Educating for Growth, Breeding Skepticism: Perceptions of Inequality and the State in China*

Micole De Vera† Javier Garcia-Brazales‡ Jiayi Lin§

June 2026

Abstract

Rapid economic growth underpins performance legitimacy in authoritarian regimes. Sustaining such growth is often associated with expanding education, which gives rise to the "autocrat's dilemma": equipping citizens with the human capital to critically scrutinize the state's performance narratives and the soaring inequality that accompanies rapid development. Exploiting China's staggered 1986 Compulsory Education Law, we demonstrate that increased schooling reduces tolerance for inequality and increases concerns about political corruption by around 6% of a standard deviation. We provide further empirical evidence that this operates in part through upward occupational mobility, as education moves individuals into better occupations that likely expose them to structural ceilings and institutional frictions. This, in turn, reshapes their distributional narratives: rather than merely fostering a preference for equality, education undermines meritocratic justifications for inequality and weakens support for state-led redistribution, driven by skepticism of state capacity. Finally, we provide suggestive evidence that these attitudinal shifts persist across generations.

JEL codes: D72, I25, O15, P26.

Keywords: Perceptions of inequality; Preferences for redistribution; Compulsory schooling; Instrumental variables; China.

*We thank Clara Martinez-Toledano and Carlos Sanz for valuable comments and suggestions. During the preparation of this work, the authors used Refine.ink in order to check the internal consistency of the paper. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Banco de España or the Eurosystem.

†Corresponding author: Banco de España. micole.devera@bde.es.

‡University of Exeter and LEAP (Bocconi). javier.g-brazales2@exeter.edu.uk.

§Universitat Pompeu Fabra. jiayi.lin@upf.edu.

1 Introduction

Rapidly growing authoritarian regimes often rely on sustained economic performance as a source of political legitimacy. One way to sustain this performance is through expanded public investment, particularly in human capital and education. Yet this creates a fundamental tension in the regime's perspective: while education is essential for economic development, it also equips citizens with the tools to scrutinize institutions, question official narratives, and assess distributive justice (Glaeser et al., 2007). This trade-off between fostering growth and maintaining social stability, discussed in Bourguignon and Verdier (2000), lies at the core of the so-called "autocrat's dilemma".

China provides a natural setting to study this tension. Following the 1978 Opening Up reforms, the country experienced a sharp increase in inequality, moving from one of the world's most egalitarian societies to inequality levels surpassing those of Continental Europe within a matter of decades (Piketty et al., 2019). To maintain stability during this transition, the state's ideological narrative actively justified these increasing income gaps as an instrumental necessity: a temporary but required engine for national economic prosperity (Naughton, 2007; Whyte, 2010). The stability of this model depends on individuals accepting both the legitimacy of inequality and the state's role in managing it. However, at the same time, the rapid expansion of education required to sustain growth may weaken these very beliefs.

In this paper, we study whether extending education causally affects support for the state's performance narratives—specifically, tolerance for inequality and perception of institutional integrity— in a context of a rapidly growing, highly centralized regime. To do so, we exploit quasi-experimental variation induced by the 1986 Chinese Compulsory Education Law which increased mandatory schooling from six to nine years, forcing a nationwide expansion of junior-high education. Its staggered implementation across provinces generated sharp cohort-by-province variation in exposure. We use this variation as an instrument for years of schooling, allowing us to isolate exogenous changes in educational attainment from underlying differences in economic development. We combine this strategy with data from the the China Family Panel Studies (CFPS) to estimate the causal effect of education on political attitudes.

Our empirical analysis proceeds in three steps. *First*, we show that education systematically erodes support for two core state performance narratives. In particular, one additional year of education reduces agreement with the macroeconomic claim that "the income gap needs to increase for economic prosperity" by 7% of a standard deviation. Moreover, it heightens concerns about public sector governance, increasing the perceived severity of government corruption by 5% of a standard deviation. These effects are identified from reform-induced variation in schooling across adjacent cohorts within provinces. We demonstrate that these results are robust to a number of possible concerns to internal validity, including threats to exogeneity and potential violations of the exclusion restriction.

Second, we uncover plausible mechanisms through which education reshapes these beliefs. We show that education induces upward occupational mobility, placing individuals in higher-status jobs as measured by the Standard International Occupational Prestige Scale (SIOPS), and therefore likely exposing them to new economic environments. This mobility is accompanied by a decline in the perceived importance of effort and talent in explaining economic success. Moreover, this skepticism extends to the political sphere: educated individuals become increasingly critical of state capacity, exhibiting lower satisfaction with government performance.

Turning to policy preferences, we uncover an apparent contradiction: despite their lower tolerance for inequality, treated individuals are less supportive of state-led redistribution. One additional year of education lowers the likelihood of agreeing that "more taxes should be levied on the wealthy to help the poor" by 4 percentage points—a 6% decrease from the baseline level. The institutional skepticism documented above offers a natural resolution to this paradox. Specifically, individuals are less willing to support redistribution when they doubt the state's capacity to implement it effectively. We bolster this interpretation by ruling out competing explanations, such as social desirability bias, selective emigration, or self-interested wealth preservation.

Third, we provide evidence of intergenerational persistence. The children of parents exposed to the reform exhibit systematically weaker agreement with the existence of meritocracy and the necessity of inequality for economic growth, alongside diminished trust in political cadres. This pattern suggests that education-induced shifts in political attitudes extend beyond the directly treated cohorts and propagate through family environments, thereby amplifying the long-run consequences of educational expansions.

Contributions to related literature. This paper contributes to four main strands of the literature. First, we contribute to the body of work that estimates causal effects of education on political attitudes. A large literature has documented correlations between educational attainment and civic engagement or political preferences, but credible causal evidence remains relatively scarce. Seminal contributions exploit compulsory schooling laws to identify the effects of education on political outcomes such as voting and political participation (e.g., Milligan et al., 2004), and more broadly highlight the non-pecuniary returns to human capital (e.g., Oreopoulos et al., 2006; Kenkel et al., 2006). A related literature examines how curriculum design can shape ideological beliefs and attitudes towards the state (e.g., Clots-Figueras and Masella, 2013; Cantoni et al., 2017). We complement this literature by focusing on education per se—through its effects on information processing and exposure to new economic and social environments—rather than on ideological content. Moreover, we consider a distinct but central dimension of political attitudes: tolerance for inequality and narratives that sustain it. In doing so, we show that expanding education can weaken beliefs that underpin performance-based legitimacy in authoritarian regimes. Finally, we add to a growing empirical literature that exploits China's 1986 Compulsory Education Law to identify the causal returns to schooling (e.g., Huang, 2015; Rawlings, 2015; Ma, 2019; Cui et al., 2019). While previous work leveraging this reform has focused primarily on health, cognition, and child development, we extend this identification strategy to demonstrate its deep implications for political economy.

Second, we contribute to the literature on inequality perceptions and meritocratic beliefs. The literature has identified a puzzling empirical regularity in which individuals at the lower end of the income distribution often display a surprising tolerance for income inequality (Kelly and Enns, 2010). This dynamic is remarkably pronounced in authoritarian regimes, where citizens remain broadly tolerant of distributive inequality, but exhibit frustration over procedural injustices, corruption, and institutional frictions (Whyte, 2010). Theoretical and empirical research emphasizes that such tolerance depends on beliefs about fairness, social mobility, and the functioning of economic systems (Piketty, 1995; Alesina and La Ferrara, 2005; Bénabou and Tirole, 2006; Stantcheva, 2024). We provide causal evidence on the formation of these beliefs by showing that education reduces the perceived importance of effort and talent as drivers of success. Importantly, we link this shift to individuals' lived experiences: education induces upward occupational mobility, which exposes individuals to structural constraints and weakens purely meritocratic interpretations of economic outcomes. In this sense, we identify a mechanism through which education reshapes the narratives that sustain tolerance for inequality. Moreover, we find that diminished trust in government efficiency helps explain why more educated individuals exhibit weaker support for redistribution, which is reminiscent of the findings in Kuziemko et al. (2015). A related strand of the literature studies the interaction between economic inequality and political cleavages. Recent work shows that the evolution of wealth inequality is closely linked to changes in political alignments and class-based voting behavior across advanced democracies (e.g., Gethin et al., 2022).

Third, we speak to the literature on political trust, institutional performance, and preferences for redistribution. Existing work highlights that the relationship between education and trust in institutions is theoretically ambiguous and empirically context-dependent (e.g., Hakhverdian and Mayne, 2012; Larreguy and Marshall, 2017; Ugur-Cinar et al., 2020). In environments with weak institutions or high perceived corruption, education may lead individuals to become more critical of state capacity. Our results provide causal evidence consistent with this channel: more educated individuals have more negative assessments of state capacity and government performance. We further show that these beliefs are crucial for understanding preferences over redistribution. Despite becoming less tolerant of inequality, more educated individuals do not support redistribution through taxation. We interpret this pattern as reflecting skepticism about the state's ability to effectively implement redistributive policies, rather than self-interested opposition to redistribution or changes in reporting behavior.

Lastly, we contribute to the literature on the intergenerational transmission of political attitudes. A large body of research in political science documents that political beliefs and behaviors persist across generations through family-based socialization processes (e.g., Jennings et al., 2009). More recent work highlights the role of parental education in shaping children's political engagement and beliefs by altering the household environment and access to opportunities (e.g., Kim and Lim, 2019; Janmaat and Hoskins, 2022). We provide suggestive causal evidence that education-induced changes in attitudes extend beyond directly treated individuals. The children of exposed cohorts exhibit systematically different views on inequality and institutions, indicating that educational expansions can have persistent and amplifying effects on political attitudes.

Outline of the paper. The rest of the paper is organized as follows. Section 2 describes the institutional context of the 1986 Compulsory Education Law. Section 3 outlines our empirical strategy. Section 4 presents the CFPS data. Section 5 reports the effects of the reform on support for the state's performance narratives, and probes the robustness of our results. Section 6 offers a detailed discussion of possible mechanisms at work and provides evidence of the intergenerational transmission of the policy's effects. Section 7 concludes.

2 Institutional Context: The 1986 Chinese Compulsory Education Law

Following the 1978 Opening Up reforms and the shift toward rapid economic modernization, the Chinese central government recognized that building a highly educated labor force was essential to sustain long-term growth. Facing significantly lower levels of educational achievement than Western countries, the 1986 Chinese Compulsory Education Law attempted to equate national educational standards to those in developed countries. This represented a major turning point in China's nation-wide educational policy.

The main feature of the law was that, rather than directly establishing a minimum legal age to leave school, it focused on the number of compulsory years of education that a person should undertake. In particular, it replaced the until-then prevailing 4–6 years of compulsory education by 9 obligatory years (starting from age 6). This effectively amounts to forbidding students to drop out of the system prior to age 15. To strengthen compliance, the law also forbade the employment of school-age youth. Important to our empirical strategy, the law was not binding for individuals above 15 years of age. Those aged 15 or younger who had left school at the time of the implementation were required to return to school until that age (rather than completing nine years of schooling, which may be infeasible for individuals who had quit school at an early age).

Because economic development varied significantly across regions, the central government permitted a staggered rollout of the reform, recognizing that not all provinces had the immediate capacity to enforce the mandate upon the Law's passage.[1] While this provides valuable time variation to account for potential existing trends, it also means that the enforcement date may correlate with provincial characteristics that independently determine educational and attitudinal trends. As we discuss in detail below, our empirical strategy will be geared towards addressing potential identification concerns due to the nonrandom timing of the implementation across provinces.

3 Identification and Estimation of the Causal Effect of Education on Attitudes Towards Inequality

We are interested in identifying the causal effect of education on attitudes toward inequality and institutional trust. However, educational attainment is endogenous as individuals with higher levels of education may differ systematically from less-educated individuals along unobserved dimensions that also shape their views on inequality, such as ability, family background, or prior beliefs. A naive regression of these outcomes on educational attainment is therefore unlikely to recover a causal relationship.

To address this concern, our identification strategy takes advantage of quasi-experimental vari-

[1] Table A.1 shows the adoption dates across provinces in our sample as reported by Cui et al. (2019). The majority of provinces implemented the policy fast (18 out of the 25 provinces did so either in 1986 or 1987). The last provinces to implement the policy were Hunan and Guangxi in 1992.

ation generated by the 1986 Compulsory Education Law. In particular, the reform extended compulsory schooling to age 15, creating a sharp cohort-based discontinuity in exposure: individuals aged 15 or younger at the time of implementation were subject to the new requirement, whereas older cohorts were not. Our design leverages this by comparing adjacent cohorts within provinces around this cut-off, holding constant province-specific conditions. To capture this margin of exposure, we define a binary instrumental variable based on cohort eligibility:

$$Z_i^C = \mathbb{1}\{\text{Age at implementation} \leq 15\} \tag{1}$$

which is an indicator equal to one if individual i was aged 15 or younger at the time of reform in province $p(i)$ and zero otherwise. In our preferred specifications, we rely on this binary instrument capturing exposure to the reform.

Alternative instruments are possible. In particular, following Huang (2015) and Ma (2019), one can exploit variation in exposure intensity by using a continuous measure based on age at the time of implementation (since younger children at the time of implementation were more exposed to the treatment) and further interact this measure with pre-reform non-completion rates (capturing the idea that provinces with lower initial education levels were more strongly affected by the reform).

We focus on the binary instrument for reasons of transparency and ease of interpretation. While the more flexible alternative instrument mentioned above may, in principle, capture additional sources of heterogeneity and improve efficiency, it also introduces potential concerns. First, it may exacerbate identification issues: if provinces' adoption timing is correlated with pre-reform characteristics, which in turn determine differential trends across cohorts, interacting the reform with initial education levels could amplify biases arising from non-random implementation. Second, results may become sensitive to the parametrization of the continuous exposure measure. For these reasons, we consider the binary instrument to provide the cleanest baseline and treat alternative specifications as complementary in our robustness checks.

Reduced-form, first-stage, and structural equations. The first estimand of interest is the reduced-form effect of the reform:

$$Y_i = \alpha + \rho Z_i^C + \gamma X_i + \lambda_{p(i)} + \varepsilon_i \tag{2}$$

where Y_i is an outcome of interest (such as tolerance for inequality or perceived corruption). We include province fixed effects (λ_p) to account for local macroeconomic conditions or policies that may affect the outcome. In addition, we include a vector of predetermined individual characteristics X_i including gender, an indicator for urban residence, and parental education.

The coefficient ρ captures the causal effect of exposure to the reform on the outcome, irrespective of whether exposure translates into additional schooling for every individual. That is, we recover an intention-to-treat effect of the reform. The validity of this interpretation relies on the exogeneity of the reform exposure but does not necessarily require an exclusion restriction, as the IV approach does. We further discuss the plausibility of these assumptions below.

Next, we establish that the reform exposure generates meaningful variation in schooling:

$$\text{Education}_i = \alpha + \pi Z_i^C + \gamma X_i + \lambda_{p(i)} + u_i \tag{3}$$

where Education_i is measured in years of schooling. The coefficient π is the first-stage effect of the reform on schooling. In our setting, the institutional design predicts a discrete increase in schooling for cohorts exposed before age 15.

Finally, we estimate the causal effect of education on outcomes using two-stage least squares corresponding to the structural equation:

$$Y_i = \alpha + \beta \times \text{Education}_i + \gamma X_i + \lambda_{p(i)} + \eta_i \tag{4}$$

where we instrument for years of education using our binary reform exposure indicator, Z_i^C. The two-stage least squares (TSLS) estimand corresponds to scaling the reduced-form effect of the policy on outcomes by the first-stage effect of the policy on education. Under conditions discussed below, this recovers a local average treatment effect.

Unless otherwise specified, standard errors are clustered at province-by-year-of-birth level to match the variation from the design of the reform (Abadie et al., 2023).

Plausibility of identification assumptions. Interpretation of the instrumental variable estimate as an average causal response (ACR) of education on our outcomes of interest relies on the four standard assumptions (Angrist and Imbens, 1995):

1. **Relevance**: we require that the instrument shifts educational attainment (that is, $\pi \neq 0$). This is predicted by the reform's compulsory nature and supported by clear cohort discontinuities and dose-response patterns in education around the age-15 cut-off. Where appropriate, we assess instrument strength using weak-instrument-robust diagnostics (e.g., effective

F-statistics by Montiel Olea and Pflueger (2013)) and report whether these diagnostics exceed conventional thresholds. Moreover, we report confidence sets that are robust to weak instruments (Andrews et al., 2019).

2. **Exogeneity**: conditional on predetermined controls and province fixed effects, we require exposure to be uncorrelated with unobserved determinants of the outcomes. The key source of identifying variation is the sharp school-age cut-off embedded in the law: within a province, adjacent cohorts on either side of the age-15 threshold experience different exposure status. For confounding to explain the results, some other factor would need to change discretely at exactly the same cohort cut-off within the same province. Because implementation was staggered, a central concern is that provinces adopting the policy earlier may differ systematically from late adopters. To assess robustness, we focus on the subset of provinces that implemented it in the first wave (1986–1987) to exclude provinces that could have self-selected into late implementation. We find quantitatively and qualitatively the same results. A related concern is the existence of cohort trends that covary with the treatment timing. To address this, we run alternative specifications that focus on cohorts narrowly around the policy implementation cut-off. We find similar results though estimated with less precision as we take narrower windows.

3. **Exclusion restriction**: we require that the reform exposure affects attitudes only through schooling; that is, that there is no direct effect of the instruments on outcomes conditional on the included controls and fixed effects. A common concern is that the reform could have altered local economic conditions or institutional environments in ways that directly affect attitudes. In our context, however, the identifying variation comes from within-province cohort cut-offs: broad provincial shocks or contemporaneous policies would typically affect many cohorts similarly, making it unlikely they would generate discontinuous changes precisely at the age-15 threshold. We complement this institutional argument with several empirical validations. First, outcomes display dose-response patterns around implementation analogous to those for education that gives support that the observed shifts are driven by incremental increases in schooling rather than by unrelated discontinuous shocks at the cut-off. Second, we assess robustness to potential violations of strict exogeneity using "plausibly exogenous" sensitivity analyses that allow the instrument to have a bounded direct effect (Conley et al., 2012). We compute the corresponding breakdown point which suggest that the estimated effects would remain qualitatively unchanged even under substantial departures from the

exclusion restriction.

4. **Monotonicity**: we assume that there are no defiers; that is, that the reform does not induce any individual to obtain less education than they would have obtained absent the reform. This assumption is highly plausible because the policy raised the minimum compulsory schooling requirement and restricted early school exit, particularly for cohorts below age 15 at implementation. While compliance may be imperfect (e.g., enforcement constraints in some regions), imperfect compliance likely generates non-compliance rather than negative compliance. It is difficult to rationalize a mechanism by which the policy would systematically cause some individuals to reduce schooling relative to the counterfactual. This institutional logic, together with the observed direction of first-stage effects across specifications, supports the monotonicity assumption underlying the ACR interpretation of TSLS.

Under the above assumptions, the TSLS estimate identifies an average causal response—a weighted average of per-unit causal effects for those whose educational attainment is affected by the reform (Angrist and Imbens, 1995). By contrast, the reduced-form and first-stage effect of the reform rely on a weaker set of assumptions, requiring only the exogeneity of the reform.

4 Data

The China Family Panel Study (CFPS) is a biennial, nationally-representative, longitudinal survey launched by Peking University in 2010. It employs a multistage probability sampling procedure to cover around 15,000 households across 25 mainland provinces and its modelling was inspired by the Panel Study of Income Dynamics. A more detailed description of the dataset can be found, for instance, in Xie and Hu (2014). We also complement our analysis of plausible mechanisms with data from the Chinese General Social Survey (CGSS).

Outcomes of interest. To explore the political consequences of the "autocrat's dilemma," we focus our main analysis on two complementary measures that map directly into the regime's performance-based narratives to legitimize the system. The first pillar is the state's macroeconomic narrative: the notion that structural disparities are a necessary engine for national prosperity. The second pillar is the perception of institutional integrity: the extent to which citizens recognize bureaucratic misconduct and unfairness within the state's management. Together, these dimensions allow us to assess whether increased education erodes the narratives through which authoritarian regimes sustain political stability.

To empirically capture these dimensions, we take advantage of the CFPS survey, which elicits respondents' views on the following two items: (i) "The income gap needs to increase for economic prosperity," which is available in the 2010 wave; and (ii) "How would you rate the severity of government corruption in China?"—a question that is not available in the 2010 wave but is in the 2012 one. Higher values indicate stronger agreement, though the scales differ: item (i) uses a 1–5 categorical scale while item (ii) uses a 0–10 numerical scale. We flip the order of item (ii) so that higher values indicate that the respondents find corruption not to be very problematic. We then standardize each of these variables to have a mean of 0 and a standard deviation of 1 for the individuals unaffected by the educational reform. Additionally, we combine the two into a composite index by taking the z-score of their average, providing a unified measure of support for government performance narratives.

To explore mechanisms underlying the main results, we first use occupation information available in the 2010 wave of the CFPS. In addition, we also leverage additional questions from the same wave that capture beliefs about determinants of success, such as the role of talent and effort, as well as measures of trust in multiple collectives. We complement this analysis with data from the 2015 Chinese General Social Survey (CGSS), which provides richer measures of institutional trust, government performance, and preferences for redistribution, enabling a more detailed investigation of how education-induced changes in beliefs translate into policy preferences.[2] We provide further details on these variables when relevant.

One important data issue to discuss upfront is the validity of the self-reported political questions. This has been extensively supported in previous work using the same dataset (e.g., Chen and Yang, 2015). Among the justifications provided are the presence of ample variability in the political attitudes recorded and that these distributions are similar to the ones found in anonymous online surveys. We additionally verify that there is internal consistency in the data by confirming significant correlations for variables that should be associated. For instance, trust in institutions and the extent to which corruption is a problem have a Pearson correlation of -0.18, $p < 0.001$.

Strengths of dataset. The CFPS offers several distinct advantages for our analysis. *First*, it is a large, nationally representative sample (collected more than two decades after the 1986 reform) that allows us to estimate the long-term effects of the policy on political attitudes, while also providing a rich set of complementary measures to explore potential mechanisms. *Second*, the CFPS records respondents' province of residence at age 12, which enables us to assign treatment

[2]Modeled after the US General Social Survey, the CGSS is China's earliest large-scale, nationally representative continuous academic survey eliciting citizens' political attitudes, institutional trust, and policy preferences.

status based on the timing and location of the reform with greater precision than is typically possible in survey data (Rawlings, 2015).

Sample selection. Throughout the analysis, we focus on cohorts born between 1968 and 1989. The 1968 lower bound ensures that the older, untreated control cohorts bypassed the severe school closures of the Cultural Revolution era[3], thereby enhancing comparability with the treated cohorts. The 1989 upper bound prevents right-censoring of final educational attainment and occupational status, as all individuals were at least 21 years old at the time of the 2010 survey.

Finally, subject to the global 1968–1989 bounds, we restrict the sample within each province to an asymmetric window encompassing a maximum of 3 untreated and 10 treated cohorts surrounding the implementation date. The untreated window is restricted to 3 cohorts to maintain a balanced pre-treatment panel across all provinces, as the earliest adopters (implementing in 1986) contribute exactly three untreated cohorts born after 1968. The treated window spans 10 cohorts to fully capture individuals exposed to the mandate, down to those who were of primary-school entry age (age 6) when the reform passed. This restriction aligns with our identification strategy, which compares adjacent cohorts differentially exposed to the reform, and is conceptually analogous to selecting an optimal bandwidth in a regression discontinuity design. We show robustness to alternative choices of this bandwidth. Summary statistics for the resulting estimation sample are reported in Table A.2.

5 Main Results

5.1 Reduced-form: Individuals exposed to the reform are systematically less supportive of performance-based narratives

We begin by documenting the reduced-form effects of exposure to the compulsory education reform on political attitudes, which is informative in its own right as it captures the intention-to-treat impact of the policy without imposing additional assumptions on the underlying mechanisms. In Table 1, we report estimates of the coefficient on the binary instrument from Equation (2). Columns (1) and (2) show that exposure to the reform is associated with a reduction in tolerance for inequality and in the agreement that corruption is not a severe problem in China, respectively.

[3]While primary schools gradually reopened in the early 1970s, the education system did not transition away from Maoist curricula and return to normalized, merit-based advancement until 1978 (Deng and Treiman, 1997; Meng and Zhao, 2021). Because the 1986 Compulsory Education Law primarily affected retention at the junior-high level (up to age 15), establishing a 1968 lower bound ensures that our control group experienced their critical secondary schooling years entirely under the post-1978 normalized system.

Table 1: Reduced-Form Effect of the Reform on Support of Government's Performance Narratives

	(1) Income Gap Needed for Prosperity	(2) Low Severity Problem with Corruption	(3) Composite Index
Treated	−0.054**	−0.048*	−0.058**
	(0.027)	(0.028)	(0.025)
Observations	5,936	5,275	6,470
R-squared	0.054	0.053	0.067

Notes: The outcomes are z-scores. All regressions control for gender, parental education, urban/rural location, and province of residence at age 12 fixed effects. Column (3) uses as an outcome the composite index, which is the z-score of the average of the available outcomes from Columns (1) and (2); when a respondent is missing data for one component, the index is computed using the non-missing item to preserve sample size. Higher values indicate greater support for institutional narratives. Standard errors clustered at the province-of-residence level × year-of-birth level in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

To summarize these effects, we construct a composite index that aggregates this two attitudinal measures into a single standardized outcome. This approach addresses measurement issues in the individual variables: each individual outcome is a noisy proxy of the underlying latent construct—tolerance for inequality and institutional support—and may contain idiosyncratic variation unrelated to the reform. By averaging across outcomes, the composite index filters out this measurement noise and isolates their common component, leading to more precise estimates. Focusing on the composite index, we find that those affected by the reform report, on average, 5.8% lower agreement with the composite index, indicating that treated individuals exhibit systematically lower tolerance for inequality and weaker support for institutional narratives justifying it.

Taken together, these results provide robust evidence that exposure to the education reform led to a systematic reduction in tolerance for inequality and institutional support. Importantly, these reduced-form estimates can be interpreted causally under the assumption that exposure to the reform is exogenous. In contrast to instrumental variable estimates, this interpretation does not rely on the exclusion restriction—that is, it does not require that the reform affects attitudes exclusively through schooling—but only on the exogeneity of exposure conditional on controls. For this reason, the reduced-form results provide a transparent benchmark, capturing the total causal effect of the policy on political attitudes across all channels. In the next section, we examine the first stage and instrumental variable estimates to isolate the specific role of education in driving these effects.

Table 2: First Stage of the Reform on Years of Education

	(1)	(2)	(3)
	Years of Education		
Treated	0.789***	0.930***	0.849***
	(0.114)	(0.119)	(0.112)
Observations	5,936	5,275	6,470
R-squared	0.396	0.393	0.402
Effective F-statistic	32.925	60.801	57.293

Notes: Replication of the specifications in Table 1 with years of education as the dependent variable. We repeat the regression separately for the samples in that table's columns (1)–(3). Effective F-statistics are obtained from the Montiel Olea and Pflueger (2013) robust test for weak instruments. The 5% critical value is 37.418. Standard errors clustered at the province-of-residence level × year-of-birth level in parenthesis. *** p<0.01, ** p<0.05, * p<0.1

5.2 First-stage: Exposure to reform increases educational attainment

We next examine whether exposure to the reform translates into meaningful increases in educational attainment. Establishing this first-stage relationship is essential for two reasons. First, it verifies that the policy had its intended effect: increasing years of schooling among the affected cohorts. Second, it is a prerequisite for the instrumental variable strategy: without a strong relationship between the instrument and education, it would not be possible to use the reform to identify the causal effect of education on political attitudes.

In Table 2, we report the effect of the reform on years of education across the different samples corresponding to availability of the different outcomes in Columns (1)–(3) of Table 1. We find that individuals exposed to the reform obtain, on average, approximately 0.79–0.93 additional years of schooling relative to non-exposed cohorts, depending on the sample. In addition, we report the effective F-statistic proposed by Montiel Olea and Pflueger (2013) which comfortably exceed the relevant critical values, indicating a strong first-stage effect of the reform on years of schooling.

5.3 IV estimates: Education reduces tolerance for inequality and reshapes institutional attitudes

Taken together, the results above confirm that the reform is a relevant instrument in the sense required for instrumental variable estimation. In this subsection, we now turn to the instrumental variable estimates to isolate the causal impact of education. The IV strategy combines the two previous pieces of evidence. Intuitively, while the reduced-form captures the total effect of the reform on attitudes, the IV estimates quantify how much of this effect operates through the increase

Table 3: IV Estimates of the Effect of Education on Attitudes

	(1) Income Gap Needed for Prosperity	(2) Low Severity Problem with Corruption	(3) Composite Index
Years of Education	−0.068**	−0.051*	−0.069**
	(0.032)	(0.030)	(0.028)
AR confidence sets	[−0.118,−0.018]	[−0.061, 0.037]	[−0.113,−0.025]
Observations	5,936	5,275	6,470

Notes: All regressions control for gender, parental education, urban/rural location and provincial fixed effects. Confidence sets from inverting the Anderson and Rubin (1949) test are reported in brackets. Standard errors clustered at the province-of-residence × year-of-birth level in parenthesis. *** p<0.01, ** p<0.05, * p<0.1

in education induced by the policy.

In Table 3, we report the IV results. We find that an additional year of education causally reduces tolerance for inequality. In particular, we find that more educated individuals are less likely to view inequality as necessary for growth and more likely to perceive corruption as a serious problem. The composite index reinforces this interpretation: an additional year of education decreases the composite index by 7% of a standard deviation.

These estimates can be interpreted as an average causal response (ACR). That is, they capture a weighted average of the causal effect of each additional year of education for the subset of individuals whose schooling decisions were affected by the reform. The validity of this interpretation relies on the standard instrumental variable assumptions discussed above. The fact that the identifying variation is driven by sharp cohort cut-offs within provinces makes it unlikely that other confounding factors generate discontinuous changes in attitudes at the same threshold. The strength of the instrument was confirmed by our first-stage regressions. To complement conventional inference, we also report confidence sets obtained by inverting the Anderson and Rubin (1949) test (Andrews et al., 2019). These confidence sets are robust to weak instruments in the sense that they remain valid even in settings where the first stage is weak. The resulting intervals for the composite index and the income gap measure exclude zero, reinforcing our findings with standard inference.

Taken together, the IV results show that the reductions in support for performance-based narratives documented in the reduced-form analysis are not only driven by exposure to the reform per se, but can be causally attributed—at least in part—to the increase in educational attainment. This provides direct evidence that education plays a central role in shaping political beliefs and preferences toward inequality.

5.4 Robustness and sensitivity analyses

In this section, we perform robustness checks to assess the sensitivity of our results to potential threats to internal validity.

Exogeneity of the reform: timing of implementation. A central requirement for our empirical strategy is the exogeneity of reform exposure. This assumption plays a dual role. First, it underpins the causal interpretation of the reduced-form estimates: if exposure to the reform is independent of unobserved determinants of political attitudes, then the reduced-form coefficients can be interpreted as intention-to-treat effects. Second, it is also a key condition for the validity of the instrumental variable strategy, as it requires that the instrument is uncorrelated with unobserved factors affecting the outcomes.

While the identifying variation exploited in our setting arises from sharp cohort cut-offs within provinces, a potential concern is that the timing of implementation across provinces may not be fully random. In particular, provinces that adopted the reform later might differ systematically from those that adopted it soon after the national announcement. To ensure our results are not driven by the endogenous timing of late adopters—or by compositional shifts in the event-time panel—we re-estimate our main specifications restricting the sample to provinces that adopted the reform in the earliest wave (1986–1987). By focusing on this subset, we substantially reduce variation driven by differential timing of adoption and limit concerns that cross-province selection into early versus late implementation is driving the results. The result is reported in Table A.3 and is of similar magnitude to our main estimate.

Pre-existing cohort trends. Another concern is the potential existence of underlying cohort trends that correlate with the timing of policy implementation. In such case, the effects we estimate might capture these secular cohort trends rather than the true causal effect of the educational reform. Controlling flexibly for cohort trends is not feasible in our setting, as treatment status is mechanically linked to cohorts, leaving limited residual variation for identification. Moreover, such an approach would shift identification toward cross-province differences in implementation timing, which may be endogenous.

We address this concern with two complementary analyses. First, we report estimates from specifications that include parametric cohort controls. In particular, we augment our baseline model with linear and quadratic cohort terms to capture smooth cohort trends. Results are reported in Table A.4. As expected, including these controls substantially reduces identifying variation and renders the estimates statistically insignificant. Reassuringly, however, the point estimates remain similar in magnitude to our baseline results, suggesting that the loss of significance reflects reduced statistical power rather than a change in the underlying effect.

Second, we perform additional analyses restricting the estimation samples to narrow cohort windows around the policy cut-off. By comparing individuals close in age on either side of the implementation threshold, this approach mitigates the influence of differential cohort trends while preserving plausibly exogenous variation in exposure. In particular, we replicate in Table A.5 our baseline analyses using windows ranging from one to five years around the cutoff. The estimates remain statistically significant when narrowing the cohort window down to as few as four years around the cutoff. For even narrower windows, the reduced-form estimates attenuate—as these narrowest boundaries capture cohorts with only very partial treatment exposure—and the smaller sample size lowers precision. However, the instrumented point estimates remain broadly stable. The stability of our estimates therefore also provides reassurance that our results are not driven by the fact that cohorts born near the 1968 boundary received early primary schooling prior to the formal 1978 educational normalization.

Alternative instrument: exploiting additional cohort and spatial heterogeneity. The binary instrument we employ may be coarse in that individuals just below the age cut-off are treated identically to much younger cohorts, despite having been exposed to the policy for a substantially shorter portion of their schooling trajectory. In Figure A.1, we present reduced-form estimates of the reform on attitudes and the first stage on education, allowing effects to vary across cohorts. The results provide additional support for our baseline results using the binary instrument. First, we find no evidence towards differences in political attitudes or education among cohorts that were already at least 15 years old when the reform was implemented.[4] Second, the coefficients display a clear dose-response pattern: cohorts more intensively exposed to the reform exhibit larger reductions in the attitudinal index and larger increases in years of schooling. This pattern reinforces the interpretation that the reform operates through the educational channel and that our empirical strategy captures variation in exposure intensity rather than secular cohort trends (further supporting our analyses above).[5]

[4]Because the 1968 global cohort bound naturally truncates the pre-treatment window for early-adopting provinces, the farthest-left pre-treatment bins are disproportionately identified by later adopters. However, the inclusion of province fixed effects absorbs cohort-invariant baseline differences between early and late implementers, ensuring these bins do not merely reflect cross-sectional compositional changes.

[5]We also replicate these patterns in specifications that expand the set of cohorts on either side of the cut-off, increasing the number of pre- and post-treatment observations. We report the estimates in Figure A.2. While these extended windows yield similar qualitative patterns, we view the narrower windows as more credible, as they better isolate variation close to the cut-off where confounding cohort dynamics are less likely to operate.

Following Huang (2015), we leverage this additional heterogeneity in an alternative instrument:

$$Z_i^I = f(\text{Age distance to } 15_i) \times \text{High Bindingness}_{p(i)} \tag{5}$$

which interacts (a) years exposed to the reform with (b) the extent to which the reform likely bound at the province level. The first term captures that potential exposure to the reform depends on the age when the reform was enacted. We choose $f(\cdot)$ so that this variable maps to $[0, 1]$ such that cohorts aged 6 or below (prior to schooling age) at implementation are fully treated, cohorts age 15 and above are not treated, and intermediate ages receive a linear dosage. The second term is the province-level pre-reform share of students not completing nine years of schooling. If prior to the reform most children in a province did not receive nine years of schooling, the reform had more bindingness.[6] We replicate Table 3 in Table A.6 using this alternative instrument alongside our baseline binary indicator. The estimates remain quantitatively and qualitatively similar to the baseline estimates.

Robustness to violations of the exclusion restriction. A key concern for our instrumental variable strategy is the validity of the exclusion restriction. While the identifying variation arises from sharp cohort cut-offs, one may still worry that the reform affected political attitudes through channels other than education. Following Conley et al. (2012), we compute the breakdown point: the smallest magnitude of the direct effect that produces bounds of the treatment effect of education on the composite measure of support for performance-based narratives to include 0. We find that this breakdown point is about -0.045, which is similar in magnitude to the baseline reduced-form effect. This exercise demonstrates that our instrumental variable estimate is robust to violations of strict exogeneity, requiring an implausibly large direct effect of the reform on political attitudes to overturn our conclusions.

Placebo tests. We implement a randomization inference exercise to assess whether our results arise simply by chance. To construct the permutation distribution while respecting the dependence structure of our data, we randomly reassign the binary instrument at the province-by-cohort level—the exact level at which we cluster our standard errors. This approach breaks the correlation with the true policy timing to effectively center the null distribution at zero, while guaranteeing that all individuals within the same region and cohort share the exact same placebo exposure. For

[6] As previously discussed, we do not adopt this specification as our baseline because the continuous instrument relies on stronger functional form assumptions and introduces additional measurement choices (e.g., the scaling of exposure). In addition, it depends on province-level bindingness which may be correlated with unobserved differential trends across cohorts, potentially inducing selection concerns.

each random reassignment, we re-estimate the reduced-form and first-stage effects. We plot the distribution of placebo effects and our observed coefficients in Figure A.3. Both our reduced-form estimate and first-stage estimate lie well outside the range of placebo estimates, and the corresponding permutation-based p-values are 0.036 and 0, respectively. These results indicate that the effects we document are highly unlikely to be driven by spurious correlations or random variation in the assignment of the instrument.

Multiple hypothesis testing. Concerns about multiple hypothesis testing are limited in our setting. Our empirical analysis is guided by a common estimating equation, rather than by an extensive search over specifications or subgroups. The robustness exercises we report, including the narrow-cohort analysis, are designed to probe specific identification concerns rather than to uncover statistically significant effects. Moreover, the stability of point estimates across specifications, together with the consistency of the results across related outcomes, suggests that our findings do not arise from chance rejections. These features mitigate the risk that inference is driven by multiple testing, and support the interpretation of our results as reflecting a genuine underlying effect rather than statistical artifacts.

6 Discussions

In this section, we discuss (1) plausible underlying mechanisms driving how education shapes political attitudes, and (2) the persistence of these effects over time. First, we explore the channels through which education induced by the reform affects individuals' perceptions of meritocracy, government performance, and redistributive policies. Second, we examine intergenerational persistence, asking whether these newly formed political attitudes extend to the next generation. Together, we find that not only does education influence political preferences, but also that its effects are long-lasting and likely amplified through family dynamics.

6.1 Exploring plausible mechanisms and preference for redistribution

In this subsection, we explore plausible mechanisms driving the reduced support for performance-based narratives. We first demonstrate that the extra schooling led to treated individuals taking higher-status occupations. This upward mobility likely exposed them to structural ceilings, which may explain the skepticism regarding the existence of true meritocracy that we also document. In addition, we show that this awareness of structural barriers is paired with a distrust of state capacity

Table 4: IV Estimates of the Effect of Education on Occupational Prestige

	(1) High ISEI Score	(2) High SIOPS Score
Years of Education	0.031**	0.037***
	(0.012)	(0.014)
Effective F-statistic	40.567	40.567
AR confidence sets	[0.012,0.052]	[0.015,0.062]
Observations	4,156	4,156

Notes: Both regressions have as outcome an indicator taking the value of 1 if the respondent's prestige score ranks at the 75th percentile or above of the ISEI and SIOPS distributions, respectively. The specification replicates those in Table 3. Standard errors clustered at the province-of-residence level × year-of-birth level in parenthesis. *** p<0.01, ** p<0.05, * p<0.1

and lower satisfaction with government performance. We argue that these results combined explain why we find that more educated individuals ultimately oppose further state-led redistribution. We end this discussion by providing evidence against competing mechanisms.

6.1.1 Plausible mechanisms

Occupational upgrading. One way in which education may reshape political attitudes is by altering individuals' socioeconomic position. If the reform induced upward occupational mobility, it may have exposed individuals to different economic environments and constraints, thereby affecting their views on inequality. To assess this channel, we examine the impact of education on occupational prestige using the reform as an instrument. We study two standardized international indices of occupational status: the International Socio-Economic Index (ISEI) and the Standard International Occupational Prestige Scale (SIOPS).[7] Estimates of how education induced by our reform affects these occupational status measures are reported in Table 4.

Using both ISEI and SIOPS, we find that education significantly increases the likelihood of attaining higher-status occupations. The estimated effects indicate genuine occupational upgrading rather than mechanical correlations with education. In particular, the use of SIOPS—which, unlike ISEI, does not embed education by construction—confirms that treated individuals achieve higher-prestige jobs.

This evidence suggests that the reform placed individuals in fundamentally different socioeconomic positions. Such upward mobility may make individuals more aware of structural barriers

[7] ISEI assigns continuous scores to occupations based on their typical levels of education and income, summarizing their socio-economic position, while SIOPS ranks occupations according to their perceived social prestige based on cross-national survey evaluations.

Table 5: Effects on Determinants of Success

	(1) Talent Key for One's Success	(2) Effort Key for One's Success	(3) Social Connections Over Capabilities	(4) Talent Key for Child's Success	(5) Effort Key for Child's Success
Panel (a): Reduced-Form Regressions					
Treated	−0.095***	−0.098***	0.020	−0.128***	−0.041
	(0.026)	(0.024)	(0.032)	(0.028)	(0.028)
Observations	6,431	6,457	6,341	6,335	6,444
R-squared	0.034	0.062	0.022	0.032	0.018
Panel (b): Instrumented Regressions					
Years of education	−0.113***	−0.117***	0.024	−0.158***	−0.050
	(0.032)	(0.032)	(0.040)	(0.037)	(0.034)
Effective F-statistic	55.861	54.397	52.900	50.827	53.201
AR confidence sets	[−0.169, −0.067]	[−0.173, −0.072]	[−0.038, 0.087]	[−0.222, −0.106]	[−0.110, −0.004]
Observations	6,431	6,457	6,341	6,335	6,444

Notes: Panel (a) replicates the reduced-form specifications from Table 1, while Panel (b) replicates the instrumental variable specifications from Table 3. All outcomes are z-scores constructed from 5-point categorical variables measuring the respondent's level of agreement with the following statements: "The most important factor affecting one's future success is his/her talent" (Column 1); "The most important factor affecting one's future success is his/her effort" (Column 2); "In today's society, having social connections is more important than having individual capability" (Column 3); "The most important factor affecting a child's future success is his/her talent" (Column 4); and "The most important factor affecting a child's future success is his/her effort" (Column 5). The 5% critical value for the effective F-statistic is 37.418. Standard errors clustered at the province-of-residence level × year-of-birth level in parenthesis. *** p<0.01, ** p<0.05, * p<0.1

and the limits of meritocratic mechanisms, providing a natural channel through which education can reshape attitudes toward inequality.

Reform affects beliefs on the drivers of success. Education may affect how individuals interpret the determinants of economic success. In particular, shifts in beliefs about the role of effort and talent can inform broader views about fairness and the legitimacy of inequality.

In Table 5, we look at the effects of the reform on beliefs about the drivers of economic success. We find that exposure to the reform reduces the belief that one's own success is driven by talent and effort (Columns (1) and (2)), and similarly lowers the perceived importance of these factors for children's future success (Columns (4) and (5), although the latter is imprecisely estimated). At the same time, we do not detect a corresponding increase nor decrease in the belief that social connections matter more than capabilities in Column (3), suggesting that the effect operates through an attenuation of meritocratic attributions rather than substitution toward alternative explanations.

Combining these results with our previous exercise, a plausible mechanism is that those affected by the reform obtained more education, which opened opportunities to certain occupations that exposed them to new environments. By moving into higher-prestige occupations and interacting with a broader set of peers and institutions, treated individuals may observe first-hand the role of structural constraints, networks, and institutional frictions. This exposure can weaken purely meritocratic interpretations of economic outcomes.

Reform increases institutional skepticism. We have thus far shown that exposure to the reform led individuals into better occupations and weakened their beliefs in meritocracy. We now examine whether this disillusionment extends to the political sphere, focusing on institutional skepticism: if improved occupations expose individuals to new environments and institutional interactions, they may become more aware of frictions, constraints, or inconsistencies in how success is realized. Since detailed measures of institutional perceptions are not available in the CFPS, we rely on the 2015 Chinese General Social Survey (CGSS) to assess how exposure to the reform shapes views of government performance and effectiveness.

In Table 6, we look at the effects of the policy and of education on measures of satisfaction with the government as well as perceptions of its effectiveness. Column (1) reports results for an index of satisfaction with government performance, where we find a negative and statistically significant effect. In Column (2), we use as an outcome an index capturing beliefs about the effectiveness of government. While the reduced-form estimate is also negative and its magnitude economically meaningful, it is not statistically significant, which is consistent with the substantially lower number of observations available for these measures.[8] Overall, the results are consistent with education being associated with lower reported satisfaction with government performance and, more suggestively, weaker perceived effectiveness of the state across several domains.

6.1.2 From views to policies

Thus far, we have established two main findings. First, education reduced support for performance-based narratives, including a reduction in tolerance for inequality. Second, education induces occupational upgrading which we argue exposes individuals to new environments and experiences that reveal limitations in meritocracy and in the efficiency of government institutions. A key remaining question is whether and how these shifts in beliefs translate into support for redistributive policies. We focus on perceptions of taxation as a tool for redistribution.

Relying again on the 2015 CGSS, we examine whether individuals exposed to the reform differ in their support for redistributive taxation. We report the results in Table 6's Column (3). The results show that treated individuals are 4 percentage points less likely to support taxing the rich to redistribute to the poor, a meaningful effect relative to the baseline agreement rate of 71%.

[8] For this outcome, the first-stage is weak due to the limited number of observations, which reduces the identifying variation available. Conventional IV inference is unreliable. Moreover, the Anderson–Rubin confidence set is unbounded, reflecting the lack of a first-stage. We caution against over-interpreting the IV coefficients for this outcome.

Table 6: Effect of the Reform and Education on Views of State Capacity and Performance and Support for Redistribution

	State Capacity and Performance		
	(1) Gov. Performance Satisfaction	(2) Gov. Effectiveness	(3) Support Progressive Taxation
Panel (a): Reduced-Form Regressions			
Treated	−0.159***	−0.084	−0.031*
	(0.047)	(0.084)	(0.018)
Observations	2,147	609	2,132
R-squared	0.088	0.072	0.028
Panel (b): Instrumented Regressions			
Years of Education	−0.206***	−0.527	−0.041*
	(0.065)	(0.749)	(0.024)
Effective F-statistic:	28.328	0.462	27.676
AR confidence sets:	[−0.330,−0.115]	$(-\infty,\infty)$	[−0.082,−0.004]
Observations	2,147	609	2,132

Notes: This table utilizes data from the 2015 Chinese General Social Survey (CGSS) to investigate evaluations of state capacity, performance, and policy preferences. Panel (a) replicates the reduced-form specifications from Table 1, while Panel (b) replicates the instrumental variable specifications from Table 3. The outcome in Column (1) is the standardized average of nine binary indicators taking the value of 1 if the respondent answered "satisfied" or "very satisfied" regarding the government's performance in: providing medical services; supporting the elderly; providing basic education; safeguarding national security; combating crime; fair law enforcement; impartial administration; environmental protection; and helping the poor/maintaining social fairness. The outcome in Column (2) is the standardized average of three binary indicators taking the value of 1 if the respondent answered "usually effective" or "always effective" regarding: the people's congress detecting and correcting unreasonable budgets; regulatory authorities supervising non-transparent bidding; and audit institutions supervising financial budgets. The number of observations in Column (2) is substantially lower because these questions address highly technical dimensions of government oversight; approximately 75% of respondents reported an inability to evaluate the state on these complex metrics. The outcome in Column (3) captures support for progressive taxation, taking the value of 1 if the respondent answered "agree" or "strongly agree" to the statement that more taxes should be levied on the wealthy to help the poor — the mean value for our estimating sample is 0.712. Standard errors clustered at the province-of-residence level × year-of-birth level in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

Taken together with the previous findings, these results shed light on the relationship between education and support for redistribution. While education reduces faith in meritocratic explanations of success and increases exposure to institutional frictions, it does not lead to stronger willingness for redistribution; indeed, the opposite holds. This pattern is consistent with a mechanism in which education reshapes beliefs about how the economy and institutions function. Rather than reflecting a heightened normative commitment to fairness, among treated individuals, opposition to progressive taxation coexists with a more critical view of state effectiveness. In this sense, education appears to simultaneously weaken purely meritocratic beliefs while fostering a more nuanced and more skeptical assessment of the state as a vehicle for redistribution.

6.1.3 Ruling out alternative mechanisms

Ruling out changes in willingness to express critical views. A potential concern is that education does not actually change underlying preferences, but merely increases individuals' willingness to express critical views of public institutions. If more educated individuals feel less constrained in voicing criticism, the observed effects could reflect changes in reporting behavior rather than genuine shifts in attitudes.

To look into this channel, in Table A.7 we examine the effect of the reform on a broader set of trust measures outside the political domain.[9] Exposure to the reform leads to a statistically significant increase in trust in doctors and Americans, while we find no significant directional effects on trust in neighbors, parents, or strangers. The fact that trust in non-state groups actually increases is difficult to reconcile with a pure "reporting" channel driven by a reduced fear of expressing negative views, which would predict a generalized decline in expressed trust. Instead, the evidence is more consistent with a reconfiguration of individuals' social beliefs, whereby increased education reshapes how individuals perceive different groups rather than simply relaxing constraints on self-expression.

Ruling out social desirability biases. The pattern of responses we document is difficult to reconcile with social desirability concerns. In the context of China's economic rise, the official narrative explicitly legitimizes income inequality as a necessary condition for growth (Whyte, 2010). If educated respondents were merely aligning with state-sanctioned views or engaging in preference falsification, we would expect them to endorse the performance narratives of the state and report

[9] In Column (1), we show that exposure to the reform also leads to less trust in cadres, which supports our above narrative that the reform led to institutional skepticism. However, the magnitude is small and not statistically significant so we do not put emphasis on this interpretation. A cadre member refers to an individual holding a relevant managing position with political affiliation in the area. The Chinese term in the survey we use is *ganbu*.

greater satisfaction with government performance. We find the opposite. Treated individuals are less likely to accept the government's justification for inequality and more critical of institutional performance. Moreover, they are less likely to agree with the widely held view that hard work alone determines success. Overall, these patterns suggest that education does not simply induce conformity to dominant narratives, but instead fosters more independent and critical assessments of both economic outcomes and state capacity.

Ruling out selective migration biases. A potential concern in this setting is differential attrition through international migration. This is unlikely to be a first-order issue in our context. The migration literature consistently shows that emigration is strongly positively selected on education (Docquier and Rapoport, 2012), and evidence for China highlights the central role of tertiary education and socioeconomic status in accessing overseas opportunities (Xiang and Shen, 2009). By contrast, the compliers to the 1986 reform are individuals whose education increased from primary to junior high, well below the margin typically associated with international migration. Moreover, international migration remains small relative to the size of the affected cohorts (Poston and Wong, 2016). If anything, insofar as the most institutionally skeptical individuals are more likely to emigrate, our estimates are likely to represent a conservative lower bound.

Ruling out pocketbook concerns. A natural competing explanation for the documented aversion towards progressive taxation reflects simple economic self-interest: if the reform moved individuals into higher-paying occupations, they may oppose progressive taxation to protect own wealth. To explicitly test this hypothesis, we re-estimate the reduced-form estimate in Column (3) of Table 6 while controlling for the previous year's individual and household (also controlling for household size) income. Since income is a post-treatment variable, this specification should be interpreted strictly as an exploratory mediation check rather than as a causal estimate. We find that the point estimate from Panel (a)'s Column (3) remains virtually unchanged (-0.033) and statistically significant (p-value $= 0.062$). This result suggests that tax aversion is not mechanically driven by income gains and self-interest. Instead, it points to an independent effect of education on attitudes toward redistribution, consistent with our broader evidence that the mechanism operates through changes in beliefs about state capacity, a dynamic reminiscent of the findings in Kuziemko et al. (2015).

Table 7: Intergenerational Transmission of Attitudes

	(1)	(2)	(3)
	Attitudinal Composite Index of Child		
Father Treated	0.020		0.075
	(0.054)		(0.021)
Mother Treated		-0.115*	-0.152**
		(0.066)	(0.010)
Observations	1,443	1,443	1,443
R-squared	0.038	0.036	0.037

Notes: The outcome is constructed by taking the average of the child's standardized responses (z-scores) to the following four statements collected in the 2012 and 2014 CFPS waves, and then standardizing the resulting average: "Hard work will be rewarded in today's society," "Intelligence and wisdom will be rewarded in today's society," "Income gap should be increased in order to achieve economic boom," and "How much do you trust: Cadre." All columns control for urban/rural location as well as for the child's gender and year-of-birth fixed effects. Column (1) additionally controls for the province-at-age-12 of the father while Columns (2) and (3) control for the province-at-age-12 of the mother. Standard errors in parentheses are clustered at the province-of-residence × year-of-birth level of the father in Column (1), of the mother in Column (2), and two-way clustered by both the mother's and the father's province-of-residence × year-of-birth levels in Column (3). *** p<0.01, ** p<0.05, * p<0.1

6.2 Intergenerational correlation of attitudes

A natural follow-up question is whether changes in political attitudes induced by the reform persist beyond the directly treated cohorts. If education reshapes individuals' political beliefs, these effects may extend to the next generation through intergenerational transmission. Such transmission can occur through multiple channels, including parental socialization, changes in household environments, and the broader socioeconomic conditions in which children are raised.

Studying intergenerational correlations is therefore important for two reasons. First, it provides evidence on whether the effects of the reform are transient or have long-lasting implications for political preferences. Second, it speaks to the broader role of education as a driver of social change: if the reform not only affects the treated individuals but also influences the attitudes of their offspring, its impacts may be substantially amplified over time.

We estimate the reduced-form effect of the reform on the political attitudes of the children of the cohorts around the reform.[10] We report these estimates in Table 7. We find that the reduced-form estimates yield a negative and statistically significant coefficient when using the mother's instrument, while the corresponding estimates for the father's instrument are small in magnitude

[10]The questions available in the child (i.e., below age 15) questionnaire are not exactly the same as those in the adult questionnaire. To parallel the analyses that we conducted for the parental generation, we employ as outcomes the answers to questions that speak to the key outcomes in Tables 1, 5 and A.7.

and not statistically significant. In particular, Column (3) shows that individuals whose mother was affected by the reform display an agreement with the attitudinal index that is lower than that of the non-affected ones by 15% of a standard deviation.

More importantly for our context, these findings suggest that the political effects of the reform extend beyond the directly treated cohorts and persist through intergenerational transmission. Our results are consistent with the idea that education-induced belief changes reshape the household environment and influence the next generation's worldview. This shows that the reform's impact is not transient, but instead has the potential to generate long-lasting changes in political preferences. Moreover, it underscores the broader role of education as a driver of social change: by altering not only individual beliefs but also the attitudes transmitted within families, educational expansions can have multiplier effects over time. Finally, the asymmetry between paternal and maternal transmission aligns with a broad literature demonstrating that maternal education consistently plays a more dominant role in shaping child outcomes and household socialization (e.g., Cui et al., 2019).

7 Conclusions

In this paper, we argue that educational expansion can generate unintended political consequences by reshaping how individuals interpret inequality and the role of the state. Exploiting China's Compulsory Education Law, we show that increased schooling reduces tolerance for inequality, increases concerns for corruption, weakens meritocratic beliefs, and erodes trust in state capacity. We document that these effects are closely tied to occupational upgrading: education moves individuals into better positions that likely expose them to structural constraints and institutional frictions, thereby altering how they make sense of economic outcomes. Rather than simply fostering a preference for greater equality, education reshapes underlying distributional narratives, leading individuals to question both meritocratic justifications for inequality and the capacity of the state to effectively address it.

While the evidence is consistent with this interpretation, we view it as a plausible mechanism supported by the data rather than a definitive causal pathway. Other channels such as changes in information, social networks, or exposure to different ideological environments may also contribute. Disentangling their relative importance remains an important avenue for future research. We also provide suggestive evidence that these attitudinal shifts persist across generations, pointing to longer-run consequences that extend beyond the directly affected cohorts.

Our results carry important policy implications. Educational expansions are often viewed as a cornerstone of development policy, with well-documented benefits for earnings, health, and productivity. Our findings emphasize that the returns to education extend beyond these economic dimensions and include profound effects on political attitudes and social cohesion. Understanding these broader impacts is crucial for evaluating the long-run consequences of education policies, particularly in contexts where economic development and political stability are closely intertwined.

References

Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge (2023) "When Should You Adjust Standard Errors for Clustering?" *The Quarterly Journal of Economics*, 138 (1), 1–35, 10.1093/qje/qjac038.

Alesina, Alberto and Eliana La Ferrara (2005) "Preferences for redistribution in the land of opportunities," *Journal of Public Economics*, 89 (5), 897–931, https://doi.org/10.1016/j.jpubeco.2004.05.009.

Anderson, T. W. and Herman Rubin (1949) "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations," *The Annals of Mathematical Statistics*, 20 (1), 46–63, http://www.jstor.org/stable/2236803.

Andrews, Isaiah, James H. Stock, and Liyang Sun (2019) "Weak Instruments in Instrumental Variables Regression: Theory and Practice," *Annual Review of Economics*, 11 (Volume 11, 2019), 727–753, https://doi.org/10.1146/annurev-economics-080218-025643.

Angrist, Joshua D. and Guido W. Imbens (1995) "Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity," *Journal of the American Statistical Association*, 90 (430), 431–442, http://www.jstor.org/stable/2291054.

Bourguignon, François and Thierry Verdier (2000) "Oligarchy, democracy, inequality and growth," *Journal of Development Economics*, 62 (2), 285–313, https://doi.org/10.1016/S0304-3878(00)00086-9.

Bénabou, Roland and Jean Tirole (2006) "Belief in a Just World and Redistributive Politics," *The Quarterly Journal of Economics*, 121 (2), 699–746, 10.1162/qjec.2006.121.2.699.

Cantoni, Davide, Yuyu Chen, David Y. Yang, Noam Yuchtman, and Y. Jane Zhang (2017) "Curriculum and Ideology," *Journal of Political Economy*, 125 (2), 338–392, 10.1086/690951.

Chen, Yuyu and David Yang (2015) "Historical Traumas and the Roots of Political Distrust: Political Inference from the Great Chinese Famine," *Available at SSRN 2652587*.

Clots-Figueras, Irma and Paolo Masella (2013) "Education, Language and Identity," *The Economic Journal*, 123 (570), F332–F357, 10.1111/ecoj.12051.

Conley, Timothy G., Christian B. Hansen, and Peter E. Rossi (2012) "Plausibly Exogenous," *The Review of Economics and Statistics*, 94 (1), 260–272, 10.1162/REST_a_00139.

Cui, Ying, Hong Liu, and Liqiu Zhao (2019) "Mother's education and child development: Evidence from the compulsory school reform in China," *Journal of Comparative Economics*, 47 (3), 669–692, https://doi.org/10.1016/j.jce.2019.04.001.

Deng, Zhong and Donald J. Treiman (1997) "The Impact of the Cultural Revolution on Trends in Educational Attainment in the People's Republic of China1," *American Journal of Sociology*, 103 (2), 391–428, 10.1086/231212.

Docquier, Frédéric and Hillel Rapoport (2012) "Globalization, Brain Drain, and Development," *Journal of Economic Literature*, 50 (3), 681–730, 10.1257/jel.50.3.681.

Gethin, Amory, Clara Martínez-Toledano, and Thomas Piketty (2022) "Brahmin Left Versus Merchant Right: Changing Political Cleavages in 21 Western Democracies, 1948–2020," *The Quarterly Journal of Economics*, 137 (1), 1–48, 10.1093/qje/qjab036.

Glaeser, Edward L., Giacomo A. M. Ponzetto, and Andrei Shleifer (2007) "Why does democracy need education?" *Journal of Economic Growth*, 12 (2), 77–99, 10.1007/s10887-007-9015-1.

Hakhverdian, Armen and Quinton Mayne (2012) "Institutional Trust, Education, and Corruption: A Micro-Macro Interactive Approach," *The Journal of Politics*, 74 (3), 739–750, 10.1017/S0022381612000412.

Huang, Wei (2015) "Understanding the effects of education on health: evidence from China."

Janmaat, Jan Germen and Bryony Hoskins (2022) "The Changing Impact of Family Background on Political Engagement During Adolescence and Early Adulthood," *Social Forces*, 101 (1), 227–251, 10.1093/sf/soab112.

Jennings, M. Kent, Laura Stoker, and Jake Bowers (2009) "Politics across Generations: Family Transmission Reexamined," *The Journal of Politics*, 71 (3), 782–799, 10.1017/S0022381609090719.

Kelly, Nathan J. and Peter K. Enns (2010) "Inequality and the Dynamics of Public Opinion: The Self-Reinforcing Link Between Economic Inequality and Mass Preferences," *American Journal of Political Science*, 54 (4), 855–870, https://doi.org/10.1111/j.1540-5907.2010.00472.x.

Kenkel, Donald, Dean Lillard, and Alan Mathios (2006) "The Roles of High School Completion and GED Receipt in Smoking and Obesity," *Journal of Labor Economics*, 24 (3), 635–660, 10.1086/504277.

Kim, Hyungryeol and Euijin Lim (2019) "A cross-national study of the influence of parental education on intention to vote in early adolescence: the roles of adolescents' educational expectations and political socialization at home," *International Journal of Adolescence and Youth*, 24 (1), 85–101, 10.1080/02673843.2018.1470993.

Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva (2015) "How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments," *American Economic Review*, 105 (4), 1478–1508, 10.1257/aer.20130360.

Larreguy, Horacio and John Marshall (2017) "The Effect of Education on Civic and Political Engagement in Nonconsolidated Democracies: Evidence from Nigeria," *The Review of Economics and Statistics*, 99 (3), 387–401, 10.1162/REST_a_00633.

Ma, Mingming (2019) "Does children's education matter for parents' health and cognition? Evidence from China," *Journal of Health Economics*, 66, 222–240, https://doi.org/10.1016/j.jhealeco.2019.06.004.

Meng, Xin and Guochang Zhao (2021) “The long shadow of a large scale education interruption: The intergenerational effect,” *Labour Economics*, 71, 102008, https://doi.org/10.1016/j.labeco.2021.102008.

Milligan, Kevin, Enrico Moretti, and Philip Oreopoulos (2004) “Does education improve citizenship? Evidence from the United States and the United Kingdom,” *Journal of Public Economics*, 88 (9), 1667–1695, https://doi.org/10.1016/j.jpubeco.2003.10.005.

Montiel Olea, José Luis and Carolin Pflueger (2013) “A Robust Test for Weak Instruments,” *Journal of Business Economic Statistics*, 31 (3), 358–369, http://www.jstor.org/stable/43702731.

Naughton, Barry (2007) *The Chinese Economy: Transitions and Growth*: MIT Press.

Oreopoulos, Philip, Marianne E. Page, and Ann Huff Stevens (2006) “The Intergenerational Effects of Compulsory Schooling,” *Journal of Labor Economics*, 24 (4), 729–760, 10.1086/506484.

Piketty, Thomas (1995) “Social Mobility and Redistributive Politics,” *The Quarterly Journal of Economics*, 110 (3), 551–584, http://www.jstor.org/stable/2946692.

Piketty, Thomas, Li Yang, and Gabriel Zucman (2019) “Capital Accumulation, Private Property, and Rising Inequality in China, 1978–2015,” *American Economic Review*, 109 (7), 2469–96, 10.1257/aer.20170973.

Poston, Dudley L and Juyin Helen Wong (2016) “The Chinese diaspora: The current distribution of the overseas Chinese population,” *Chinese Journal of Sociology*, 2 (3), 348–373, 10.1177/2057150x16655077.

Rawlings, Samantha B. (2015) “Parental education and child health: Evidence from an education reform in China,” CINCH Working Paper Series 1511, Universitaet Duisburg-Essen, Competent in Competition and Health, https://ideas.repec.org/p/duh/wpaper/1511.html.

Stantcheva, Stefanie (2024) “Perceptions and preferences for redistribution,” *Oxford Open Economics*, 3 (Supplement 1), 96–100, 10.1093/ooec/odad038.

Ugur-Cinar, Meral, Kursat Cinar, and Tekin Kose (2020) “How Does Education Affect Political Trust?: An Analysis of Moderating Factors,” *Social Indicators Research*, 152 (2), 779–808, 10.1007/s11205-020-02463-z.

Whyte, Martin (2010) *Myth of the social volcano: Perceptions of inequality and distributive injustice in contemporary China*: Stanford University Press.

Xiang, Biao and Wei Shen (2009) “International student migration and social stratification in China,” *International Journal of Educational Development*, 29 (5), 513–522, https://doi.org/10.1016/j.ijedudev.2009.04.006.

Xie, Yu and Jingwei Hu (2014) “An Introduction to the China Family Panel Studies (CFPS),” *Chinese Sociological Review*, 47 (1), 3–29, 10.2753/CSA2162-0555470101.2014.11082908.

A Additional Tables and Figures

Table A.1: Time of Implementation by Province

(1) Province	(2) Implementation Year
Beijing	1986
Tianjin	1987
Hebei	1986
Shanxi	1986
Liaoning	1986
Jilin	1987
Heilongjiang	1986
Shanghai	1987
Jiangsu	1987
Zhejiang	1986
Anhui	1988
Fujian	1988
Jiangxi	1986
Shandong	1987
Henan	1987
Hubei	1987
Hunan	1992
Guangdong	1987
Guangxi	1992
Chongqing	1986
Sichuan	1986
Guizhou	1988
Yunnan	1987
Shaanxi	1988
Gansu	1991

Notes: Implementation years of the Compulsory Education Law are based on Cui et al. (2019).

Table A.2: Summary Statistics for Main Analysis Sample

Variable	**Obs.**	**Mean**	**Std. Dev.**	**Min**	**Max**
Female	5,936	0.507	0.500	0	1
Father's education (1-6)	5,936	2.197	1.133	1	6
Mother's education (1-6)	5,936	1.686	0.920	1	6
Urban residence	5,936	0.500	0.500	0	1

Notes: Descriptive statistics of the estimating sample in Table 1's Column (1). Sample consists of individuals born between 1968 and 1989 observed in the 2010 round of the CFPS who resided across 25 provinces at age 12. Education categories: 1=illiterate/semi-literate, 2=primary school, 3=junior high school, 4=senior high school, 5=three-year college, 6=four-year college and above.

Table A.3: Robustness: Early Adopters

	(1)
	Composite Index
Treated	−0.050*
	(0.028)
Observations	4,890
R-squared	0.063

Notes: Replication of the specification in Table 1's Column (3) where the sample is restricted to individuals who, at age 12, lived in provinces that implemented the reform no later than 1987. Standard errors clustered at the province of residence level × year-of-birth level in parenthesis. *** p<0.01, ** p<0.05, * p<0.1

Table A.4: Robustness: Cohort Controls

	(1)	(2)	(3)
	Composite Index		
Years of Education	-0.069**	-0.052	-0.077
	(0.028)	(0.223)	(0.493)
Cohort Controls	Original	Linear	Quadratic
Observations	6,470	6,470	6,470
R-squared	0.083	0.089	0.077

Notes: Column (1) replicates the specification in Table 3's Column (3). Columns (2) and (3) control for linear and quadratic cohorts, respectively. Standard errors clustered at the province-of-residence level × year-of-birth level in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

Table A.5: Robustness: Narrow Cohort Window

	(1)	(2)	(3)	(4)	(5)
			Composite Index		
Panel (a): Reduced Form					
Treated	−0.030	−0.031	−0.022	−0.044	−0.046*
	(0.043)	(0.033)	(0.028)	(0.026)	(0.026)
Age window	[-1,1]	[-2,2]	[-3,3]	[-4,4]	[-5,5]
Observations	1,682	2,815	3,923	4,365	4,811
R-squared	0.095	0.069	0.069	0.069	0.072
Panel (b): Instrumented Regressions					
Years of education	−0.069	−0.071	−0.047	−0.079*	−0.074*
	(0.086)	(0.072)	(0.056)	(0.046)	(0.041)
Age window	[-1,1]	[-2,2]	[-3,3]	[-3,4]	[-3,5]
Observations	1,682	2,815	3,923	4,365	4,811

Notes: Replication of Tables 1 and 3 for subsamples with different windows around the province-specific implementation year of the reform. In particular, Column (1) uses a 1-year window (i.e., individuals whose age is between one year above and one year below age 15), Column (2) does it for a two-year window, and so on. Columns (4) and (5) still enforce the three-year left-side window for the reasons explained in our sample selection discussion of Section 4. Standard errors clustered at the province of residence level × year-of-birth level in parenthesis. *** p<0.01, ** p<0.05, * p<0.1

Table A.6: Robustness: Use of Continuous Instrument

	(1) Income Gap Needed for Prosperity	(2) Low Severity Problem with Corruption	(3) Composite Index
Years of Education	−0.054**	−0.088***	−0.084***
	(0.026)	(0.027)	(0.025)
P-value Hansen J-test	0.551	0.022	0.374
Effective F-statistic:	50.759	60.609	56.992
AR confidence sets:	[−0.109,−0.004]	∅	[−0.131, −0.037]
Observations	5,936	5,275	6,470

Notes: Replication of Table 3 where, apart from using our binary instrument, we also include the alternative instrument interacting age-specific dosages with a measure of provincial bindingness, as proposed by Huang (2015). The table additionally reports the p-value of the Hansen J test of overidentifying restrictions. In the Column (2), we reject the null hypothesis in the test of overidentifying restrictions which explains why the confidence set implied by inverting the AR test is null. Standard errors clustered at the province-of-residence level × year-of-birth level in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

Table A.7: Robustness: Trust in Other Collectives

	(1)	(2)	(3)	(4)	(5)	(6)
Trust in...	Cadres	Parents	Neighbors	Americans	Strangers	Doctors
Treated	−0.020	0.042	−0.013	0.072**	0.036	0.052*
	(0.029)	(0.031)	(0.028)	(0.030)	(0.025)	(0.028)
Observations	5,432	5,438	5,444	5,306	5,438	5,439
R-squared	0.035	0.062	0.033	0.066	0.049	0.034

Notes: Replication of the specification in Table 1 where the outcomes measure the strength of trust in various collectives (z-scores after transforming the original 0–10 scale). Higher values indicate more trust. Standard errors clustered at the province-of-residence level × year-of-birth level in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

Figure A.1: Effects of the Reform on Attitudes and Education: Cohort Heterogeneity

(a) Attitudes

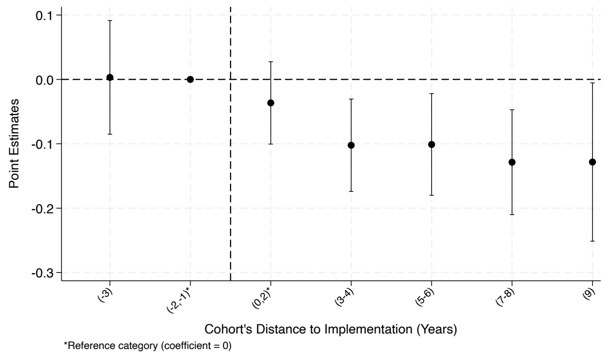

(b) Education

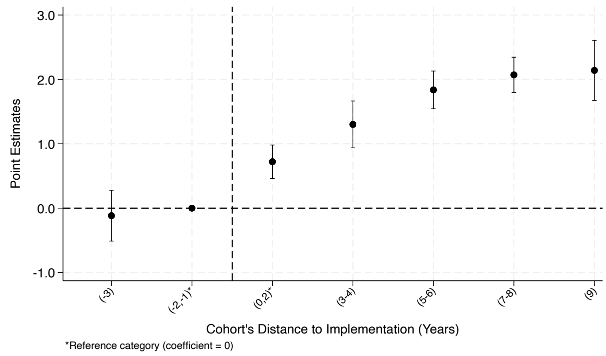

Notes: Both panels follow a similar structure: the outcome of interest is regressed on distance-to-treatment and province fixed effects. Numbers in parentheses on the x-axis indicate distance in years to age 15 at the time of the province-specific implementation. Negative numbers indicate individuals older than 15 at that time (e.g., -1 indicates individuals aged 16). The cohort bin at distance (0 to 2) comprises low-dosage, partially treated individuals who were in middle school when the law passed (aged 13-15). All dots to the left of the dashed vertical line represent never-treated individuals, while dots to the right represent treated individuals. Estimates are plotted relative to the last cohorts that were not treated at all, which serve as the omitted reference category (distance -2 to -1). Shaded areas depict 95% confidence intervals.

Figure A.2: Cohort Heterogeneity: Extended Windows

(a) Attitudes

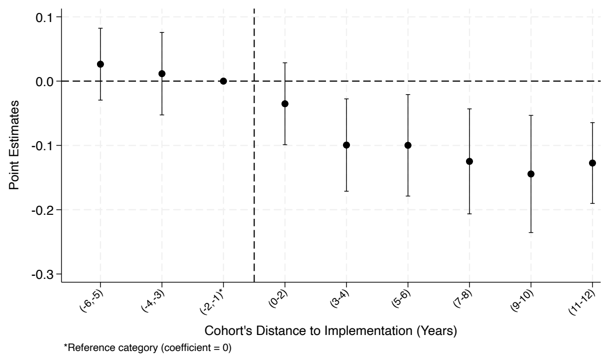

(b) Education

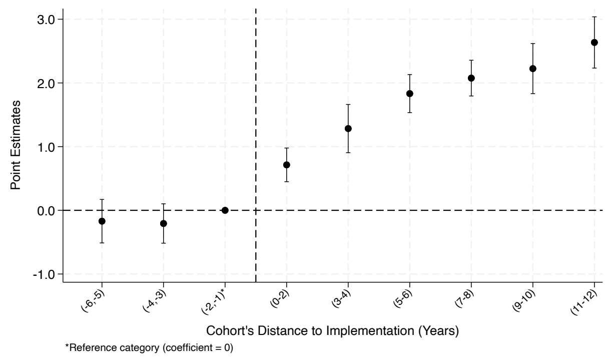

Notes: Replication of Figure A.1 extending the cohort windows around the province-specific implementation time.

Figure A.3: Robustness: Randomization Inference

(a) Reduced-form

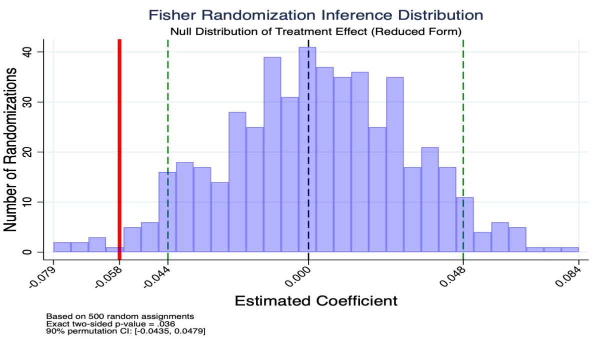

(b) First-stage

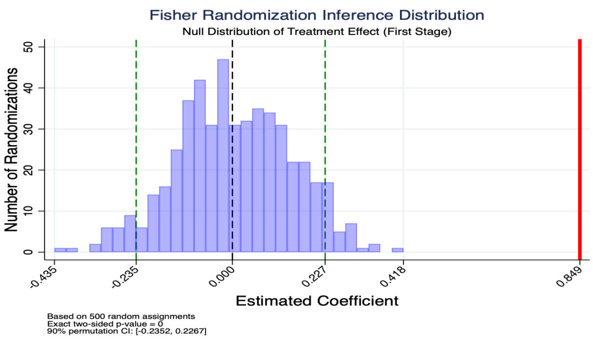

Notes: Histogram of point estimates over 500 replications of Table 1's Column (3) in panel (a) and of Table 2's Column (3) in panel (b), where binary treatment is randomly reassigned at the province-by-cohort level to preserve the original clustered structure of the data while ensuring a properly centered null distribution. Dashed green lines indicate the 95% empirical confidence interval. The red solid line indicates the original point estimate.