sample, that requirement leaves 53,083 judge-identified judgments out of 82,334 judgments with non-missing main outcomes. Cases usually leave the judge-identified subsample because the published judgment does not expose a clean author field or because the observed author string cannot be linked uniquely to the TCPD biography file. These variables capture the observable channels through which the post-1993 appointments regime changed who entered the bench.

Text-based reasoning measures. The reasoning outcomes in Panel A of Table 3 are text-based measures constructed from pre-specified phrase lists and fixed text windows. They measure whether opinions become less deferential to the state and more focused on the individual claimant.

Columns (1) and (2) are court-relative phrase-axis measures. For each judgment, we count matches to two opposing phrase dictionaries and form an axis score as the log difference between the support-side count and the counter-side count. For *scrutiny versus deference*, the support dictionary contains scrutiny and active-review language, such as references to judicial review, searching scrutiny, or the need to examine the legality of state action, while the counter dictionary contains restraint and deference language, such as policy discretion, judicial restraint, or limited review. For *individual versus systemic*, the support side contains individualized-relief and claimant-specific language, while the counter side contains systemic-remedy, institutional-management, and generalized public-order language. These judgment-level axis scores are then expressed relative to the corresponding High Court's pre-1988 baseline mean, so that higher values indicate a shift toward more scrutiny in column (1) and more claimant-specific reasoning in column (2) relative to that court's pre-1988 reasoning baseline.

Column (3), *human-rights and individual-rights contrast*, is constructed on a narrower conclusion-focused reasoning slice rather than on the full opinion text. The text window is taken from the final portion of the judgment, where Indian High Court decisions usually state the operative reason for granting or denying relief. Within that window, we count support phrases referring to human rights, individual rights, scrutiny, personal liberty, and petitioner-specific relief, and we count counter phrases referring to deference, public order, state-capacity justifications, and systemic remedies. The reported measure is the log difference between the support count and the counter count. This variable is therefore narrower than columns (1) and (2): it is designed to isolate the language of the operative holding rather than the broader tone of the full opinion.

Column (4), *pro-individual-rights index*, is the standardized average of columns (1) through (3) in the common estimation sample. It summarizes whether a judgment moves in a more rights-protective and claimant-oriented direction across these three distinct reasoning margins.

The first two variables capture court-relative movement along stable reasoning axes, whereas the third captures the operative conclusion of the judgment using a separate conclusion-window contrast. Together they let the paper ask not only whether the state wins less often, but also whether opinions increasingly justify those outcomes through less deference and more claimant-centered reasoning.

Case-merit and procedural measures. The case-merit outcomes in Panel B of Table 3 are built to distinguish judgments that engage the dispute on the merits from judgments that end the matter on threshold procedural grounds. The explicit-evidentiary-weighing measure counts language that evaluates facts, evidence, or the record. The unsupported-government-claim rejection measure counts language that explicitly rejects government assertions as unproven or unsupported. The procedural-closure measure counts language indicating dismissal or disposal on maintainability, jurisdiction, limitation, or other threshold grounds. These objects are then standardized and combined into a merits-oriented adjudication index. These measures are informative because they reveal which aspect of adjudication changes: whether judges weigh disputes more carefully, push back against unsupported state claims, and rely less on procedural closure.

The text inputs for these measures come from pre-specified phrase dictionaries grouped into substantive-merits categories and procedural-closure categories. The continuous measures used in the paper are therefore transparent dictionary-based objects. In addition, the broader data construction uses a binary procedural-disposition classifier trained on a labeled sample of judgments to distinguish merit-based dispositions from threshold procedural resolutions. The results in Table 3 use the phrase-based continuous measures because they are constructed from fixed, pre-specified dictionaries and are therefore directly reproducible. The classifier provides a complementary check.

Machine-assisted coding and reproducibility. Some judicial variables rely on automated text classification, but the replication logic is fixed. Disposition labels are generated from the final 6,000 characters of each judgment under a fixed five-label prompt that permits only `ALLOWED`, `PARTLY_ALLOWED`, `DISMISSED`, `DISPOSED`, or `UNCLEAR`. The 1998–2010 extension uses the OpenAI mini model `gpt-4o-mini` with deterministic decoding (temperature 0 and a 20-token response cap). The earlier archived extraction wave uses the same task structure and deterministic decoding, and the analysis uses those frozen archived labels rather than live API calls. The phrase-axis reasoning measures and the dictionary-based case-merit measures are then constructed deterministically from the stored judgment text and the fixed phrase lists described above. Any manual correction applied to an ambiguous case is recorded in a transparent audit log. The empirical results therefore use frozen coded variables whose legal and economic meaning is defined in this appendix.

A.2 Outcome-Definition Robustness

The main judicial outcome in the paper is *Government Wins*, a binary indicator equal to one if the government fully wins or partly wins the case and equal to zero if the government loses. This coding maps most directly to the question of whether faster transition into the collegium regime makes High Courts less favorable to the state in state-involving land and property-rights disputes.

Table 2A tests whether the main result depends on that exact coding choice. It reports the same judicial treatment effect under three outcome definitions: the main coding, a continuous coding in which partly allowed cases are assigned value 0.5, and a conservative binary coding in which partly allowed cases are excluded. The table reports three estimators for each coding: the preferred continuous treatment based on the retirement-turnover ratio,

the staggered-adoption average treatment effect, and the alternative continuous treatment based on accumulated new appointments.

If the negative relationship between judicial turnover and government wins is truly a feature of adjudication rather than an artifact of how mixed outcomes are coded, then the sign should survive across these alternative definitions. That is exactly what Table 2A shows: all three treatment constructions continue to point to lower government-win rates in faster-turnover courts. The result is therefore not driven by a fragile coding decision at the margin of partly allowed cases.

- *Government Wins:* partly allowed cases are coded as government wins.
- *Continuous coding:* partly allowed cases are coded as 0.5 rather than 1.
- *Conservative binary coding:* partly allowed cases are excluded from the sample.

Table 2A: Outcome-Definition Robustness for the Judicial-Outcome Designs

Dependent variable: Government Wins under the coding noted in each column			
	(1)	(2)	(3)
	Main coding	Partial wins coded as 0.5	Partial wins excluded
Retirement-turnover ratio	-0.0636***	-0.0593***	-0.0629***
	(0.0187)	(0.0178)	(0.0188)
Staggered-adoption average treatment effect	-0.0417***	-0.0278**	-0.0299**
	(0.0132)	(0.0124)	(0.0135)
Accumulated new appointments ratio	-0.0525***	-0.0502***	-0.0539***
	(0.0132)	(0.0126)	(0.0136)
Observations	82,334	82,334	75,682

Notes: Column (1) repeats the paper's main outcome, which counts partly allowed cases as government wins. Column (2) uses a continuous coding in which partly allowed cases are assigned value 0.5. Column (3) excludes partly allowed cases. The first row reports the main continuous-treatment specification using the retirement-turnover ratio with High Court fixed effects, year fixed effects, case controls, and High Court×year clustered standard errors. The retirement-turnover ratio equals cumulative post-October-1993 retirements in High Court c divided by that court's sanctioned bench strength and is identically zero before the reform. The second row reports the Sun and Abraham (2021) staggered-adoption average treatment effect, with standard errors clustered at the High Court level. The third row reports estimates using the accumulated new appointments ratio, which is less precisely identified because realized appointments can reflect administrative delay, with High Court fixed effects, year fixed effects, case controls, and High Court×year clustered standard errors. Standard errors are shown in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

A.3 Alternative Continuous Treatment Definition

The main text uses the retirement-turnover ratio because retirements map the speed of institutional replacement after 1993 in the cleanest way. A natural alternative is the *accumulated new appointments ratio*, defined as the cumulative number of appointment spells that begin in court c on or after October 6, 1993, divided by the sanctioned bench strength of court c. Like the retirement-turnover ratio, this measure is zero before the reform and rises afterward. Unlike the retirement-turnover ratio, however, it can also reflect administrative delays between collegium recommendation and realized appointment.

Table 3A reports the same three-step specification ladder used in the main judicial table, but with the accumulated new appointments ratio replacing the retirement-turnover ratio. The table therefore answers a narrow and important question: if exposure to the post-1993 appointments regime is measured with realized appointments rather than retirements, does the court-level result continue to point in the same direction?

The answer is yes, but with the interpretation that the alternative treatment is weaker. Table 3A continues to show a negative and statistically significant relationship between post-reform institutional replacement and government-win rates. This reinforces the main finding by showing that it is not unique to one continuous treatment construction. At the same time, the paper does not elevate the appointments-based measure to the main design because realized appointments can be influenced by post-recommendation executive timing, whereas retirements are mechanically determined by age.

Table 3A: Government Wins and the Accumulated New Appointments Ratio

Dependent variable: Government Wins			
	(1)	(2)	(3)
Accumulated new appointments ratio	-0.0528***	-0.0541***	-0.0513***
	(0.0124)	(0.0129)	(0.0131)
Year fixed effects	Yes	Yes	Yes
High Court fixed effects		Yes	Yes
Case controls			Yes
Observations	82,334	82,334	82,334
R-squared	0.020	0.024	0.026

Notes: This table reports estimates using an alternative continuous treatment based on accumulated new appointments. The identification is weaker than for the main retirement-turnover ratio because the timing of realized appointments can be affected by delays in final executive action even after collegium nomination. The dependent variable is Government Wins. The unit of observation is a state-involving High Court judgment over 1988–2010 in which the government appears on one side and a private party appears on the other. The treatment is the accumulated new appointments ratio, where the numerator is the cumulative number of appointment spells that begin in the High Court on or after October 6, 1993 and the denominator is sanctioned strength; the variable is identically zero before the reform. For reference, the paper's preferred retirement-turnover ratio is cumulative post-October-1993 retirements divided by sanctioned bench strength. Column (1) includes year fixed effects. Column (2) adds High Court fixed effects. Column (3) adds case controls, log judgment length and a state-petitioner indicator. Standard errors are shown in parentheses and are clustered at the High Court×year level. * p<0.10, ** p<0.05, *** p<0.01.

A.4 Pre-Reform Judicial Balance Checks

The main judicial design relies on cross-court differences in how quickly benches turn over after the 1993 reform. A natural concern is that courts above and below the median of eventual retirement exposure might already have looked systematically different before the reform in ways that would predict later government-win rates even without the appointments shock. The two most immediate observable threats are pre-reform case composition and pre-reform judge composition.

Appendix Figure 1A collects two judicial balance exercises in a common standardized-difference format. Panel A uses a pre-reform High Court×year panel of state-involving land and property judgments to compare case-composition margins across Above-Median and Below-Median retirement-exposure courts. Panel B uses a pre-reform High Court×month panel to compare judge-composition margins across the same court partition. In both cases, the Above-Median versus Below-Median split is the paper's binary partition based on whether the retirement-turnover ratio reaches one by 2000; as discussed in the main text, this split coincides with the court-level median partition in eventual retirement exposure. The displayed coefficients are standardized mean differences with 95 percent confidence intervals. No displayed coefficient remains significant at the 10 percent level after Holm correction within the relevant family of tests.

Table 4A then provides the numeric judge-composition coefficients underlying the judge-balance part of Figure 1A. The logic of the table is diagnostic rather than causal: it asks whether eventual retirement exposure was already predictable from observable bench composition before October 1993. The answer is no. No attribute remains significant after Holm correction, which supports the interpretation that retirement-driven treatment intensity is not simply proxying for a pre-existing composition difference across High Courts.

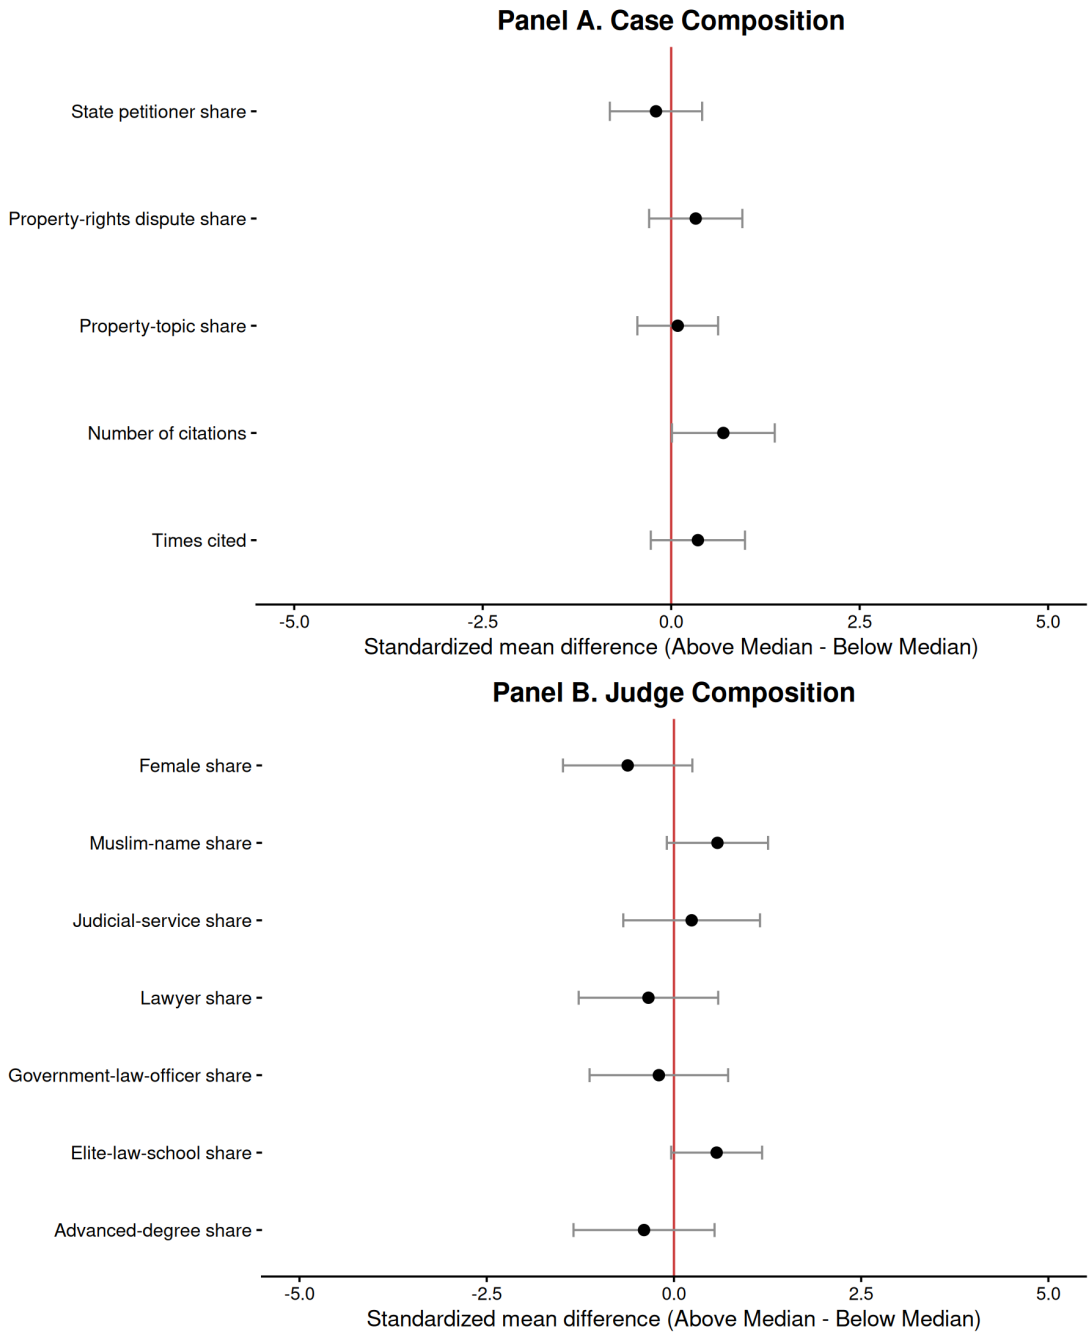


Figure 1A: Pre-Reform Judicial Balance: Case Composition and Judge Composition

Notes: The figure compares High Courts above and below the median of eventual retirement exposure, using the paper's binary split based on whether the retirement-turnover ratio reaches one by 2000. This split coincides with the court-level median partition in retirement exposure. Panel A reports standardized mean differences for pre-reform case-composition margins from a High Court×year panel of state-involving land and property judgments. Panel B reports standardized mean differences for pre-reform judge-composition margins from a High Court×month panel. Vertical bars denote 95% confidence intervals. Judgment length is excluded because it appears too sensitive to procedural and formatting noise to function as a clean balance variable. No displayed coefficient remains statistically significant at the 10 percent level after Holm correction within the relevant family of tests.

Table 4A: Balance Test: Pre-Reform Judge Composition and Treatment Assignment

	Female	Muslim	Judicial service	Lawyer	Govt. law officer	Elite law school	Advanced degree
Above-Median retirement exposure	−0.0168	0.0197*	0.0456	−0.0821	−0.0369	0.0358*	−0.0446
	(0.0120)	(0.0117)	(0.0897)	(0.1146)	(0.0864)	(0.0195)	(0.0536)
Holm p-value	[0.805]	[0.552]	[1.000]	[1.000]	[1.000]	[0.460]	[1.000]
Dependent-variable mean	0.013	0.016	0.213	0.437	0.286	0.031	0.148
Year-month fixed effects	Yes						
Cluster level	High Court						
Observations	1,242						
High Courts	18						

Notes: Each column reports a separate OLS regression of a pre-reform High Court bench-composition variable on an indicator for Above-Median retirement exposure, with year-month fixed effects. The unit of observation is a High Court × year-month cell over 1988:1–1993:9. The Above-Median indicator equals one for High Courts whose retirement-turnover ratio reaches one by 2000; this split coincides with the court-level median partition in eventual retirement exposure. Dependent variables are the share of sitting judges with the indicated attribute in a given court-month. Standard errors in parentheses are clustered at the High Court level. Holm p-values in brackets control the family-wise error rate across all seven balance tests. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Appendix B. Economic Outcomes and Alternative Economic Identification

This section defines the district-level outcome panels and exposure measures, explains the sample-support differences across tables, and documents the robustness version of the buildability identification.

B.1 District-Level Economic Panels, Exposure Measures, and Sample Support

The main text combines three district-level outcome families and two district-exposure objects. Appendix Table 5A reports the distributions of the district-level economic outcomes. Appendix Table 6A reports the distributions of the exposure variables and the exact observation, district, and High Court support behind the main and alternative economic designs.

Table 5A: Descriptive Statistics for District-Level Economic Outcomes

	N	Mean	SD	Min	P25	Median	P75	Max
Economic Census (district × census round)								
Formal employment (log)	2,468	8.907	1.761	0.000	7.914	8.962	9.974	15.156
Total employment (log)	2,468	11.013	1.382	3.871	10.335	11.097	11.805	16.062
Average establishment size	2,468	2.398	0.802	1.401	1.945	2.186	2.590	13.490
Hired labor share	2,468	0.385	0.227	0.000	0.259	0.438	0.541	0.925
Formal establishment share	2,468	0.007	0.006	0.000	0.003	0.005	0.008	0.082
GHSL built environment (district × five-year epoch)								
Built volume (log m^3)	6,410	18.303	1.316	12.897	17.620	18.493	19.190	21.741
Built height (m)	6,410	4.316	1.768	2.500	3.231	3.814	4.801	18.729
Built surface (log km^2)	6,410	3.179	1.044	0.137	2.574	3.323	3.926	6.110
Nightlights (district × year)								
Total light (log)	14,102	8.854	1.886	0.000	7.999	9.305	10.117	12.344
Mean light (log)	14,102	1.339	0.888	0.000	0.640	1.294	1.898	4.159

Notes: The table reports summary statistics for the district-level outcomes used in the economic analysis. Economic Census outcomes are observed in four rounds (1990, 1998, 2005, and 2013). GHSL built-environment outcomes are observed in ten five-year epochs from 1975 to 2020 and measure district-wide built stock rather than direct construction flows. Nightlights are annual DMSP-OLS stable-lights aggregates from 1992 to 2013. Built volume equals district-level built surface multiplied by district-level average gross building height. Built surface and built volume therefore measure different margins of physical capital accumulation than nightlights, which proxy visible local economic activity rather than the stock of structures.

Table 6A: District Exposure Measures and Estimation-Sample Support

	N	Mean	SD	Min	P25	Median	P75	Max
Panel A: Exposure distributions								
Backlog count	640	277.484	303.901	1.000	94.750	164.500	436.000	3,469.000
Backlog share within parent High Court	640	0.027	0.024	0.000	0.014	0.019	0.035	0.251
Backlog exposure (within-parent-High-Court z-score)	640	0.000	0.987	-3.993	-0.437	0.155	0.636	2.698
Physical buildability index	641	0.000	0.640	-2.272	-0.092	0.228	0.404	0.599
Preferred buildability exposure	641	0.000	1.000	-7.257	-0.229	-0.001	0.253	8.823
Accessibility-adjusted buildability exposure	606	0.000	1.000	-7.830	-0.246	-0.014	0.237	6.594

Panel B: Outcome-panel support Sample	Unit	Obs.	Districts	High Courts	Coverage Coverage rule
Raw Economic Census panel	District × census round	2,560	640	20	1990, 1998, 2005, and 2013 rounds
Raw GHSL panel	District × five-year epoch	6,410	641	21	1975–2020, ten epochs
Raw nightlights panel	District × year	14,102	641	21	1992–2013 annual series
Backlog design support: Economic Census	Estimation sample	2,349	640	20	Main backlog-exposure specification
Backlog design support: GHSL	Estimation sample	6,009	640	21	Main backlog-exposure specification
Backlog design support: nightlights	Estimation sample	13,338	640	21	Main backlog-exposure specification
Preferred buildability design: Economic Census	Estimation sample	2,452	636	17	Preferred physical-buildability exposure
Preferred buildability design: GHSL	Estimation sample	6,410	641	18	Preferred physical-buildability exposure
Preferred buildability design: nightlights	Estimation sample	14,102	641	18	Preferred physical-buildability exposure
Robustness buildability design: Economic Census	Estimation sample	2,353	602	15	Accessibility-adjusted exposure robustness
Robustness buildability design: GHSL	Estimation sample	6,060	606	16	Accessibility-adjusted exposure robustness
Robustness buildability design: nightlights	Estimation sample	13,332	606	16	Accessibility-adjusted exposure robustness

Notes: Panel A reports the distributions of the district exposure variables used in the main and alternative economic designs. Backlog exposure is the within-parent-High-Court standardized count of pending state-involving property-rights cases linked to district d immediately before October 6, 1993. The physical buildability index is the district-level geophysical feasibility measure defined in Appendix B.3. Preferred buildability exposure equals physical buildability multiplied by underbuilt slack, where underbuilt slack is the negative standardized residual from a 1990 built-volume regression on physical buildability alone. The robustness exposure replaces that slack term with the negative standardized residual from a 1990 built-volume regression on physical buildability and accessibility.
Panel B explains why the number of observations, districts, and High Courts varies across economic tables. The raw Economic Census, GHSL, and nightlights panels are harmonized to common district geography but differ in temporal coverage and district support. The backlog design is available only where districts can be linked to the pre-reform backlog inventory. The buildability designs additionally require district-level geophysical inputs, 1990 GHSL built-volume support, and, in the appendix robustness version, non-missing accessibility. The buildability specifications also remap post-1993 court splits to their pre-reform parent High Courts, which is why the number of High Courts falls from the raw 20–21 labels to 17–18 in the preferred buildability design and to 15–16 in the accessibility-adjusted robustness design.

Two features of the outcome distributions deserve emphasis. First, the GHSL outcomes are stock outcomes, not direct administrative measures of newly permitted construction. Built surface and built volume therefore measure accumulated physical capital, while built height is an average-intensity object that can move differently from total built mass. Second, the nightlights outcomes live on a very different scale from the built-environment variables because they measure visible luminosity rather than structures.

Appendix Table 6A reports the observation, district, and High Court support behind each economic design. The raw Economic Census panel covers 640 districts in 20 High Courts, while the raw GHSL and nightlights panels cover 641 districts and 21 High Court labels. The backlog design is available only where a district can be linked to the pre-reform backlog inventory. The preferred buildability design additionally requires district-level geophysical support and a 1990 GHSL built-volume baseline. The appendix robustness buildability design further requires non-missing accessibility. These requirements are why the district and High Court counts shrink from the raw panels to the estimation samples.

Backlog Exposure. Backlog Exposure is the main district heterogeneity term in Table 4. For each district, the project counts the number of pending state-involving property-rights cases linked to that district immediately before October 6, 1993 and standardizes that count

within the district's parent High Court. Panel A of Appendix Table 6A reports the raw backlog count, the district's share of the parent High Court's backlog, and the standardized within-parent-High-Court z-score used in the regressions. Within-court standardization matters because the main economic design compares districts that were more versus less exposed to unresolved state-property conflict within the same High Court rather than across courts of very different scale.

Buildability Exposure. Buildability Exposure is the alternative district heterogeneity term in Table 5. The preferred version is the standardized product of a physical buildability index and an underbuilt-slack term. Physical buildability is a district-level geophysical feasibility index. It averages standardized terrain and hydrology components, with signs oriented so that higher values mean easier broad construction feasibility: mean slope, median slope, the share of district area with slope below five degrees, the share above fifteen degrees, local terrain ruggedness, elevation dispersion, elevation range, permanent-water share, wetland share, floodplain share, and major-river surface share. Preferred underbuilt slack is the negative standardized residual from a 1990 cross-sectional regression of district built volume on physical buildability alone. Intuitively, this object is high in districts that were physically feasible for development but still had unusually low built stock at baseline.

Why the sample counts differ. The support differences in Appendix Table 6A are not reporting accidents. They reflect the underlying design. The backlog design follows the raw outcome panels as long as backlog exposure is available. The buildability designs require a 1990 built-volume baseline because underbuilt slack is defined off that baseline. The appendix robustness version further requires the accessibility variable used in the residual-slack construction. Finally, the buildability specifications remap later court splits to their pre-reform parent High Courts, which is why the number of High Courts falls from 20–21 in the raw outcome panels to 17–18 in the preferred buildability design and to 15–16 in the appendix robustness design.

Pre-reform regional and economic balance. The main economic interpretation does not require High Courts above and below the median of eventual retirement exposure to be identical in every respect, but it does require that the court-level treatment not simply proxy for obvious pre-reform regional differences. Appendix Figure 2A therefore compares court jurisdictions on baseline population, urban population, rural population, employment, backlog concentration, and district-level conflict salience before the reform. The unit of observation is the High Court jurisdiction. The plotted coefficients are standardized mean differences, and none remains statistically significant at the 10 percent level after Holm correction within this family of tests.

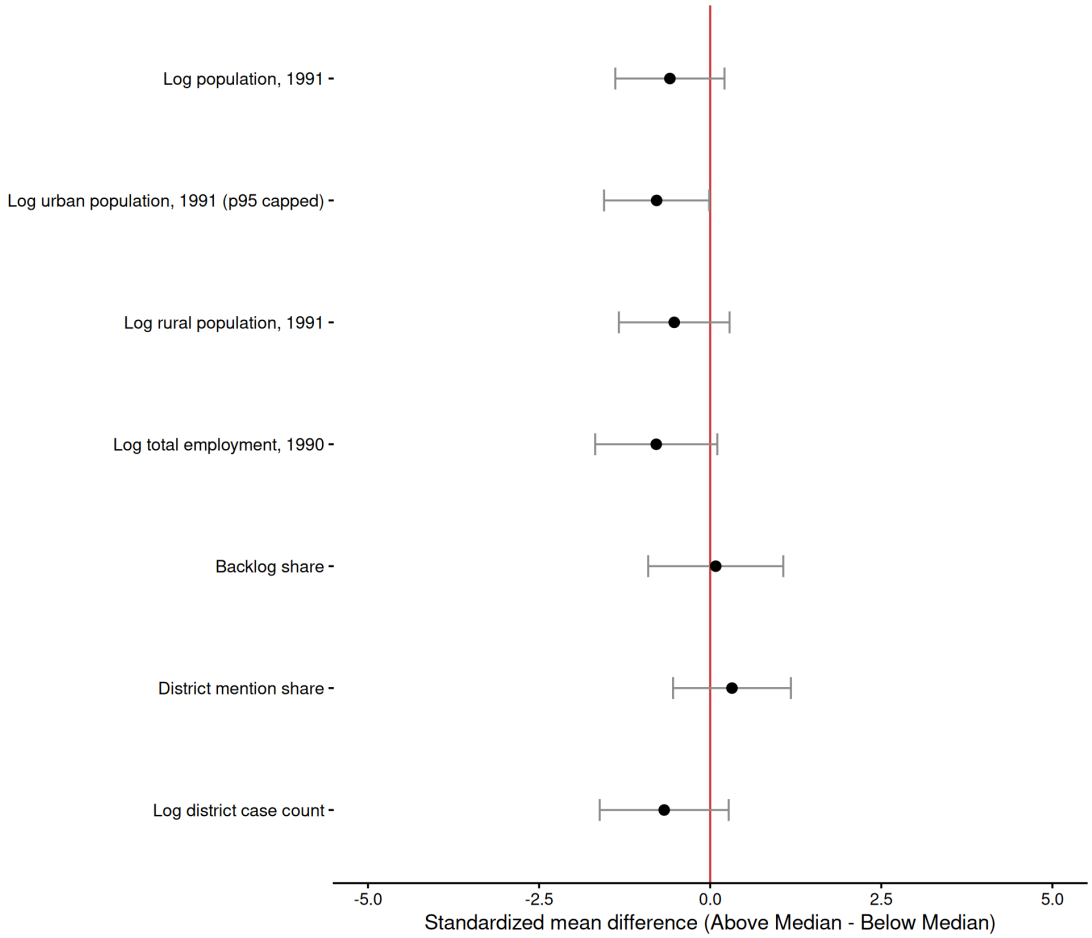


Figure 2A: Pre-Reform Regional and Economic Balance

Notes: The figure compares High Courts above and below the median of eventual retirement exposure, using the same binary split based on whether the retirement-turnover ratio reaches one by 2000. The unit of observation is the High Court jurisdiction. The plotted variables are court-jurisdiction averages of log population, log urban population with the urban-population margin capped at the 95th percentile before log transformation, log rural population, log total employment, backlog share, district mention share, and log district case count. Coefficients are standardized mean differences, and vertical bars denote 95% confidence intervals. No displayed coefficient remains statistically significant at the 10 percent level after Holm correction within this family of tests.

B.2 Satellite Built-Environment and Nightlights Methodology

The district-level economic analysis uses two satellite-based outcome families. The first is the built-environment data from the Global Human Settlement Layer (GHSL R2023A). The second is the luminosity data from DMSP-OLS stable nightlights. This subsection explains how these data are constructed and what the district-level variables measure.

Overview. GHSL R2023A is produced by the European Commission's Joint Research Centre. The paper uses three GHSL products: built-up surface, average building height, and built volume. Built surface records how many square meters within each 100-meter grid cell are classified as built-up. Built height records average gross building height in

meters at the same 100-meter resolution. Built volume is defined mechanically as built surface multiplied by built height. In the regression analysis, these gridded measures are aggregated to common district boundaries. Figures 3A–6A reproduce the technical figures used to understand how these satellite variables are constructed and what they look like after aggregation to district boundaries.

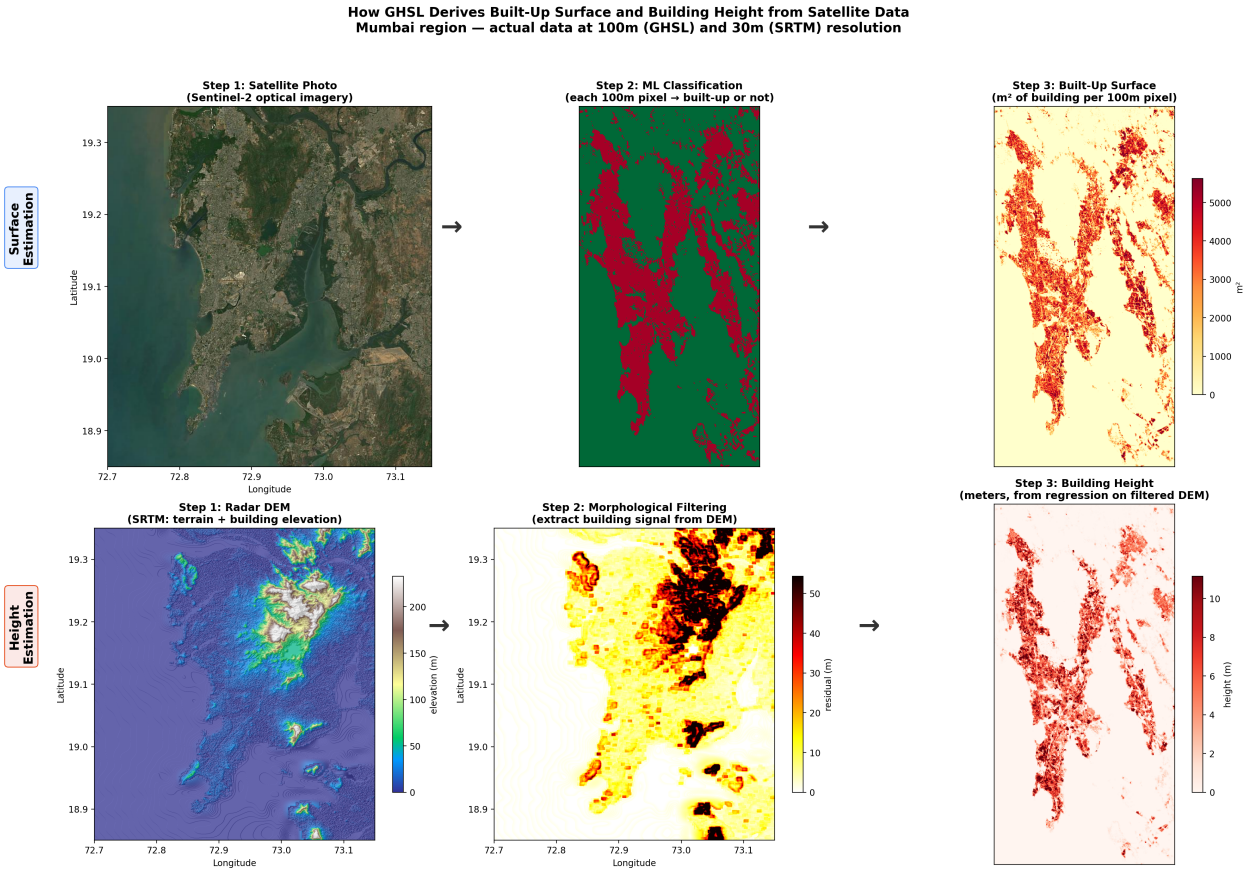


Figure 3A: GHSL Construction Pipeline for Built Surface and Built Height

Notes: The figure illustrates the GHSL construction pipeline using actual satellite data for the Mumbai region. The top row shows the optical-imagery workflow for built surface: Sentinel-style optical imagery, pixel-level built-up classification, and the resulting built-up surface measure. The bottom row shows the elevation-based workflow for built height: the radar-based elevation model, filtering steps that isolate the building signal, and the resulting building-height measure. All panels refer to the same geographic area, and the final GHSL outputs are expressed on a 100-meter grid.

GHSL built-environment measures. GHSL R2023A is produced by the European Commission's Joint Research Centre and provides gridded measures of built-up surface, building height, and built volume at 100-meter resolution. We use three GHSL products. *Built surface* measures the amount of built-up area within each 100-meter cell. *Built height* measures average gross building height in meters. *Built volume* is the product of built surface and height and therefore summarizes the total mass of built structures within the cell.

The built-surface measure is estimated from multispectral optical imagery. GHSL uses Landsat data and a supervised global classification procedure that distinguishes built-up

materials from vegetation, water, and bare soil on the basis of their spectral signatures. The output is a continuous built-up surface measure ranging from zero to the full 10,000 square meters of a 100-meter cell.

The building-height measure is estimated differently. GHSL combines digital elevation models from satellite radar and stereo imagery, principally SRTM and AW3D30, to obtain a digital surface model that includes rooftops rather than only bare terrain. It then applies multi-scale filters designed to isolate sharp local peaks and rough vertical texture associated with buildings, and uses those filtered features in a multivariate regression trained on high-resolution building-height data from a set of reference cities. The resulting quantity is average gross building height for each 100-meter cell. Built volume is then defined mechanically as built surface multiplied by built height. The description here follows the GHSL R2023A technical documentation and the validation exercises reproduced in this appendix.

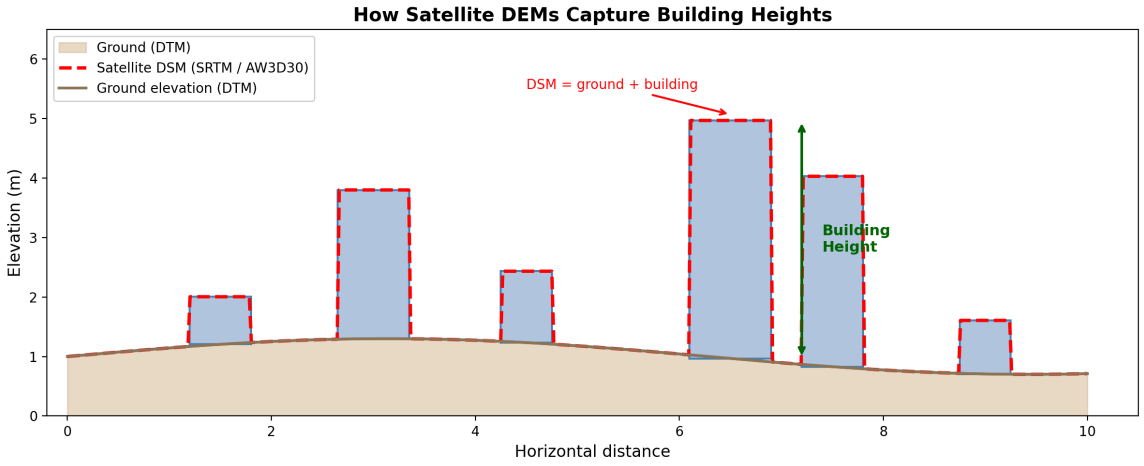


Figure 4A: How Elevation Data Translate into Building Height

Notes: The figure gives a stylized cross-section of the elevation logic behind GHSL built height. The digital surface model follows rooftops and other above-ground objects, while ground elevation follows the underlying terrain. Building height is the vertical gap between the two. In practice, GHSL approximates this gap by combining multiple elevation sources and filtering procedures rather than by observing ground and rooftop height directly for every pixel.

Stage 1: built-surface estimation from multispectral imagery. Built surface comes from optical satellite imagery. Landsat photographs the earth's surface in multiple wavelength bands, including visible light, near-infrared, and shortwave infrared. Built-up materials such as concrete, asphalt, and metal roofing have characteristic spectral signatures that differ from vegetation, water, and bare soil. GHSL uses these spectral signatures in a symbolic machine-learning classification procedure trained on global reference data. The classifier assigns each location a built-up probability and converts that information into a continuous built-up surface measure at 100-meter resolution, ranging from zero to 10,000 square meters per cell.

Stage 2: digital-surface-model construction for building height. Built height uses different underlying information. Satellite radar and optical stereo imagery measure the

height of the earth's surface, including rooftops rather than only bare ground. GHSL combines the main global elevation products used for this purpose, principally SRTM and AW3D30, to create a digital surface model that preserves the signal from buildings. Because these elevation models contain both terrain and structures, the raw signal cannot be interpreted as building height without additional processing.

Stage 3: filtering and multivariate prediction. GHSL then applies several families of filters to separate the building signal from terrain. Convolutional filters emphasize sharp local elevation changes characteristic of building edges. Morphological filters isolate compact elevated structures and their boundaries. Textural filters capture the local roughness patterns that distinguish dense clusters of buildings from smooth terrain. In the technical note, this stage generates a large bank of features summarizing the "vertical texture" of each cell. GHSL selects the most predictive features and enters them into a multivariate regression trained on external high-resolution building-height data from reference cities. The output is average gross building height at 100-meter resolution. Built volume is defined mechanically as built surface multiplied by the predicted height, so volume should be interpreted as a composite outcome rather than as an independently observed satellite product.

In our district-level panel, these cell-level GHSL quantities are aggregated to common district boundaries for ten five-year epochs from 1975 to 2020. Built surface is the district-wide total built footprint. Built volume is the district-wide total built stock measured in cubic meters. Built height is implied as district-wide built volume divided by district-wide built surface, so it should be interpreted as average gross building height among built structures rather than as a skyline measure for the entire district. This distinction matters because a district can change through horizontal spread, vertical intensification, or both.

Some specifications also use *new built surface* and *new built volume*. These are not direct administrative measures of newly permitted construction. They are epoch-to-epoch differences in the underlying GHSL stock measures. We therefore interpret them as supplementary flow proxies derived from the built stock series, not as exact counts of newly completed projects.

Interpretation and limitations of GHSL outcomes. The GHSL outcomes are well suited to the paper because they measure physical capital accumulation directly and do so consistently over long horizons. At the same time, they have limits that matter for interpretation. First, the measures are estimated from global remote-sensing algorithms rather than from administrative building records, so they should be interpreted as systematic proxies rather than exact engineering quantities. Second, the height measure is noisier than the surface measure because it is inferred from elevation data and filtering rather than directly observed. Third, changes in district-average built height can be negative even when development is expanding, because rapid outward growth in low-rise built surface can mechanically lower average height. For this reason, the paper places particular weight on the combined evidence from built volume, built surface, built height, and nightlights rather than on any single satellite outcome in isolation.

Temporal comparability also matters. The early GHSL epochs rely on older Landsat imagery, while later epochs incorporate newer optical and elevation inputs. The district

fixed effects and epoch fixed effects in the regression design absorb common level differences across periods, so identification comes from relative within-district changes rather than from comparing raw pixel values across sensor vintages. This reinforces the importance of reading the GHSL results jointly across built surface, built volume, built height, and nightlights.

The technical note also reports the benchmark accuracy of the GHSL height model. In that note, average gross building height achieves a root mean squared error of roughly 2.4 meters with an adjusted R^2 around 0.57 in the validation sample. These statistics are useful in setting expectations: GHSL height is informative and economically meaningful, but it is inherently noisier than the built-surface classification and should therefore be read together with built volume and nightlights rather than as a stand-alone measure.

Figures 5A and 6A move from the construction pipeline to the district-level objects used in the regression analysis. They aggregate the GHSL grids to 2011 Census district boundaries and show side-by-side maps of built surface, built height, and built volume. The figures illustrate why the three built-environment outcomes are not redundant. A district can have large built surface but low height because development is horizontally extensive and low-rise. Another district can have modest surface but high height because development is vertically intensive. Built volume combines these two margins and therefore comes closest to a district-wide stock measure of built capital.

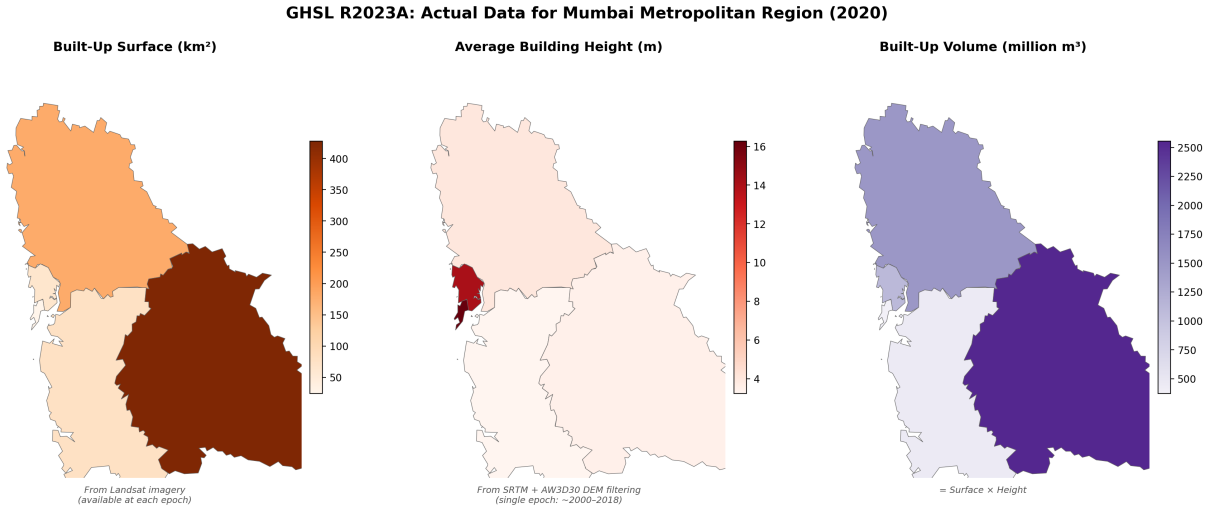


Figure 5A: Actual GHSL District Aggregates: Mumbai Metropolitan Region

Notes: The figure shows actual GHSL district aggregates for the Mumbai metropolitan region in 2020. The left panel maps built surface, the middle panel maps average building height, and the right panel maps built volume. The figure illustrates why surface, height, and volume are distinct: large sprawled districts can have very large built surface with modest height, while compact central districts can have high height but much smaller total surface. Built volume combines both margins.

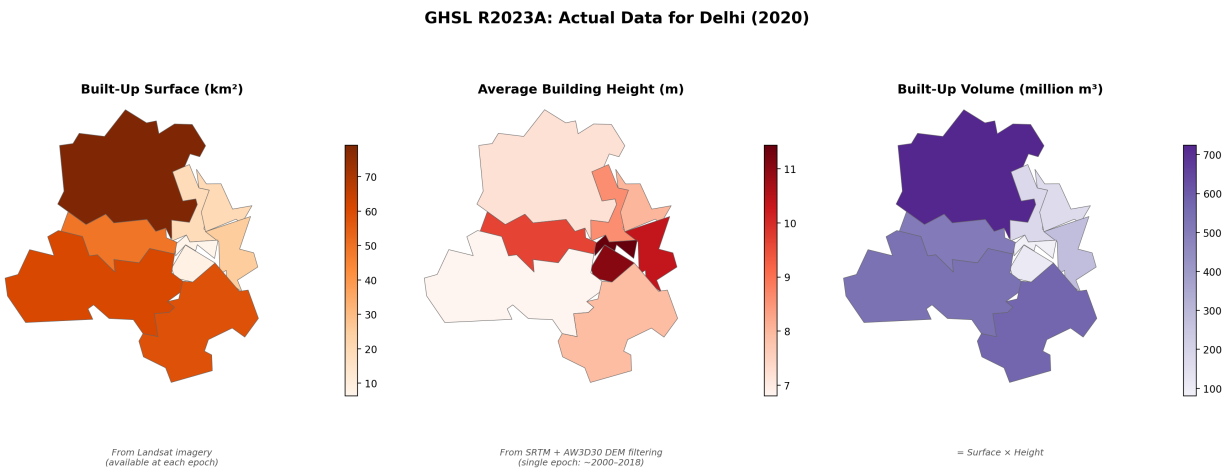


Figure 6A: Actual GHSL District Aggregates: Delhi

Notes: The figure shows actual GHSL district aggregates for Delhi in 2020. As in the Mumbai example, the left panel maps built surface, the middle panel maps average building height, and the right panel maps built volume. The Delhi example illustrates the same point in a different urban form: commercial-core districts exhibit higher height, while suburban districts exhibit larger horizontal footprints. The district-level regression outcomes are aggregates of this kind of gridded information.

Nightlights. The luminosity outcomes use DMSP-OLS stable lights, aggregated to the same district geography over 1992–2013. We report both total light and mean light. Total light is informative about aggregate visible economic activity but is mechanically larger in bigger and more spatially extensive districts. Mean light is therefore useful as a complement because it captures average luminosity intensity rather than total luminous area. In the paper, positive coefficients on nightlights are interpreted as evidence of greater visible economic activity, electrification, or commercial intensity, rather than as direct measures of construction.

Because the nightlights series begins only shortly before the 1993 reform, it provides limited pre-reform variation and is therefore best interpreted alongside the longer-horizon Economic Census and GHSL panels.

Complementary roles of built-environment and nightlights data. The GHSL and nightlights data answer related but distinct questions. GHSL measures the physical stock of built structures. Nightlights measure visible local activity. If judicial independence affects the allocation and protection of land rights, one should expect it to matter for both the built environment and for observed local economic intensity, but not necessarily to the same degree. The joint use of these outcomes therefore helps separate physical capital accumulation from broader economic activation.

For this reason, the GHSL outcomes and the nightlights outcomes receive different interpretive weight. Positive coefficients on built surface, built height, or built volume are read as evidence about construction, densification, and the physical stock of capital. Positive coefficients on total or mean nightlights are read as evidence about visible local activity, commercial intensity, and electrified economic use of space.

B.3 Alternative Economic Identification Robustness

The main text introduces a narrower alternative economic identification strategy in which district exposure is based on physical buildability and baseline underbuilding rather than on the pre-reform distribution of litigation backlog. The preferred version of this design defines

$$\text{Buildability exposure}_d = \text{Physical buildability}_d \times \text{Underbuilt slack}_d,$$

where physical buildability is the geophysical index described in Appendix B.1 and underbuilt slack captures whether the district had unusually low built volume in 1990 relative to what would be expected from that geographic feasibility. Intuitively, the preferred exposure is high in districts that were physically feasible for development but still had unusually low built stock at baseline.

Table 7A asks whether the main alternative-design result survives when that slack term is made stricter. In the appendix version, the underbuilt component is defined as a residual measure that partials out not only physical buildability but also pre-reform accessibility:

$$\text{Robustness buildability exposure}_d = \text{Physical buildability}_d \times \text{Residual underbuilt slack}_d,$$

where residual underbuilt slack is the part of low baseline built stock that remains after accounting for both geographic feasibility and accessibility. This is a useful robustness check because it asks whether the positive buildability-based pattern in the main text is really about underbuilt districts, rather than simply about districts that were both buildable and well connected.

This table reinforces the main alternative design rather than replacing it. If the preferred buildability-based specification is informative, then a stricter exposure term should preserve the same broad sign pattern even if the point estimates shrink. That is what Table 7A shows: the built-stock outcomes remain positive and precisely estimated, and mean night-lights also remain positive. The buildability-based results are therefore not driven solely by one particular way of defining underbuilt slack.

Appendix C. Supplementary Binary-Design Robustness

The main judicial analysis uses the retirement-turnover ratio as the continuous treatment. The binary designs summarize treatment onset as a discrete threshold, which supports staggered-adoption inference. The results below test whether the sign and magnitude of the effect survive under these designs and under alternative clustering rules.

Table 8A contains four variants. Column (1) uses a stricter treatment-onset threshold based on each court's 1993 baseline collegium fraction plus one pooled standard deviation of within-court annual changes. Columns (2)–(4) keep the paper's main threshold-crossing binary treatment but change the clustering rule used for inference. The rows vary the coding of Government Wins in the same way as Appendix A.2: the main coding, the 0.5 coding for partial wins, and the exclusion of partial wins.

The core question is whether the sign and rough magnitude of the binary estimates remain stable when treatment onset is defined more conservatively or when the clustering rule changes. Table 8A shows that they generally do. The continuous retirement-turnover ratio remains the preferred treatment because it uses the full path of bench replacement rather than a discrete threshold, but the binary designs move in the same direction.

Table 7A: Appendix Robustness: Alternative Economic Effects Under Accessibility-Adjusted Buildability Exposure

Panel A: Built Environment (GHSL R2023A), District × 5-Year Epoch			
	(1) Built Volume (log m^3)	(2) Built Height (m)	(3) Built Surface (log km^2)
Retirement-turnover ratio × buildability exposure	0.0175***	0.0172***	0.0151***
	(0.0053)	(0.0051)	(0.0048)
District FE	Yes	Yes	Yes
Epoch FE	Yes	Yes	Yes
Constituent design terms	Yes	Yes	Yes
Observations	6,060	6,060	6,060
Districts	606	606	606
High Courts	16	16	16
R^2	0.987	0.973	0.981
Panel B: Nightlights (DMSP-OLS), District × Year			
	(4) Total Light (log)	(5) Mean Light (log)	
Retirement-turnover ratio × buildability exposure	0.0299	0.0183**	
	(0.0214)	(0.0082)	
District FE	Yes	Yes	
Year FE	Yes	Yes	
Constituent design terms	Yes	Yes	
Observations	13,332	13,332	
Districts	606	606	
High Courts	16	16	
R^2	0.963	0.965	

Notes: Each column reports the coefficient on the interaction between the parent High Court's retirement-turnover ratio and district buildability exposure when the exposure term is rebuilt using the same physical buildability index as in Table 5, but replacing the preferred slack term with the stricter accessibility-adjusted residual underbuilding measure.

Panel definitions, samples, and fixed effects follow Table 5. All columns include district fixed effects, period fixed effects, the court-level treatment, the district exposure term, and the constituent two-way interaction terms implied by the design. Standard errors clustered at the High Court level are in parentheses.

* p<0.10, ** p<0.05, *** p<0.01.

Table 8A: Supplementary Binary-Design Robustness

	Alternative treatment onset threshold	Main treatment onset threshold; High Court clustering	Main treatment onset threshold; High Court × post-1993 clustering	Main treatment onset threshold; High Court × year clustering
Main coding (partial wins counted as government wins)	-0.0748**	-0.0278**	-0.0278*	-0.0278
	(0.0291)	(0.0124)	(0.0139)	(0.0200)
Partial wins coded as 0.5	-0.0095	-0.0417***	-0.0417**	-0.0417**
	(0.0285)	(0.0132)	(0.0160)	(0.0209)
Partial wins excluded	-0.1033***	-0.0299**	-0.0299*	-0.0299
	(0.0305)	(0.0135)	(0.0157)	(0.0211)
Case controls	Yes	Yes	Yes	Yes
Court fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations (continuous / expanded)	81,661	82,334	82,334	82,334
Observations (conservative)	75,052	75,682	75,682	75,682

Notes: Column (1) uses the alternative binary treatment defined as the 1993 baseline collegium fraction plus one pooled standard deviation of within-court annual changes. Columns (2)–(4) use the paper's main threshold-crossing binary treatment and differ only in the clustering rule. The coefficient is therefore mechanically identical across columns (2)–(4); only the standard errors change. The first row uses the paper's main outcome coding, which counts partly allowed cases as government wins. The second row assigns partly allowed cases a value of 0.5. The third row excludes partly allowed cases. All columns report the full-specification Sun and Abraham (2021) staggered-adoption average treatment effect with log judgment length and a state-petitioner indicator. Standard errors are shown in parentheses. * p<0.10, ** p<0.05, *** p<0.01.