Social Insurance Without a Job: School-Linked Benefits and Early Formalization in Mexico[1]

Alain Pineda[2]

August 2026

Abstract

In 2016 Mexico assigned more than six million public high school and college students a permanent social security number and registered each of them individually for national social insurance (IMSS) medical coverage, a benefit tied to school enrollment rather than to a formal job. A long-standing objection to such policies is that benefits delinked from formal contracts reduce the relative value of a formal job and subsidize informality at the margin. I study the high-school component, comparing birth cohorts still in school at rollout with slightly older cohorts and letting the contrast vary with each state's pre-policy informality rate; the design recovers how the response scales with initial informality rather than a national average. Where informality was initially higher, precisely where the subsidy concern is sharpest, exposed cohorts pulled further ahead of adjacent unexposed cohorts in formal employment, and informal employment fell rather than rose. The response is concentrated entirely in wage employment and absent in self-employment. Delinking did not tilt young workers toward informality. The entitlement predated the reform, and schools had long listed their students; what 2016 changed was that coverage was attached to each student as a person, with the social security number that formal hiring requires, and made salient. Delivered that way, a benefit outside a formal contract worked as an on-ramp to formal employment rather than as a substitute for it.

[1]The views and conclusions presented in this paper are exclusively of the author and do not necessarily reflect those of Banco de México or its Board of Governors. I am grateful to Nicholas Bloom, Caroline Hoxby, Ran Abramitzky, Aldo Heffner, Alfonso Cebreros, Daniela Puggioni, and two anonymous Banco de México referees for valuable comments. All errors are my own.

[2]Banco de México, Dirección General de Investigación Económica. alain.pineda@banxico.org.mx.

1 Introduction

More than two billion workers hold jobs without legal or social protection, over 60% of the global workforce (ILO, 2018), and informality is a defining feature of labor markets in developing countries (Breza and Kaur, 2025, Ulyssea, 2020). Its costs run through poverty and job insecurity at the individual level and through tax collection and social insurance provision at the aggregate level (Bachas et al., 2023, Brockmeyer et al., 2025, Donovan et al., 2023, Jensen, 2022, Ulyssea et al., 2025). As governments extend social protection beyond formal employment, a long-standing objection has gained force: benefits that do not require a formal contract reduce the relative value of a formal job and subsidize informality at the margin (Levy, 2008). This paper asks whether that objection is borne out for a policy that attached social insurance to school enrollment rather than to employment, and finds that it is not, for a reason the delinking debate has largely overlooked: what mattered was not that the benefit was delinked from a formal contract but how it was delivered, through the contributory system's own registry and before the first job.

In 2016, Mexico registered more than six million students in public high schools and universities individually for the medical coverage of the national social insurance system (IMSS), independently of their employment status, and assigned or matched each of them to a permanent social security number, the administrative prerequisite for formal hiring. The coverage entitlement was not new: public-school students had been eligible since 1998, and schools registered them in bulk under a conventional number; what 2016 changed was that coverage was attached to a person, with a number the student keeps for life, and a national campaign told students so. Using the quarterly labor force survey (ENOE), I evaluate how access to these benefits shaped early-career labor market outcomes, at a stage where informal spells persist and scar later employment and earnings (Akay and Khamis, 2012, Cruces et al., 2012), in a country where roughly 60% of employment is informal and state-level rates range from 40% to 80%.[1]

Using a cohort-by-state difference-in-differences design, I compare birth cohorts that were still in school in 2016 (and therefore eligible for school-linked IMSS benefits) with slightly older adjacent cohorts serving as controls, and interact exposure with pre-policy state informality. Two features of the design should be stated at the outset. First, because eligibility was rolled out nationwide within a short window, there is no untreated compar-

[1] National informality fell modestly from 57% in 2016 to 55% in 2025. Informality at the state level ranges roughly from 40% to 80%. Source: INEGI.

ison group of the same age and no staggered geographic adoption; identification comes from interacting cohort exposure with pre-existing state conditions, and the estimates should be read as intention-to-treat differentials. Second, what that recovers is how the policy's effect scales with initial informality, not a national average: any component common across states is absorbed by cohort and time effects. Intuitively, the design asks whether the exposure gap between adjacent cohorts is larger where informality was initially higher, net of cohort, state, and time effects. The negative answer to the delinking concern is therefore a statement about the informality gradient of the response, that informal employment fell relative to less informal states precisely where the substitution concern is sharpest, rather than about the national level of informal employment.

Where informality was initially higher, exposed cohorts pulled further ahead of adjacent unexposed cohorts in participation and in formal employment than they did in less informal states. Each percentage point of initial informality is associated with a 0.8 to 1.3 percentage-point wider cohort gap in the formality rate. Decomposing labor market status into exhaustive shares of the cohort population shows where the response comes from: employment absorbs the entire participation response, formal employment absorbs more than the full increase, informal employment declines, and unemployment does not move. The formalization response therefore combines entry into formal jobs with reallocation out of informal ones, and it is concentrated in wage employment with no response of self-employment. Formality is the more robust of the two headline margins; the participation response attenuates under specifications that let the informality gradient vary smoothly with cohort and age.

The identifying assumption survives direct scrutiny. Cohort-specific informality gradients sit in a narrow band among cohorts too old to be exposed and break only at the eligibility ramp; placebo versions of both designs run on untreated cohorts yield null or opposite-signed estimates; and a formal sensitivity analysis following Rambachan and Roth (2023) shows the formality result survives violations of parallel cohort profiles larger than the largest departure observed among untreated cohorts. Effects are similar for women and men, for students and non-students, and for youths in households with formal and informal heads, and are robust to excluding the self-employed. ENILEMS, a survey of high-school graduates that records whether each respondent attended a public school and therefore permits a sharper exposure assignment, independently corroborates the participation response.

The pattern of results, effects concentrated in formal wage employment, absent in self-employment, and present among youth no longer enrolled in school, is most consistent with a salience and administrative-frictions channel: enrollment made benefits concrete and assigned the social security number that formal hiring requires. Delinking did not tilt young workers toward informality where informality was highest to begin with; it strengthened early formal attachment. Attaching the administrative prerequisite for formal hiring to a school-linked benefit converted a potential informality subsidy into an instrument of early formalization.

This paper speaks most directly to the literature on social protection that is not conditioned on formal employment. In high-income countries, means-tested transfers and income support generally reduce work incentives (Meghir and Phillips, 2008, Moffitt, 2002), and the Affordable Care Act's extension of dependent coverage to age 26 reduced hours and full-time work among young adults (Akosa Antwi et al., 2013). In Latin America, conditional cash transfers have small and typically insignificant labor supply effects (Alzúa et al., 2013), while an expansion of *Bolsa Família* in Brazil raised local formal employment (Gerard et al., 2021). In Mexico, the debate has centered on *Seguro Popular*: Levy (2007, 2008) argues that it subsidized informal employment, and Bosch and Campos-Vazquez (2014) find a reduction in formal registration, whereas most subsequent work finds no significant impact on informality (Alonso-Ortiz and Leal, 2018, Arias et al., 2010, Azuara and Marinescu, 2013, Campos-Vázquez and Knox, 2013, Esquivel and Ordaz-Díaz, 2009), including Seira et al. (2023) using administrative data. This paper studies a benefit that is likewise delinked from employment but delivered at labor market entry, and shows that its effect on formality depended on what the benefit was bundled with rather than on the delinking itself.

It also relates to work on the persistence of early labor market experiences. Youth unemployment spells scar later employment and earnings (Arulampalam, 2001, Bell and Blanchflower, 2011, Schwandt and von Wachter, 2019); in Latin America, early informal spells have similar persistent effects (Cruces et al., 2012), with gender differences consistent with an added-worker response (Berniell et al., 2023, Guner et al., 2025). Because a non-trivial share of high-school students combine work and study, the evidence on working while in school is also relevant: experimental evidence from Uruguay finds positive income effects two years out (Le Barbanchon et al., 2023), against mixed non-experimental findings (Ashworth et al., 2021, Hotz et al., 2002, Ruhm, 1997).

The remainder of the paper is organized as follows. Section 2 describes the policy and sets out the channels through which it could affect early labor market outcomes and the predictions that distinguish them. Section 3 describes the data and research design, presents the long-difference and panel estimates, and tests the identifying assumption. Section 4 examines mechanisms and robustness: school enrollment, heterogeneity by gender and household formality, an alternative definition of informality, the measurement of exposure and the partial treatment of the comparison cohorts, and corroborating evidence from a survey of high-school graduates, before drawing the channels together. Section 5 concludes.

2 Setting and Predictions

2.1 The Policy

Students enrolled in public upper-secondary and higher-education institutions have been entitled to IMSS medical benefits, independently of their employment status, since a presidential agreement of 1987, placed on its current legal basis by decree in 1998.[2] Before 2016, however, coverage was attached to institutions rather than to persons. Schools registered their students in bulk under a conventional, school-level social security number, and a student who wished to use the services had to apply individually, obtain a number, and register at a clinic. IMSS registers listed 7.4 million student affiliations at the end of 2015, of the order of the eligible population, but these were roster entries without a personal number or a date of birth.[3] In October 2015 the IMSS Technical Council made public institutions responsible for enrolling and de-enrolling their students electronically, each under an ordinary social security number, and in 2016 IMSS generated and assigned ordinary, unique, permanent social security numbers to students, the number required to hold a formal job in Mexico, and re-registered them individually.[4] Students obtained free access to medical consultations, lab tests, medicines, certain surgical interventions, pregnancy support, and guidance on nutrition, addictions, and sexual education.

[2] *Diario Oficial de la Federación*, June 10, 1987, and September 14, 1998. The decree covers students of public institutions who lack equivalent protection from IMSS or another social security institution; the federal government pays a fixed daily contribution per registered student.

[3] IMSS, *Informe al Ejecutivo Federal y al Congreso de la Unión sobre la Situación Financiera y los Riesgos del IMSS 2019–2020*; IMSS, *Glosario de datos abiertos, asegurados*. The count fell to 6.8 million in 2016 as records were regularized.

[4] Acuerdo ACDO.SA1.HCT.281015/246.P.DIR, *Diario Oficial de la Federación*, December 16, 2015; IMSS, *Glosario de datos abiertos, asegurados*. The social security number in Mexico is not used as the main identification number, but it is required for formal employment and remains fixed over the life cycle.

In 2016, the federal and state governments launched the "Tienes IMSS" (You Have IMSS) campaign to register students under the new procedure and promote benefit take-up. There was a widespread media campaign and big political events with the presence of the president, cabinet members, and governors of the states announcing the enrollment progress. Between March and August 2016, over six million students were registered, effectively covering the universe of students in public high schools and colleges. Although administrative efforts were coordinated by states and thus staggered, the window from launch to near-universal coverage was about five months. Given the quarterly frequency of ENOE, this rapid nationwide rollout leaves insufficient cross-state timing variation to exploit geographic staggering. The empirical strategy therefore relies on cross-cohort exposure interacted with pre-policy state conditions rather than a direct treated-versus-untreated comparison. Figure A.1 in the Appendix shows how enrollment into the program evolved. Section B in the Appendix includes additional details about the operation of the program.

In the Mexican context, the core distinction between formal and informal employment is access to social insurance (principally healthcare and retirement contributions). Survey evidence indicates that health coverage is the most valued component of formal jobs.[5] Public-school enrollment rates are high: the coverage rate of the population in high-school age is 74% and roughly 81% of high-school students attend public institutions.[6] Thus, this policy affected a substantial portion of Mexican youth and provides a setting to study how early, school-linked access to health insurance (decoupled from holding a formal job) shapes early-career labor market outcomes.

2.2 Conceptual Framework and Predictions

Ex-ante, this policy has an uncertain impact on the early labor market outcomes of the exposed population. I focus on three channels:

1. **Schooling margin (stay in school longer).**
   By attaching medical benefits to current enrollment in school, the policy raises the return to staying in school. If this dominates, exposed students delay labor market entry: lower participation while enrolled, possibly higher enrollment, and ambiguous effects on formality conditional on working.

[5] *Módulo de Trayectorias Laborales (MOTRAL) 2015.*
[6] Secretaría de Educación Pública. *Principales Cifras del Sistema Educativo Nacional 2015–2016.*

2. **Benefit substitution while enrolled (relative price of informality).**
While most high-school students are full-time, a non-negligible minority works alongside school. Using ENOE, I document that among individuals of typical high-school age (15–18) that report going to school, the average labor force participation rate over 2015–2019 is about 14%. Conditional on being part of the labor force, 95% are employed. Consistently, ENILEMS 2019 reports that 26% of recent high-school graduates had some work experience before finishing school. These magnitudes establish that the "study-and-work" margin is empirically relevant. Accordingly, the student IMSS expansion can plausibly reduce the incremental value of formal jobs' benefits while enrolled.

   Formal jobs typically bundle non-wage benefits B^F (IMSS coverage) with wages w^F. Informal jobs pay w^I with $B^I \approx 0$. Let $b^S > 0$ denote the value of school-linked benefits while enrolled. For students, the formal premium becomes:

$$(w^F - w^I) + (B^F - b^S).$$

   If b^S appreciably closes the benefit gap $(B^F - b^S)$, the relative attractiveness of informal work rises. Among the currently enrolled, we could expect a reallocation toward informal employment. Effects should dissipate after leaving school (when $b^S = 0$). Wage effects are a priori ambiguous, as employers may substitute wages for benefits. This is precisely the concern raised by policymakers and international institutions when expanding social protection outside formal employment: delinking benefits from contracts may unintentionally subsidize informality at the margin.

3. **Salience, rights, and reduced frictions toward formality during the school-to-work transition.**
The program made IMSS benefits concrete (healthcare) and assigned a permanent social security number, lowering informational and administrative frictions to formal hiring. If students learn and value these benefits, they may seek formal jobs to retain them after school, and employers face lower onboarding frictions. We could expect higher formal employment at exit (through both entry into the labor force and reallocation away from informal jobs) concentrated in wage employment rather than self-employment, with no rise in informality, and stronger effects where pre-policy informality was high.

Given these opposing forces, the net effect is empirical. For my conceptual framework, it is important to consider the *intensive margin* of informality. My setting highlights the worker's choice between formal and informal jobs at labor market entry. School-linked benefits change two features during enrollment: a benefit substitution component (some medical coverage is available regardless of formal status), which could lower the relative cost of informal work while in school; and a salience/administrative-frictions component (benefits become concrete and a permanent social security number is assigned), which can ease formal hiring at exit. Consistent with the relevance of the intensive margin in Mexico, formal firms post both formal and informal vacancies (Samaniego de la Parra and Fernández Bujanda, 2024).

Accordingly, I test three predictions: (1) effects confined to currently enrolled youth, dissipating at school exit, would indicate benefit substitution; (2) formal-sector entry concentrated at the school-to-work transition (employment expanding through formal wage jobs, with no rise in informal employment) is consistent with salience and reduced frictions; and (3) results should be stronger where pre-policy informality was higher. The evidence below aligns with (2) and (3): total employment among exposed cohorts expands entirely through formal wage employment while informal employment declines, the opposite of what net substitution toward informality would produce, and the effects are present among youth no longer enrolled in school rather than dissipating at school exit. One heterogeneity pattern is weakly consistent with the substitution margin operating for students who combine work and school, without generating a net shift toward informality.

3 Impact on Labor Market Outcomes

3.1 Data and Empirical Strategy

The main data source for this paper is the Mexican quarterly survey on employment called ENOE (*Encuesta Nacional de Ocupación y Empleo*).[7] This survey has information on approximately 200 thousand workers each quarter, is representative at the national and state level, and allows me to have information for both formal and informal workers. The design of the survey is a rotating panel, such that each individual is followed for 5 quarters, and each quarter, one fifth of the sample is replaced. Crucially, ENOE

[7] Due to the COVID-19 pandemic there is no information for the second quarter of 2020. For the period between the third quarter of 2020 and the fourth quarter of 2022 the information comes from the ENOE New Edition (ENOE^{N}). This temporary survey maintains the conceptual, statistical, and methodological design of ENOE, but combines in-person interviews with phone interviews. Throughout this period, in-person interviews represent between 79% and 100% of all interviews in ENOE^{N}.

allows me to implement INEGI's official definition of employment informality, ensuring a consistent and comparable measurement across time and states. I adopt INEGI's employment-informality concept (worker/job status) and construct outcomes using the pre-coded survey variables.[8]

The cleanest design would randomize (or stagger exogenously) access to school-linked IMSS benefits across schools or states and track named individuals longitudinally through school exit and early careers, observing school type (public/private), eligibility, take-up and subsequent entry into the labor market. In practice, the 2016 rollout reached the universe of eligible public-school students nationwide within about five months, leaving no exploitable same-age untreated group or multi-period geographic staggering at ENOE's quarterly frequency. Moreover, ENOE's short rotating panel does not support tracking individuals from school to work at scale, and it does not identify whether a respondent attended a public vs. private high school nor student-IMSS take-up. I therefore adopt a cohort-by-state difference-in-differences approach: compare birth cohorts exposed to the 2016 program (still in school at rollout) with slightly older adjacent cohorts, and interact exposure with pre-policy state informality to recover a differential intention-to-treat effect. Intuitively, I ask whether the exposure gap between adjacent cohorts is larger in states that were initially more informal. The main analysis centers on 2019, when the first set of fully exposed cohorts completed high school and before the COVID-19 shock.

I complement my main results with information from a survey that measures the labor market insertion of high-school graduates (aged 18–20), called ENILEMS (*Encuesta Nacional de Inserción Laboral de los Egresados de la Educación Media Superior*). This survey provides more detailed information about an individual's high school and labor market experiences. Crucially, I can identify whether they attended a public or private high school, making the identification of exposure to the policy cleaner. ENILEMS also provides information on the educational attainment of parents, whether the individuals attempted to, and were ever enrolled in a higher education institution, and whether the individuals were employed during high school. Unfortunately, this survey is not administered regularly. However, I can use the 2016 and 2019 surveys (the two most recent ones), that capture outcomes for cohorts both exposed to and unaffected by the policy. It is important

[8] Specifically INEGI's definition of informal employment considers the following groups: self-employed and independent workers in the informal sector (economic units operating with household's resources that are not identifiable and independent of those households), self-employed workers in the agricultural sector, workers without pay, domestic workers without social insurance, employees with wages or non-wage perceptions but no social insurance, and workers with non-specified occupation position.

to mention that ENILEMS started as a special module of ENOE and for its 2019 wave it became an independent survey but the sampling follows households identified by ENOE, making both surveys fully comparable. Given the limited number of waves and smaller sample, ENILEMS plays a corroborating role: it sharpens the exposure proxy and enables mechanism checks, but it is not the basis for identification. Two features of its design bound what it can corroborate: ENILEMS observes only individuals who completed high school, at ages 18–20, a population already substantially more formally attached than the cohorts driving the main results, and its two waves yield a single before-after contrast rather than a quarterly panel.

I focus on the high-school component of the policy (roughly U.S. grades 10–12) for two reasons. First, the mapping from birth cohorts to school grade/age is tight: most high-school students are 15–18, so cohort of birth is an informative proxy for exposure. Second, coverage at high-school ages is high (74% of the age group enrolled, 81% in public schools), whereas college enrollment in Mexico is approximately 30% with much greater age dispersion.[9] Focusing on high school therefore reduces composition noise and improves interpretability. For these reasons, the birth cohort determines eligibility exposure in my design. Analyzing the college component and longer-run outcomes is left for future work.[10]

Two further features rule out a conventional before-after design. For the exposed cohorts there are no pre-policy labor market outcomes to observe, since they were below working age before 2016; and even if school type were observed, public–private comparisons would not necessarily be credible counterfactuals.[11]

I therefore exploit the eligibility cutoff in cohort space. I measure exposure continuously as the share of grades 10–12 that coincided with the implementation of the policy. Figure 1 plots the exposure measure. Cohorts aged 18 or older in 2016 (born in 1998 or earlier) had already graduated from high school before the policy and are fully unexposed. Co-

[9]Secretaría de Educación Pública. *Principales Cifras del Sistema Educativo Nacional 2015–2016*.

[10]Because the policy also applied to public universities, some individuals in the control cohorts could have become eligible upon entering public college, and ENOE does not identify public versus private attendance. This contamination is quantitatively relevant (39% of the 1995–1998 cohorts are observed attending school during 2016–2019, see Appendix Table A.6) and it attenuates the cohort contrast in every state, so the estimates may be read as lower bounds on the combined high-school-plus-college schooling-linked impact. Section 4.4 develops the argument, states the exact condition, and quantifies the implied magnitudes per enrolled-eligible student.

[11]Arguably public and private schools differ systematically in selection, resources, and trajectories.

horts aged 15 or younger in 2016 (born in 2001 or later) entered high school after full implementation and are fully exposed. The two intermediate cohorts (born in 1999 and 2000, aged 17 and 16 in 2016) are partially exposed, having had access to benefits for one or two years of high school, respectively. The long-difference specification in Section 3.2.1 compares only the fully exposed cohorts with the fully unexposed cohorts. The panel specification in Section 3.2.2 additionally incorporates the partially exposed cohorts through the continuous exposure measure.

Because the reform's bite plausibly depends on local frictions and returns to formality, I interact this cohort exposure with pre-policy state informality. Figure 2 plots the pre-policy informality rate by state (2013–2015 average), showing substantial heterogeneity. I treat the pre-policy informality as an intensity shifter of the reduced-form effect, which I am unable to estimate in the conventional before-after design. Intuitively, the design asks whether the gap between adjacent cohorts is larger in states that were initially more informal, after absorbing cohort, state, and time effects, as well as common shocks.

Figure 1: Exposure to high-school-linked social insurance by cohort

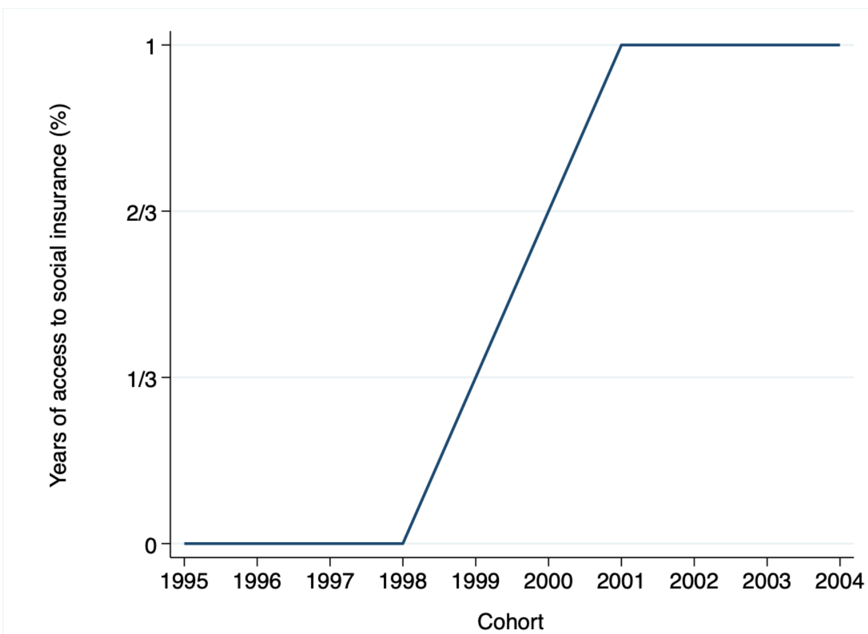

Note: The figure shows the exposure to the policy by cohort, defined as the percentage of years of high school when students had access to social insurance benefits.
Source: Author's calculation.

Figure 2: Pre-policy informality rate by state (%)
2013–2015 average

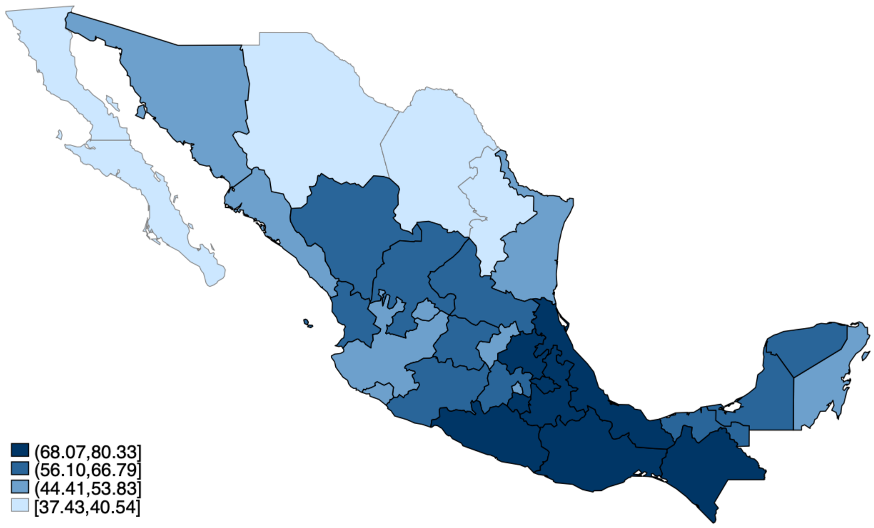

Source: Author's calculations with data from ENOE.

This strategy identifies an intention-to-treat (ITT) differential: how the policy's effect scales with initial informality, rather than an aggregate average effect or a within-age treated vs. untreated effect. This mirrors the logic in Duflo (2001) and Bleakley (2010): exploit timing by cohort and variation in pre-policy intensity, and demonstrate that only the exposed cohorts exhibit breaks aligned with exposure. Effects common across states are absorbed by cohort and time effects and are not identified here.

Outcomes are constructed at the cohort-state-quarter level. Intuitively, the difference-in-differences (DiD) design is outlined in Table 1. For simplicity, consider two state groups by pre-policy informality: low (L) and high (H). Within each group compare exposed cohorts (still in high school at rollout) to unexposed adjacent cohorts. Let Y_g^E and Y_g^U denote mean outcomes for group $g \in \{L, H\}$. Each $\Delta_g = Y_g^E - Y_g^U$ is the exposed-unexposed gap within a given state environment. The difference-in-differences estimand across state groups is $\beta = \Delta_H - \Delta_L$. This is the ITT differential: how the exposure effect varies when moving from low to high pre-policy informality. In the continuous specification used in the paper "L vs. H" becomes the slope of Δ_s with respect to $I_{s,pre}$ (pre-policy informality at the state level), and exposure is measured continuously (share of grades 10–12 under the policy).

Table 1: Cohort-by-state DiD with intensity

	Low pre-policy informality	High pre-policy informality
Exposed cohorts	Y_L^E	Y_H^E
Unexposed adjacent cohorts	Y_L^U	Y_H^U
Within-state cohort difference	$\Delta_L = Y_L^E - Y_L^U$	$\Delta_H = Y_H^E - Y_H^U$
Estimand (ITT differential): $\beta = \Delta_H - \Delta_L$		

The unexposed controls are older than exposed cohorts, so any comparison $Y^E - Y^U$ mixes (i) a life-cycle/cohort difference and (ii) the policy effect on the younger, exposed cohort. The design removes the life-cycle component by differencing across state environments. Let D_s denote the baseline cohort/age gap (unrelated to the policy) and τ_s denote the policy effect on exposed cohorts within state s. We can express the within-state cohort difference as:

$$\Delta_s = D_s + \tau_s.$$

We only observe the sum. However, taking the difference between high- and low-pre-policy informality environments:

$$\beta = \Delta_H - \Delta_L = (D_H - D_L) + (\tau_H - \tau_L).$$

If, absent the policy, the baseline cohort/age gap D_s is the same across state environments (or, more generally, uncorrelated with pre-policy informality after conditioning on fixed effects and controls), then $D_H - D_L = 0$ and the DiD estimand identifies $\tau_H - \tau_L$: the differential policy effect. Equivalently, we allow exposed cohorts to be younger than unexposed, we only require that age/cohort profile of outcomes does not systematically differ by pre-policy informality once we net out cohort, state, and time effects and standard controls. That is the parallel trends in cohort space assumption necessary for identification of the ITT differential. In other words, the identifying assumption is that, in the absence of the policy, the exposed-minus-unexposed cohort gap D_s would have been the same across states with different pre-policy informality, that is, $D_H = D_L$. Only then does $\beta = \Delta_H - \Delta_L$ recover $\tau_H - \tau_L$, the differential policy effect. After conditioning on cohort, state, and time fixed effects and controls, the cohort/age profile of outcomes must not co-vary systematically with pre-policy informality in the absence of the reform.

There are several features of my empirical design that make this assumption plausible.

First, I use slightly older adjacent cohorts as controls, minimizing life-cycle spacing. Second, cohort fixed effects absorb national age/cohort profiles; state fixed effects absorb time-invariant differences; time fixed effects absorb national shocks. I consider alternative specifications that include state-level trends to consider differentiated shocks. Third, I estimate β_k (slope of outcome vs. pre-policy informality for each birth cohort) and overlay exposure timing. Identification requires β_k flat for unexposed cohorts and breaking only when exposure is positive. Fourth, I consider placebo cohorts. For cohorts too old to be exposed, $\Delta_H - \Delta_L$ should be zero. Cohorts are tracked by state of residence rather than state of schooling, so selective migration could in principle generate a cross-state age gradient. Any such gradients should, however, appear equally among the placebo cohorts, which yield null or opposite-signed estimates, and pre-2016 cohort gaps evolve in parallel across state groups. Results are robust across gender, school attendance, formality of the household head, and focusing on workers excluding the self-employed.

3.2 Results

3.2.1 Long Differences

First, I simplify the analysis and aggregate the different cohorts into two groups: those who were fully exposed to the policy and those who were not exposed at all. I compare the main labor market outcomes (participation rate, employment rate, formality rate, monthly income) across these groups of cohorts for each state to obtain long difference estimates. Therefore, the basic equation to be estimated is:

$$Y_{s,E,t} - Y_{s,U,t} = \beta I_{s,pre} + X_{s,pre}\Gamma + \delta_t + \varepsilon_{st},$$

where $Y_{s,g,t}$ is a labor market outcome of cohort group g in state s and quarter t, and the cohort subscripts E/U denote fully exposed and fully unexposed cohort groups. In the estimation, outcomes are computed as cohort-state-quarter cell means; the equation is estimated on 32 states $\times$ 4 quarters = 128 observations, with quarter fixed effects δ_t included to absorb seasonal variation. The preceding informality rate is $I_{s,pre}$, defined as the average state informality rate between 2013 and 2015. The X variables are a set of controls measuring conditions at the state level before the implementation of the policy. The main outcomes of interest are the participation, employment, and formality rates, and monthly labor income. Throughout, the participation rate is the share of the cohort population in the labor force; the employment rate is the share of the labor force participants who are employed; and the formality rate is the share of the employed holding formal jobs.

Because the latter two are conditional rates, movements in them do not by themselves reveal how a participation response is absorbed.

For the set of controls I include state GDP per capita in 2013, mean educational attainment in 2015 and regional controls, grouping the 32 states into 8 different regions, outlined in figure A.2 of the Appendix (see Table A.1 for descriptive statistics). The equation is estimated using weighted least squares with data from the four quarters of the employment survey of 2019. Weighting proceeds in two layers. Cell-level rates and incomes are computed using ENOE expansion factors, so each cell estimates its population quantity. Regressions then weight cells by the square root of cell population, the efficient weighting when cell means are estimated with sampling variance inversely proportional to cell size. Results are virtually unchanged using population weights or no weights (Appendix Table A.2). Following Solon et al. (2015), this stability also indicates that the estimates are not driven by heterogeneity across differently sized states. The coefficient β is the ITT differential in long differences: the extra exposed–unexposed gap associated with a 1-p.p. increase in initial informality.

The long differencing estimates are presented in Table 2. Panel A shows the estimates with the preceding informality rate as the only independent variable, and Panel B includes the full set of controls. From the mean differences between cohorts fully exposed and not exposed at all, we can see that the exposed cohorts have significantly lower participation, and formality rates, as well as average income. These patterns are expected as fully exposed cohorts are younger. Because the treatment-intensity variable varies only across states, the effective number of independent units is 32 regardless of the number of individuals. I therefore cluster standard errors at the state level and, given the modest number of clusters, base inference for the coefficient of interest on wild-cluster bootstrap p-values with Webb weights and 9,999 replications (Cameron et al., 2008; Roodman et al., 2019). Appendix Table A.2 shows that the main estimates are insensitive to the inference method: cluster-robust, block-bootstrap (resampling states), cluster-jackknife (CV3; MacKinnon et al., 2023), and wild-bootstrap intervals all exclude zero by wide margins for participation and formality in both research designs, as does a specification collapsed to one observation per state.

From the estimated coefficients in panel B, one additional percentage point in the preceding informality rate is associated with a wider exposed-minus-unexposed cohort gap of 0.59 percentage points in the participation rate, 0.15 percentage points in the employment rate, and 0.87 percentage points in the formality rate. The estimated impact on labor income, while positive, is not statistically significant. The sample of income results is smaller due to higher non-response rates in this question of the survey. These reduced-form results suggest that in states with higher informality, the exposed cohorts pulled further ahead of the unexposed cohorts in participation and formality, relative to the analogous gap in less informal states.

Table 2: Cross Cohort Long Differences in 2019

	(1) Participation rate	(2) Employment rate	(3) Formality rate	(4) Monthly income
Panel A				
Pre-policy informality rate	0.50***	0.16***	0.79***	29.87***
	(0.05)	(0.02)	(0.04)	(3.88)
	[<0.001]	[<0.001]	[<0.001]	[<0.001]
Panel B				
Pre-policy informality rate	0.59***	0.15*	0.87***	11.18
	(0.09)	(0.06)	(0.16)	(10.00)
	[<0.001]	[0.053]	[0.003]	[0.278]
Log GDP per capita in 2013	6.04***	1.58	1.91	-134.53
	(0.92)	(0.99)	(1.33)	(128.86)
Educational attainment in 2015	-3.29***	-0.37	2.25	-328.57**
	(1.10)	(0.82)	(1.88)	(137.75)
Regional controls	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Mean dependent variable	-37.89	0.08	-32.75	-1,930.18
Observations	128	128	128	128

Note: Each column reports a weighted least squares regression of the 2019 exposed-minus-unexposed cohort gap on the state's pre-policy informality rate on 32 states $\times$ 4 quarters $=$ 128 state-quarter cells. Cells are weighted by the square root of the cell population (the square root of the positive-income cell population in column (4), which restricts to individuals reporting strictly positive income). Standard errors clustered at the state level in parentheses; wild cluster bootstrap p-values for the pre-policy informality coefficient in brackets (Webb weights, $9,999$ replications; Cameron et al., 2008; Roodman et al., 2019). Stars on the pre-policy informality coefficient are based on the wild-bootstrap p-values. Appendix Table A.2 shows that inference is unchanged under a block bootstrap resampling states, cluster-jackknife (CV3) standard errors, and a specification collapsed to one observation per state.
***p<0.01, **p<0.05, *p<0.1

The employment and formality rates in Table 2 are conditional rates: the employment rate is defined among participants and the formality rate among the employed, so a rise in either could in principle reflect selection rather than formalization. To separate these, Table 3 decomposes labor-market status as shares of the cohort population, so that the categories are exhaustive and the coefficients add up mechanically: the formal and informal employment responses sum to the employment response, and employment, unemployment, and non-participation sum to the participation response. The decomposition shows that the participation response is absorbed entirely by employment, that formal employment absorbs more than the full employment increase and that informal employment declines.

The formalization response therefore combines new entrants taking formal jobs with reallocation away from informal work; it is not a composition artifact of a changing employment base.[12] Of the 0.90 percentage-point differential increase in formal employment per point of initial informality, 0.63 points come from net entry into employment (almost all of it, 0.59 points, from reduced non-participation, with the remainder from a small and imprecisely estimated decline in unemployment) and 0.28 points from reallocation out of informal jobs. Roughly two thirds of the formalization response is therefore new entry and one third the formalization of work that would have been informal.

Panel C of Table 3 splits formal and informal employment aggregates by position in the job. The entire formal-employment response is subordinate wage employment (0.89 of the 0.90), with no response of formal self-employment; the decline in informal employment is likewise entirely in informal wage jobs (-0.28), with no movement in informal self-employment. Because the policy operated on the administrative prerequisites for being hired rather than on the returns to working for oneself, this is the split the salience and administrative-frictions channel predicts, and it is not what a change in the relative price of informal work would produce.

[12] Three checks support reading the formal-employment coefficient as policy-driven rather than as persistent age structure in high-informality states. First, the placebo version of the decomposition, assigning exposure to the 1995–1998 cohorts with the 1991–1994 cohorts as controls, yields a formal-employment coefficient of -0.06 (wild-bootstrap $p = 0.61$); the raw placebo gradient without controls is positive, as persistent age structure implies, and is eliminated by the controls. Second, conditioning on each state's 2013 cross-cohort gap in the same population shares (the analogue of the lagged-gap check in Table A.5) reduces the coefficient to 0.55 (wild-bootstrap $p = 0.009$), so roughly a third of the raw gradient reflects pre-existing structure. Third, in the panel with an informality-specific linear cohort trend, the coefficient is 0.47 ($p < 0.001$).

Table 3: Population-Based Decomposition of Labor Market Status, Long Differences 2019

	(1) Employment	(2) Formal employment	(3) Informal employment	(4) Unemployment	(5) Non-participation
Panel A: baseline (exposed 2001–2004 vs. unexposed 1995–1998)					
Pre-policy	0.63***	0.90***	−0.28**	−0.03	−0.59***
informality rate	(0.08)	(0.09)	(0.08)	(0.04)	(0.09)
	[<0.001]	[<0.001]	[0.016]	[0.514]	[<0.001]
Panel B: placebo ("exposed" 1995–1998 vs. 1991–1994)					
Pre-policy	−0.10	−0.06	−0.04	−0.01	0.11
informality rate	(0.08)	(0.10)	(0.08)	(0.03)	(0.07)
	[0.235]	[0.608]	[0.583]	[0.834]	[0.171]
Panel C: formal and informal employment by type of job					
	Formal, wage emp.	Formal, self-emp. & other	Informal, wage emp.	Informal, self-emp. & other	
Pre-policy	0.89***	0.01	−0.28**	0.00	
informality rate	(0.09)	(0.01)	(0.07)	(0.05)	
	[<0.001]	[0.237]	[0.012]	[0.953]	
Panel D: formal employment under alternative specifications					
Conditioning on 2013 cross-cohort gap					0.55 [0.009]
Panel, 2016–2019 (cohorts 1991–2004)					0.96 [<0.001]
Panel, informality × cohort trend					0.47 [<0.001]
Excluding Chiapas, Guerrero, Oaxaca					0.92 [<0.001]
Controls	✓	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓	✓
Observations	128	128	128	128	128

Note: Each outcome is a share of the cohort-group population, so the columns are exhaustive, up to rounding: (2) and (3) sum to (1), and (1), (4), and (5) sum to zero. Each observation is a state-quarter cell in 2019; the dependent variable is the exposed-minus-unexposed cohort gap in the indicated share. Panel B assigns placebo exposure to the 1995–1998 cohorts with the 1991–1994 cohorts as controls; without controls the placebo formal-employment gradient is positive (0.12, wild-bootstrap $p = 0.047$), reflecting persistent cross-state age structure that the controls absorb. Panel C splits formal and informal employment into subordinate remunerated (wage) employment and self-employment and other positions (employers, own-account workers, and unpaid workers); no employed respondent has an unspecified position, and the wage and self-employment coefficients sum to the Panel A aggregates. Panel D reports the formal-employment coefficient conditioning on the state's 2013 cross-cohort gap in the same population share; in the 2016–2019 panel of cohorts 1991–2004 (7,168 cohort-state-quarter cells) without and with an informality-specific linear cohort trend; and excluding the three southernmost states (116 cells). Baseline 2019 levels (percent of cohort population): employment 21.3 (exposed) and 56.2 (unexposed); formal employment 1.6 and 23.9; informal employment 19.7 and 32.3; unemployment 1.5 and 4.5; non-participation 77.1 and 39.4. Controls: log state GDP per capita (2013), educational attainment (2015), and regional dummies. Weighted least squares with square root of cell population weights. Standard errors clustered by state in parentheses; wild cluster bootstrap p-values (Webb weights, 9,999 replications) in brackets. ***p<0.01, **p<0.05, *p<0.1

The analysis is done with information from 2019 to evaluate the medium-term effects of the policy for the fully exposed cohorts in 2016, right before the onset of the COVID-19 pandemic. To consider the full picture, Figure 3 plots the estimated coefficients ($\hat{\beta}_t$) of the long differences equation for each quarter t from 2016 to 2022, keeping the same exposed and unexposed cohort groups throughout. Two caveats apply: as cohorts age over this window the comparison captures increasingly medium- and longer-run effects rather than a fixed early-career snapshot, and some control cohorts may have gained partial exposure through public college after 2016. Pandemic-related disruptions after 2019 introduce an additional break in comparability. The 2016–2019 portion of the figure is therefore the most directly comparable to the long-difference estimates above and the panel estimates below. Within that window, the effects on participation and formality rates are positive throughout, though the formality coefficient declines from mid-2019.[13] The post-2019 path cannot by itself distinguish a transitory entry effect from the pandemic, and it should be read alongside the age-profile specification in Table A.4, column (5), which nets out the convergence of the exposed cohorts' gradient toward that of older cohorts as they age. After the COVID-19 pandemic, the estimated coefficients for both of these indicators are close to zero and no longer statistically significant.

It is not clear how to interpret these estimates as the pandemic severely disrupted labor markets. There is evidence that in Latin America there was an initial drop in participation, employment, and informality (followed by a rebound in the latter), with the informal sector acting as an important margin of adjustment (Alvarez and Pizzinelli, 2021; Leyva and Urrutia, 2023). Furthermore, Osuna-Gómez (2023) finds that people that entered the labor market after the Great Recession lost formal employment during the pandemic at higher rates than previous cohorts. It could be the case that the pandemic shock reverted any effects that were driven by the school-linked social insurance policy on labor market outcomes. However, there are additional concerns for more recent estimates as we are tracking cohorts at the state level and not individuals, and we could worry about potential migration patterns and readjustments between different local labor markets. The

[13]Two federal policies changed in January 2019 and warrant comment. First, *Jóvenes Construyendo el Futuro*, an apprenticeship program providing a stipend and IMSS medical coverage to 18–29 year-olds who are neither studying nor working. The comparison cohorts (aged 21–24 in 2019) fall in the eligible age range while the exposed cohorts (aged 15–18) do not. To the extent apprenticeships are recorded as formal employment the program raises measured formal attachment among the comparison cohorts and narrows the estimated gap. Second, the general minimum wage rose 16 percent and the minimum wage in the northern border free zone doubled; the free zone covers municipalities accounting for under 9 percent of the 2019 sample, and those municipalities lie in states at the low end of the informality distribution. Both would have to inflate the 2019 estimates to threaten the results, and neither does. Re-estimating the long differences on 2018 gives 1.16 for formality against 0.87 in 2019 and 0.55 for participation against 0.59, so 2019 is if anything the more conservative year for formality. Excluding the six states containing free-zone municipalities raises rather than lowers both coefficients (to 0.95 and 0.63 for formality and participation, respectively), and the informality gradient does not differ significantly in those states in 2019. Because the exposed cohorts are one year younger in 2018, the two years are not interchangeable estimates of the same quantity; the comparison bears on the direction of any 2019-specific bias, not on which year is preferred.

employment and income estimates are mostly small and not significant.

Figure 3: Estimated coefficients of pre-policy informality rate on cross cohort long differences, $\hat{\beta}_t$

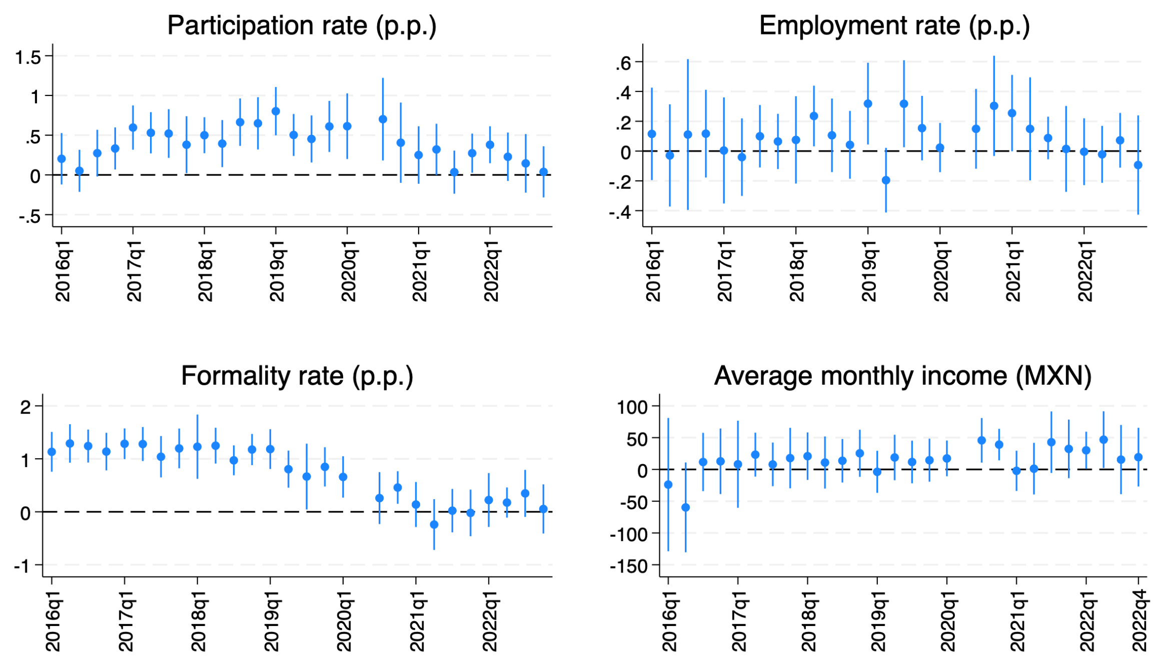

Note: Each point plots the estimated coefficient from quarterly cross-cohort long difference regressions (exposed minus unexposed) on pre-policy state informality (2013–2015). No available data for 2020q2.
Source: Author's calculations with data from ENOE.

3.2.2 Panel Estimates

The long difference estimates transparently summarize medium-run gaps between fully exposed and unexposed adjacent cohorts across states with different initial informality. Their limitation is that they do not, by construction, demonstrate that such gaps were not already drifting with pre-policy conditions. To address this, I turn to a panel framework that (a) uses all quarters from 2016 to 2019 to absorb common shocks with survey-wave fixed effects and (b) follows Bleakley's cohort-space logic: for each birth cohort I recover the cross-state gradient with respect to pre-policy informality and show that this gradient is within a narrow band for cohorts too old to be exposed and turns on only at the eligibility cutoff. In addition, I report robustness with state-specific linear time trends to soak up differential growth by state unrelated to the policy.

First, to incorporate cohorts that were partially exposed to the policy I estimate pooled regressions of the form:

$$Y_{kst} = \tilde{\beta} I_{s,pre} \times \text{Exp}_k + \delta_k + \delta_s + \delta_t + \sum_i \left(x_s^i \times \text{Exp}_k \right) \gamma_i + \nu_{skt},$$

where Y_{kst} is the labor market outcome of interest (participation rate, employment rate, formality rate, monthly income) for cohort k in state s and quarter t, Exp_k is exposure to the policy for cohort k as defined by figure 1, $I_{s,pre}$ is the pre-policy informality rate in state s, the x_s^i are state-specific controls, and $\delta_k, \delta_s, \delta_t$ are fixed effects for cohort, state, and quarter, respectively. Because the controls enter interacted with exposure, the coefficient of interest is identified from variation in pre-policy informality that is orthogonal to state GDP per capita, educational attainment, and regional cohort profiles.

Second, I compare the evolution of labor market outcomes of interest by cohort across states to assess the contribution of the policy to the observed changes. The regressions to be estimated are:

$$Y_{kst} = \beta_k I_{s,pre} + \delta_k + X_s \Gamma_k + \nu_{skt},$$

where β_k is a cohort specific coefficient on pre-policy informality, X_s is a vector of state controls, and δ_k and Γ_k are cohort-specific intercept and slope coefficients. By estimating the equation through weighted least squares for each cohort, I generate a series of estimates across cohorts. The series of cohort estimates helps to show graphical evidence that the results are capturing the impact of the policy by comparing it with the pattern of the exposure measure.

Before presenting the estimates of the pooled and cohort-specific regressions, figure 4 plots the summary statistics for the panel constructed by cohort in states and quarters between 2016 and 2019. Each observation represents the value of the dependent variable for each cohort-state-quarter tuple. The size of the observations is determined by the size of the population represented. I single out in red the three poorest states (Guerrero, Oaxaca, and Chiapas), which also happen to be the southernmost states.

We observe expected patterns in the data, such as lower participation rates and lower formality rates for younger cohorts. Among the poorest states it is interesting to observe that for the oldest cohorts analyzed, the participation rate is lower relative to richer states but for the youngest cohorts, relative to richer states, the participation rate is higher. It is also worth noting that the employment rate is consistently higher for the poorest states across all cohorts, but the quality of jobs within these states is worse, as seen by the lower formality rate and lower monthly income.

Figure 4: Summary statistics of panel by cohort across states and quarters, 2016–2019

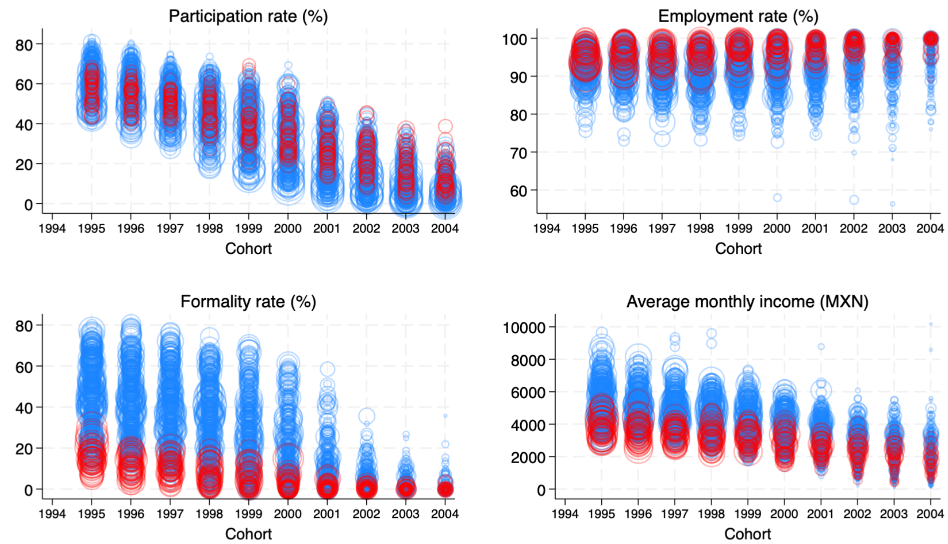

Note: Each observation represents the value of the dependent variable for each cohort by state-quarter. The red observations denote the three poorest states (Guerrero, Oaxaca, and Chiapas). The size of the observations is determined by the size of the population for each cohort-state-quarter tuple.
Source: Author's calculations with data from ENOE.

Table 4 presents the estimates for the pooled regressions. The regressions are estimated through weighted least squares with the weights constructed from the square root of cell size. Standard errors are clustered at the state level. These estimates are consistent with the results of the long differencing. The magnitude of the estimate for the participation rate decreased but remains significant. The estimate for the employment rate is very similar, whereas the estimate for the formality rate increased to almost 1.3 percentage points associated with an increase of 1 percentage point in the initial informality rate. This means that taking into account the panel structure, which adds quarters in which the exposed cohorts are younger, and adds the partially exposed cohorts through the continuous measure, increases the differential positive impact of the policy on the formalization of exposed cohorts in high informal states relative to low informal states. Income effects are statistically detectable but economically small (≈ 1.1 USD per month), consistent with a reallocation toward formal jobs without large short-run gains in earnings.

Table 4: Panel Estimates of the Effect of Exposure to Policy, 2016–2019

	(1) Participation rate	(2) Employment rate	(3) Formality rate	(4) Monthly income
Pre-policy informality rate × Exposure	0.44***	0.15***	1.28***	21.33**
	(0.05)	(0.03)	(0.06)	(7.05)
	[<0.001]	[0.001]	[<0.001]	[0.021]
Log GDP per capita in 2013 × Exposure	2.01	0.43	3.67**	-61.73
	(1.01)	(1.02)	(1.01)	(114.12)
Educational attainment in 2015 × Exposure	-0.33	1.86**	1.86	-72.11
	(0.67)	(0.55)	(1.22)	(86.32)
Cohort FE	✓	✓	✓	✓
State FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Mean dependent variable	33.31	93.84	17.99	4,299.21
Observations	5,094	5,091	5,091	5,046

Note: Standard errors clustered by state in parentheses. Cells are weighted by the square root of the cell population (the square root of the positive-income cell population in column (4), which restricts to individuals reporting strictly positive income). Wild cluster bootstrap p-values for the informality-by-exposure coefficient in brackets (Webb weights, 9,999 replications).
***p<0.01, **p<0.05, *p<0.1

A natural concern is that labor-demand fluctuations (national or state-specific) could differentially affect youth and mimic treatment heterogeneity. The panel specification mitigates this along several layers, and I add targeted robustness checks. Time fixed effects absorb nationwide demand movements (e.g. macro cycles, fuel-price spike), ensuring identification does not come from aggregate time shocks. State fixed effects soak up permanent gaps in demand, institutions, and composition across states (industrial structure, enforcement). Cohort fixed effects net out nationwide age/cohort profiles in youth demand or school-to-work transitions that are common across states. The identifying variation compares exposed vs. unexposed adjacent cohorts within the same state and quarter, scaled by the state's pre-policy informality rate. Any state-quarter demand shock that affects young cohorts symmetrically cancels in this within-state contrast. Additionally, I augment the panel with state-specific linear time trends so that smooth state-level growth is absorbed. Table A.3 in the Appendix shows that the results are materially unchanged. The stability of the estimated coefficients strongly mitigates concerns for a demand-driven explanation, though non-linear state-specific shocks that interact with both quarter and cohort cannot be fully excluded by design.

Figure 5: Cohort-specific estimates on pre-policy informality rate, $\hat{\beta}_k$

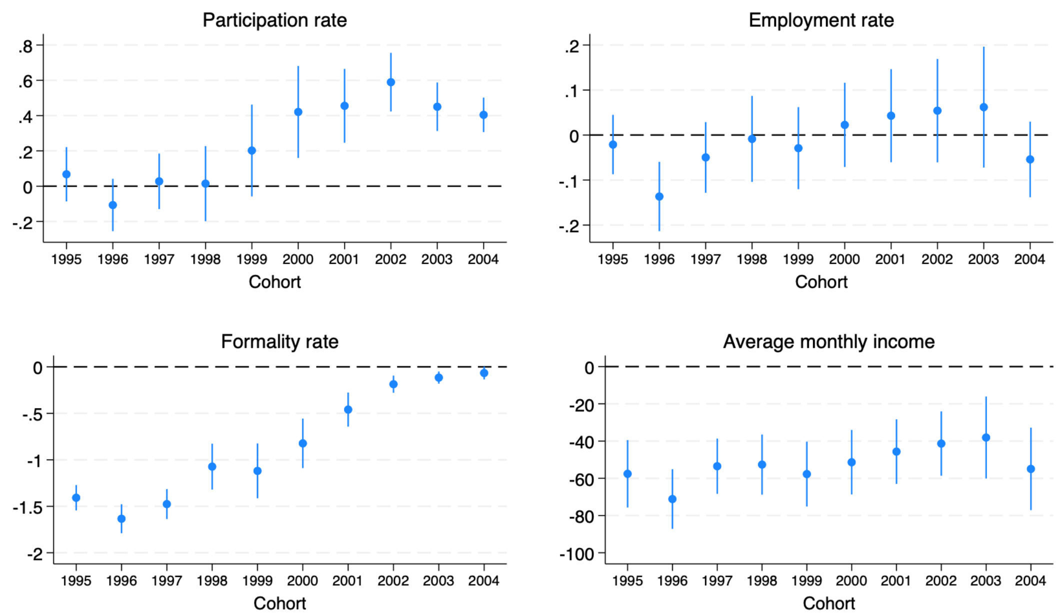

Note: Each point plots the cohort-specific slope of the outcome with respect to the state pre-policy informality rate (2013–2015). Consistent with the exposure schedule, estimates are within a narrow band for unexposed cohorts and shift at the eligibility cutoffs for participation and formality.
Source: Author's calculations with data from ENOE.

Figure 5 plots the cohort-specific estimates on the pre-policy informality rate for the different outcomes of interest. The pattern of estimates is broadly consistent with the exposure to the policy, as defined in figure 1: the gradients sit within a narrow band for cohorts too old to be exposed and shift sharply across the partial-exposure cohorts (upward for participation and toward zero for formality). This graphical evidence points to the relevance of the policy in explaining the results of the long-differencing and pooled regression analyses. The next section develops this cohort-space evidence into formal tests of the identifying assumption.

3.3 Identification

In this section I test the validity of the identifying assumption underlying the previous results. Two related threats could provide alternative explanations for the main estimates. The first is convergence: even absent the policy, participation and formality of younger cohorts in high-informality states might have risen as these poorer, more informal areas catch up with richer, more formal ones. The second is heterogeneous cohort-age profiles: if the age at which young people enter work and formalize varies systematically with a state's informality, cross-cohort gaps would correlate with pre-policy informality mechanically. I address these threats in five steps: formal tests on the cohort-specific gradients, placebo versions of both research designs on untreated cohorts, specifications that allow informality-specific cohort and age profiles, a sensitivity analysis that bounds rather than tests violations of the identifying assumption, and checks on influential states and on pre-policy trends over time.

To move from the visual evidence in figure 5 to formal tests, I estimate an alternative, fully interacted version of the cohort-specific equation:

$$Y_{kst} = \alpha_k + \beta_k I_{s,pre} + X_s \Gamma_k + \delta_{kt} + \nu_{kst}$$

in which the intercept, the informality gradient, the control coefficients, and the quarter effects are all allowed to differ by cohort and the equation is estimated jointly on cohort-state-quarter cells over 2016–2019, weighted by the square root of cell population. Joint estimation delivers cohort-specific gradients numerically identical to estimating each cohort separately, but yields a single state-clustered covariance matrix, so hypotheses that compare gradients across cohorts can be tested formally. I focus on participation and formality, the two outcomes for which figure 5 shows systematic cross-cohort patterns. Figure 6 plots the resulting gradients with 95% confidence intervals.

The untreated cohorts' gradients lie in a narrow band (-0.46 to -0.26 for participation, -1.58 to -1.29 for formality) whose width is roughly one fifth of the treatment-induced movement for formality and two fifths for participation. A formal test rejects their exact equality (participation: $F(7, 31) = 7.4$; formality: $F(7, 31) = 10.8$; both $p < 0.001$), reflecting the precision of the cell-level estimates, but the direction of the variation is the informative object: the estimated linear drift across untreated cohorts is negative for both outcomes (participation: -0.022 per cohort, wild-bootstrap $p = 0.03$; formality: -0.020, $p = 0.15$), opposite in sign to the treatment effect. Smooth cohort trending in high-informality states would push exposed cohorts' gradients further down, not sharply up as observed; under a linear trend counterfactual the baseline estimates are, if anything, conservative. The uptick in the 1998 cohort's gradient relative to 1995–1997 is consistent with partial exposure through the public college channel, which affects the youngest "untreated" cohorts first (39% of the 1995–1998 cohorts are observed in school, either public or private, after the policy).

Figure 6: Cohort-specific pre-policy informality gradient

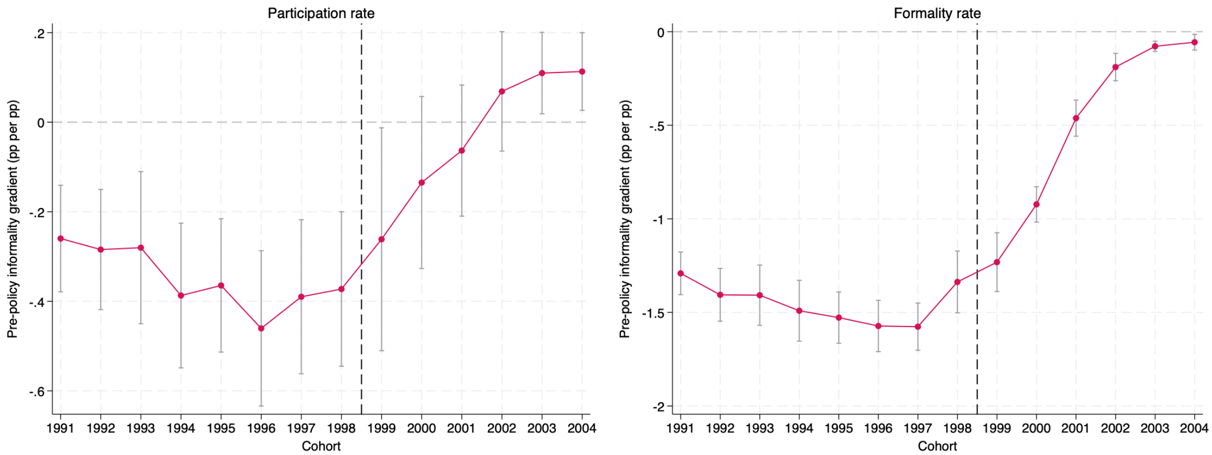

Note: Each point is the cohort-specific informality gradient β_k from the fully interacted specification described in Section 3.3, in which intercepts, informality gradients, control coefficients, and quarter effects all vary by cohort, estimated jointly on cohort-state-quarter cells over 2016–2019 with square-root-of-cell population weights. Bands are 95% confidence intervals based on standard errors clustered by state. The dashed line separates untreated cohorts (1991–1998) from cohorts with positive exposure.
Source: Author's calculation using data from ENOE.

As a falsification test aimed directly at convergence, I replicate both research designs on cohorts largely untreated by the high-school component of the policy, assigning placebo exposure to the 1995–1998 cohorts with the 1991–1994 cohorts as controls. Although some individuals in the 1995–1998 cohorts may have had partial exposure through public college, they were not subject to the high-school benefit assignment that drives the main analysis. If differential catch-up in high informality states were driving the main estimates, the placebo coefficients should be positive and of comparable magnitude to the baseline. Table 5 shows the opposite for the two key outcomes: for participation and formality, placebo coefficients are statistically indistinguishable from zero or negative, pointing to divergence, with confidence intervals excluding effects as large as the baseline. Panels B and C do show positive and significant placebo estimates for employment and income, but these are not the outcomes driving the paper's main conclusions on formalization and likely reflect underlying age-earnings profiles among these older cohort groups; consistent with this reading, the analogous long-difference income placebo is statistically indistinguishable from zero once state controls are included. This pattern undercuts the convergence hypothesis and supports a causal differential ITT interpretation for participation and formality. The population-based decomposition passes the same falsification (Table 3).

Table 5: Placebo Regressions

	(1) Participation rate	(2) Employment rate	(3) Formality rate	(4) Monthly income
Panel A: Long Differences, 2019				
Pre-policy informality rate	-0.11	-0.01	-0.24*	-7.22
	(0.07)	(0.04)	(0.12)	(7.86)
	[0.171]	[0.811]	[0.099]	[0.433]
Full set of controls	✓	✓	✓	✓
Mean dependent variable	-12.47	-2.17	-8.07	-1,050.38
Observations	128	128	128	128
Panel B: Pooled Regressions, 2016–2019				
Pre-policy informality rate $\times$ Exposure	-0.09**	0.04*	-0.10	10.83**
	(0.03)	(0.02)	(0.05)	(3.23)
Full set of controls	✓	✓	✓	✓
Mean dependent variable	60.42	93.00	40.45	5,288.04
Observations	4,096	4,096	4,096	4,096
Panel C: Pooled Regressions, 2016				
Pre-policy informality rate $\times$ Exposure	-0.01	0.08**	0.14	10.11**
	(0.04)	(0.03)	(0.09)	(2.94)
Full set of controls	✓	✓	✓	✓
Mean dependent variable	54.39	91.88	34.81	4,490.38
Observations	1,024	1,024	1,024	1,024

Note: Placebo exposed cohorts are 1995–1998, with 1991–1994 as controls. Standard errors clustered by states in parentheses, wild cluster bootstrap p-values for the pre-policy informality coefficient in brackets for the long differences. The square root of cell sizes is used to construct weights for the observations. Cells are weighted by the square root of the cell population (the square root of the positive-income cell population in column (4), which restricts to individuals reporting strictly positive income).
***p<0.01, **p<0.05, *p<0.1

The placebo evidence rules out catch-up among older cohorts, but heterogeneous profiles could in principle operate smoothly across all cohorts. Table A.4 therefore augments the panel specification with interactions of pre-policy informality and smooth cohort or age profiles, extending the sample to cohorts 1991–2004 so that the added untreated cohorts help pin down the profiles. The formality effect is robust throughout, ranging from 0.78 to 1.21 and always significant at the 1 percent level by wild cluster bootstrap. The participation effect attenuates to 0.15–0.22 but remains positive and significant at conventional levels. Given the exposure function ramps from zero to one over only four cohorts, flexible informality-specific trends absorb part of the exposure variation by construction, so these estimates should be read as conservative bounds rather than preferred point estimates. These specifications also address a mechanical concern specific to comparing cohorts at different points of the age-formality profile. Because formal employment among 15–18 year-olds is close to zero in every state (0.6 to 3.1 percent of the cohort population across informality quartiles), one might worry that the cross-state gradient in the cohort gap simply tracks the steeper decline of formal employment among the older comparison cohorts. Any such base-rate mechanism implies an informality gradient that varies smoothly with age, which is exactly what columns (4) and (5) absorb; the exposure coefficient is identified from the discrete exposure ramp net of those profiles, and remains between 0.78 and 0.93.

The tests above establish that the untreated cohorts' gradients are not exactly equal, and that the direction of their variation is opposite in sign to the estimated policy effect. Treating this as a pass/fail screen would be unsatisfactory on two counts: with 32 clusters the test has limited power against violations large enough to matter, and conditioning the analysis on having passed a pre-test distorts subsequent inference (Roth, 2022). I therefore report how large a violation of parallel cohort profiles would have to be to overturn the results, following Rambachan and Roth (2023). Treating the birth cohort as event time and normalizing on the 1997 cohort, I bound the informality-specific cohort profile by $\Delta^{RM}(\bar{M})$, which permits the post-1997 violation to be at most $\bar{M}$ times the largest violation observed among untreated cohorts. For formality the exposure differential remains significantly positive for all $\bar{M} \leq 1.15$: the estimate survives even if the counterfactual cohort profile departs from parallel trends by somewhat more than the largest departure visible in the untreated data. For participation the breakdown value is 0.43, consistent with the greater sensitivity of the participation margin already apparent in Table A.4. Under $\Delta^{SD}(0)$, which extrapolates the untreated cohorts' linear drift forward rather than assuming it away, the formality estimate rises rather than falls, as the negative drift documented above implies. Appendix C describes the construction and reports the full set of results.[14]

[14]Two implementation choices are worth noting. First, the procedure requires a covariance matrix and cannot accept the wild cluster bootstrap used elsewhere in the paper. I therefore use cluster-jackknife (CV3) standard errors, which are the natural analogue with 32 clusters (MacKinnon et al., 2023). Using CRV1 instead would give a breakdown value of 1.38 for formality rather than 1.15. Second, the reference is the 1997 rather than the 1998 cohort, since the latter carries more of the partial exposure through the public college channel; normalizing on 1998 gives 0.78. Both alternatives are reported in Table C.1.

The estimates are also not driven by influential states. Re-estimating the main specifications excluding one state at a time yields formality coefficients ranging from 0.79 to 0.98 in the long differences and from 1.24 to 1.31 in the panel, always significant at the 1 percent level; participation coefficients are similarly stable. In particular, excluding jointly the three southern states with the highest informality and the most atypical age-participation profiles (Chiapas, Guerrero and Oaxaca) leaves the estimates essentially unchanged (Appendix Table A.2 reports the full set of inference and leverage checks).

Finally, I estimate pre-policy trends over time. Table A.5 conditions the long differences on each state's cross-cohort gaps measured three years before the start of the policy. The participation estimate is unchanged (0.63 versus 0.59 at baseline), and the formality estimate remains large and significant at the 5 percent level (0.58, wild-bootstrap $p = 0.02$), about two thirds of the baseline. The lagged formality gap itself enters positively, indicating that part of the raw formality gradient reflects persistent state-level differences in young-versus-old formality gaps. The estimate net of this persistence, 0.58, sits somewhat below the profile-augmented panel range of 0.78 to 1.21, so two independent adjustments (conditioning on lagged gaps in the cross-section and allowing informality-specific cohort and age profiles in the panel) both reduce the formality effect and leave it positive and significant. Figure A.3 in the Appendix provides the corresponding visual evidence: quarterly cohort gaps between the treated (15–18) and control (21–24) age groups, computed by state since 2010, smoothed by locally weighted regression, and split at the median pre-policy informality rate, evolve in parallel across the two groups of states before 2016 and diverge only after policy implementation.

4 Mechanisms and Robustness

An important consideration is whether this policy had an impact on the schooling decisions of individuals: medical benefits tied to enrollment could raise schooling to the detriment of early participation, an outcome that could be desirable for the individual and for society. The increased participation among exposed cohorts in high-informality states already suggests otherwise. This section uses the margins along which the response varies to distinguish among the channels of Section 2.2 and to establish robustness. Section 4.1 examines the enrollment margin directly and splits the estimates by current school attendance; Section 4.2 splits them by gender and by the formality of the household head; Section 4.3 re-estimates them under an alternative, employees-only definition of informality. Figure 7 collects these splits for both research designs, with the underlying estimates in Appendix Table A.7. Section 4.4 examines the measurement of exposure and the partial treatment of the comparison cohorts; Section 4.5 brings in ENILEMS, whose public-school indicator permits a sharper exposure assignment; and Section 4.6 draws the evidence together.

4.1 School Enrollment

First, I explore the impact of the social insurance expansion policy on school enrollment rates of the cohorts analyzed. Table 6 shows estimates of the long differences and panel equations presented in sections 3.2.1 and 3.2.2, using school enrollment rates as the dependent variable. There is no evidence of an impact of the policy on schooling rates: neither the long-difference estimate nor the panel estimate is statistically significant.

Table 6: Long Differences and Panel Estimates of School Enrollment Rates

	(1) Long Differences	(2) Panel Estimates
Pre-policy informality rate	-0.21	-0.12
	(0.15)	(0.09)
	[0.220]	[0.390]
Log GDP per capita in 2013	-9.36***	-1.76
	(3.27)	(1.60)
Educational attainment in 2015	4.41*	-2.55*
	(2.28)	(1.38)
Regional controls	✓	
Cohort FE		✓
State FE		✓
Quarter FE	✓	✓
Mean dependent variable	46.42	63.97
Observations	128	5,120

Note: Standard errors clustered by state in parentheses; wild cluster bootstrap p-values for the pre-policy informality coefficient in brackets (Webb weights, 9,999 replications). The square root of cell population is used as weights for the observations. For column (2) the independent variables are interacted with the policy exposure function as in Table 4.
***p<0.01, **p<0.05, *p<0.1

Now, I distinguish the main results of the previous section between workers going to school at the time of the survey and those not going. Table 7 presents the estimates of long differences on the main outcomes of the labor market distinguishing by school enrollment. Panel A shows the main results, considering the full sample, which are exactly those presented previously in Table 2. Panel B shows the long differences estimates using the sample of people not going to school and Panel C shows the results using the sample of students at the time of the survey. Approximately 53% of the full sample studied is enrolled in school. The long-difference estimates are larger among those attending school for all four outcomes. In the case of the formality rate, a one-percentage-point increase in the pre-policy informality rate is associated with a 0.95-percentage-point wider exposed-minus-unexposed cohort gap in the formality rate among those attending school. Formality is significant at the 1 percent level by wild cluster bootstrap in both subsamples, and participation in both as well (at the 5 percent level among those not in school); the smaller employment and income estimates are at most marginally significant.

Figure 7 places these long-difference splits alongside their panel counterparts; Table A.7 reports every estimate. The two designs agree on the central point: both groups respond, with participation and formality positive and significant for students and non-students alike in every specification. They differ only in ordering: in the long differences the in-school formality estimate (0.95) exceeds the not-in-school one (0.83), while in the panel the ordering reverses (1.00 against 1.15). What matters for mechanism is the absence of concentration rather than the ordering: the formalization response is at least as large among youths who have already left school as among current students, which is not the signature benefit substitution would leave, since its effects should be confined to the currently enrolled and dissipate at school exit.

An important caveat is that school enrollment in Table 7 is measured contemporaneously with labor market outcomes (that is, after policy exposure) so the in-school and not-in-school subgroups are not fixed populations. Even though Table 6 finds no clear average enrollment response, conditioning on current enrollment may still shift subgroup composition in ways that vary with pre-policy informality. The in-school and not-in-school estimates should therefore be read as suggestive heterogeneity rather than clean effects for fixed student and non-student populations. With that caveat in mind, the results are consistent across both subsamples. In particular, they provide suggestive evidence that students affected by the social insurance expansion policy are not more likely to be employed informally, despite the reduction in the relative cost of an informal job while also being enrolled in school. This is more relevant as there is evidence that there was an increase in the participation rate among exposed cohorts in states with higher informality. The absence of an enrollment response also means the participation increase documented in Section 3 does not come at the expense of schooling.

Table 7: Cross Cohort Long Differences by School Enrollment, 2019

	(1) Participation rate	(2) Employment rate	(3) Formality rate	(4) Monthly income
Panel A: Full Sample				
Pre-policy informality rate	0.59***	0.15*	0.87***	11.18
	(0.09)	(0.06)	(0.16)	(10.00)
	[<0.001]	[0.053]	[0.003]	[0.278]
Log GDP per capita in 2013	6.04***	1.58	1.91	-134.53
	(0.92)	(0.99)	(1.33)	(128.86)
Educational attainment in 2015	-3.29***	-0.37	2.25	-328.57**
	(1.10)	(0.82)	(1.88)	(137.75)
Regional controls	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Mean dependent variable	-37.89	0.08	-32.75	-1,930.18
Panel B: Sample Not In School				
Pre-policy informality rate	0.37**	0.10	0.83***	18.86
	(0.13)	(0.07)	(0.17)	(12.54)
	[0.039]	[0.255]	[0.007]	[0.155]
Log GDP per capita in 2013	-1.02	2.43**	2.32	-247.01*
	(3.49)	(1.11)	(1.58)	(135.31)
Educational attainment in 2015	-2.65	-1.80	2.59	-110.39
	(2.17)	(1.07)	(1.86)	(175.33)
Regional controls	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Mean dependent variable	-17.25	-1.13	-31.43	-1,493.67
Panel C: Sample In School				
Pre-policy informality rate	0.56***	0.20*	0.95***	29.16
	(0.08)	(0.09)	(0.19)	(15.33)
	[<0.001]	[0.076]	[0.001]	[0.115]
Log GDP per capita in 2013	4.06*	1.20	-2.91	-363.24
	(2.13)	(1.33)	(4.34)	(494.58)
Educational attainment in 2015	-0.79	1.52	5.07	49.21
	(1.70)	(0.92)	(3.80)	(269.29)
Regional controls	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Mean dependent variable	-21.95	0.68	-30.78	-2,120.29
Observations	128	128	128	128

Note: Standard errors clustered by state in parentheses; wild cluster bootstrap p-values for the pre-policy informality coefficient in brackets (Webb weights, 9,999 replications); stars on that coefficient are based on the wild-bootstrap p-values. Panels B and C restrict the sample to individuals not attending and attending school at the time of the survey, respectively, with cell rates computed over each subsample's population. Cells are weighted by the square root of the cell population (the square root of the positive-income cell population in column (4), which restricts to individuals reporting strictly positive income). Regional controls and quarter fixed effects included in all panels.
***p<0.01, **p<0.05, *p<0.1

It is uncertain whether it is desirable to have more individuals start working while they study. Employment spells while studying could provide valuable training experiences. On the other hand, it could deter from further investing in schooling (Atkin, 2016). However, if anything, it is desirable to have students work formally vs. informally at the beginning of their labor market trajectories as there are negative employment and earning effects associated with youth informality spells.

4.2 Heterogeneous Results

I next split the estimates by gender and by the formality of the household head; Figure 7 plots both splits and Table A.7 reports the estimates. The sample is roughly evenly divided by gender. In the long differences the participation and formality point estimates are larger for men (0.90 and 0.97) than for women (0.27 and 0.70); in the panel the two are close (0.47 against 0.40 for participation, and 1.27 against 1.30 for formality), and the employment and income estimates favor women in both designs. The formalization response is therefore not a single-gender phenomenon: formality is positive and significant for men and women in both designs, and I do not find significant differences by gender.

The split by the formality of the household head speaks to whether the response requires the coverage itself to be new. Children of insured workers are already covered as dependents, automatically below age 16 and, in the IMSS case, up to age 25 while enrolled in the national education system, so most youths in households with a formal head had access to medical coverage before 2016 and the student program could add little to it; what it added was a permanent social security number in their own name, since students holding equivalent protection were not enrolled in the student scheme (Appendix B). If the response ran through the value of newly acquired coverage, it should be concentrated among youths whose household head was informal. The panel estimates show otherwise: the formality gradient is nearly identical for the two groups (0.92 with an informal head, 0.96 with a formal head), and the participation gradient is larger, not smaller, where the head is formal (0.59 against 0.36). The response is therefore not local to households without prior coverage, which is consistent with the registration and salience channel and difficult to reconcile with a coverage-valuation channel alone.

4.3 Alternative Sample of Informal Workers

Up to this point, I have been using the official definition of informal employment established by the Mexican National Institute of Statistics and Geography (INEGI). This definition includes both employees and self-employed individuals. I explore how the main results change if we focus only on subordinate employment, i.e., employees without access to social insurance. I use this alternative sample that only considers employees for the pre-policy period as well as for the main period of analysis (see figure A.4 for the spatial distribution of the alternative pre-policy informality across states).

For these estimates both the outcome samples and the regressor change: the pre-policy informality rate is recomputed among employees only, and its cross-state range is roughly half that of the baseline measure, so coefficient magnitudes are not directly comparable across definitions. The informative content is that sign and significance survive the redefinition: the estimates are noisier, as expected, and the panel estimates show a positive and significant response of participation and formality among exposed cohorts in high-informality states (formality 1.27, wild-bootstrap $p = 0.008$; Table A.7).

Figure 7 summarizes all four checks in this section – school enrollment, gender, formality of the household head, and the alternative informality definition – in one place, for both the long-difference and panel specifications. Point estimates, standard errors, and wild-cluster-bootstrap p-values for every split shown are reported in Table A.7.

Figure 7: Heterogeneity and robustness of the pre-policy informality rate coefficient

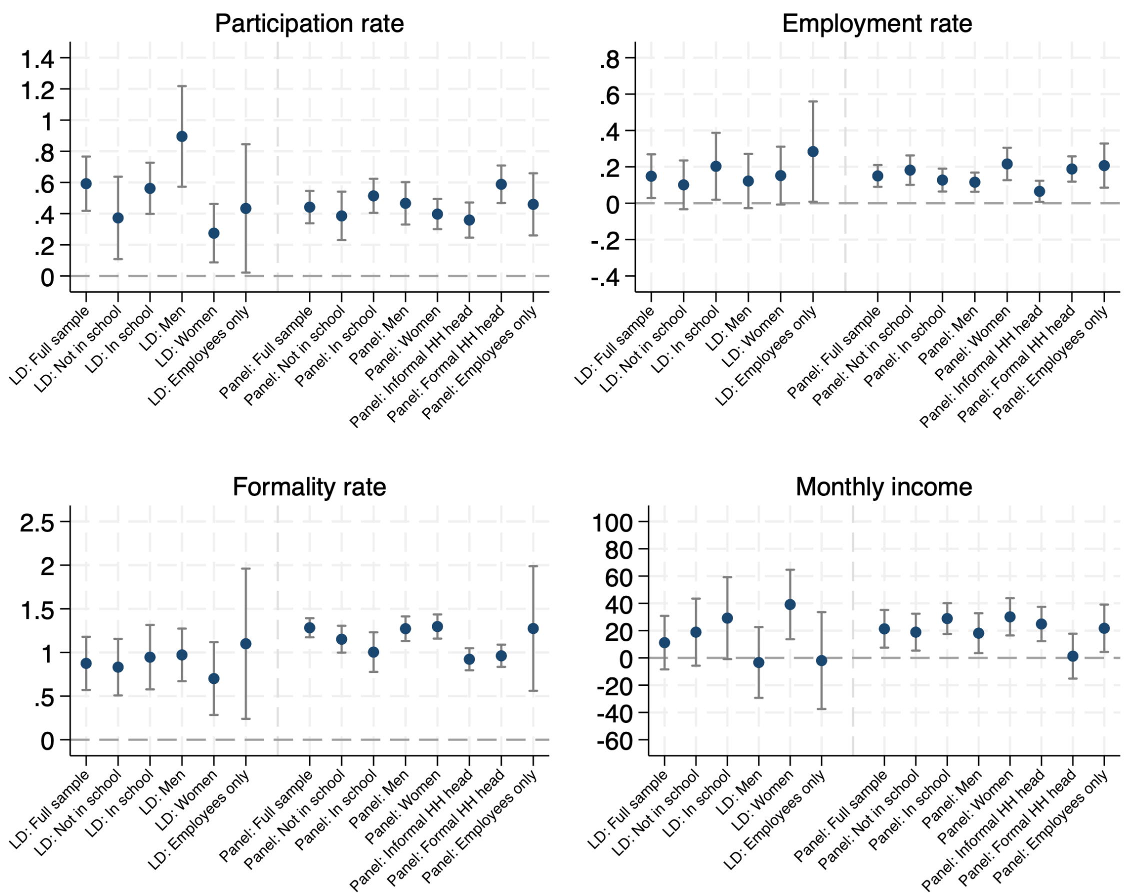

Note: Points are coefficients on the pre-policy informality rate (long differences, LD) or on its interaction with the policy exposure function (panel). Long-difference estimates (2019) regress the exposed-minus-unexposed cohort gap on state pre-policy informality (2013–2015); panel estimates (2016–2019) report the coefficient on the interaction between pre-policy informality and the exposure function. Bars are 95% confidence intervals based on standard errors clustered by state; the wild-cluster-bootstrap p-values used for inference throughout the paper are reported, together with the point estimates and clustered standard errors, in Appendix Table A.7. "Full sample" entries reproduce the corresponding estimates from Tables 2 and 4. The household-head split is estimated in the panel design only. In the "Employees only" entries both the outcome samples and the pre-policy informality measure are restricted to subordinate employees (figure A.4), so their magnitudes are not directly comparable to the other entries.
Source: Author's calculations with data from ENOE.

4.4 Measurement of Exposure and Contamination

The baseline design assigns exposure at the cohort level, but effective exposure to the school-linked benefit required being enrolled, and enrolled in a public institution, both of which vary across states. To assess whether this measurement matters, I construct a predicted exposure measure that scales cohort exposure by each state's pre-policy school-attendance rate of 15–18 year-olds (ENOE, 2013–2015; mean 71.5% ranging from 58% to 83% across states) and by the public share of upper-secondary enrollment in the 2015–2016 school year (roughly 81% nationally, see figure A.5). Table A.6 re-estimates the panel specification replacing cohort exposure with predicted exposure. The cross-state pattern is unchanged in sign and significance, and the coefficients rise when the treatment variable is scaled toward actual eligibility, indicating that the findings are not an artifact of assuming uniform exposure within cohorts.

The same ingredients discipline the interpretation of magnitudes and of the contamination discussed in the research design. Two facts stand out. First, partial treatment of the comparison group is substantial: 39% of the 1995–1998 cohorts are observed attending school during 2016–2019. Second, neither ingredient is orthogonal to initial informality: the pre-policy enrollment rate and the college contamination rate both decline significantly with pre-policy informality (Table A.6; figure A.6). On net, these correlations work in a conservative direction. Lower enrollment in high informality states means the intention-to-treat estimates understate per-student effects precisely where the identifying variation is largest. The contamination gradient requires a longer argument which the remainder of this subsection provides.

Contamination of the comparison group Let c_s denote the share of the comparison cohorts in state s that is partially treated through the public college channel, measured as the share of the 1995–1998 cohorts observed attending school (either public or private, undistinguishable in ENOE) during 2016–2019. It averages 39% and ranges from 28% to 56% across states. Two features of c_s bear on the estimates, and they pull in opposite directions.

The first is its *level*. Because the comparison group is itself partly treated, part of the schooling-linked effect is present on both sides of the cohort contrast and differences out, so the exposed-minus-unexposed gap understates the effect. This attenuates the estimated gradient, and attenuates it more where c_s is larger.

The second is its *gradient* with respect to informality. Contamination is not evenly distributed across states: regressing c_s on the pre-policy informality rate across the 32 states yields a slope of -0.21 percentage points of contamination per percentage point of informality (Table A.6; wild cluster bootstrap $p = 0.013$). Comparison groups in more informal states are less contaminated, and are therefore attenuated less than comparison groups in less informal states. Because the design reads the policy effect off the difference between high- and low-informality states, this differential attenuation makes the gradient appear steeper than it is. Where the level of contamination biases the estimate down, its gradient biases it up.

Which force dominates can be stated exactly, and Appendix D does so. Under the premise that motivates the design, the college effect scales with informality as the high-school effect does, and the estimate is attenuated whenever the state's contamination share exceeds a threshold that reaches 0.17 at the most informal state; the least contaminated state has a share of 0.28, so contamination attenuates in every state and the estimates read as a lower bound. If instead the college effect were a constant across states, only the uneven distribution of contamination would survive, and Appendix D shows it inflates the gradient by about one percent of the baseline. Finally, scaling by measured enrollment and public-school shares, the estimates imply effects per enrolled-eligible student of at least 1.7 times the intention-to-treat magnitudes.

4.5 Corroborating Evidence from High-School Graduates

I complement the main results with ENILEMS, a specialized survey of 18–20 year-old high-school graduates that is fully comparable with ENOE and, crucially, records whether each respondent attended a public or private high school. Only public-school students were eligible for the program, so exposure can be assigned at the cohort-by-school-type level rather than by birth cohort alone. This is a sharper exposure measure than the main design permits, and it provides an independent check on the participation response. I use the 2016 and 2019 waves, which bracket the policy; the cohorts surveyed in 2016 (born 1996–1998) were unexposed, while those surveyed in 2019 (born 1999–2001) were partially or fully exposed if they attended a public school. Table 8 presents summary statistics for both waves.

I estimate the analogue of the panel specification in Table 4 on cells defined by wave, state, cohort, and school type. Because exposure varies at the cohort-by-school-type level, the specification absorbs cohort-by-school-type fixed effects, which play the role that cohort fixed effects play in Table 4: they absorb the main effect of exposure, so that the coefficient on the interaction of exposure with pre-policy informality is identified from cross-state variation within cohort-by-school-type cells. State fixed effects absorb the main effect of informality, and the controls (state GDP, the share with a parent holding higher education, and the share that worked during high school) enter interacted with exposure, as in Table 4.[15]

Table 9 reports the estimates. The participation response is corroborated: the coefficient of 0.29 (wild bootstrap $p = 0.03$) is close to the panel estimate of 0.44 in Table 4, and its confidence interval contains it. Given the sharper exposure measure, the smaller sample, and the restriction to graduates, this is meaningful independent confirmation of the participation margin. The employment and income estimates carry the same signs as their ENOE counterparts but are imprecise.

Table 8: Labor Market Insertion Data (ENILEMS) - Summary Statistics

	(1) 2016 survey	(2) 2019 survey	(3) Difference
Female	0.54	0.53	-0.00 (-0.14)
Public high school	0.83	0.83	0.01 (0.71)
Had job during high school	0.16	0.20	0.04*** (4.61)
Continues in first job	0.35	0.41	0.06*** (3.96)
Participation rate	0.42	0.56	0.13*** (10.32)
Employment rate	0.89	0.92	0.03*** (2.76)
Formality rate	0.34	0.32	-0.02 (-1.09)
Real monthly income	3,066	3,765	699*** (6.25)
Attempted higher education	0.40	0.45	0.05*** (4.18)
Ever enrolled in higher education	0.26	0.34	0.08*** (6.77)
Either parent with higher education	0.21	0.23	0.02 (1.45)
Observations	7,191	7,799	

Note: t-statistic in parentheses. Survey weights used.
***p<0.01, **p<0.05, *p<0.1

[15] Each birth cohort appears in exactly one survey wave, so wave effects are subsumed by the cohort-by-school-type effects. Inference uses wild cluster bootstrap p-values with Webb weights and 9,999 replications, clustering by state.

Table 9: Labor Market Insertion Panel Estimates (ENILEMS), 2016–2019

	(1) Participation rate	(2) Employment rate	(3) Formality rate	(4) Monthly income	(5) Enrolled in higher ed
Pre-policy	0.29**	−0.18	−0.12	12.96	−0.15
informality	(0.13)	(0.16)	(0.20)	(12.89)	(0.16)
rate × Exposure	[0.027]	[0.285]	[0.550]	[0.342]	[0.356]
Cohort × school-type FE	✓	✓	✓	✓	✓
State FE	✓	✓	✓	✓	✓
Controls × Exposure	✓	✓	✓	✓	✓
Mean dependent variable	50.2	90.7	33.9	4,212.6	31.2
Observations	383	376	373	372	383

Note: Cells are defined by survey wave, state, birth cohort, and public/private high school. Exposure is the share of grades 10–12 under the policy for public-school graduates and zero for private-school graduates. Cohort-by-school-type fixed effects absorb the main effect of exposure; state fixed effects absorb the main effect of pre-policy informality. Controls (log state GDP, the share with a parent holding higher education, and the share that worked during high school) enter interacted with exposure. Cells are weighted by the square root of cell population (of the positive-income cell population in column (4)). Standard errors clustered by state in parentheses; wild cluster bootstrap p-values (Webb weights, 9,999 replications) in brackets; stars follow the bootstrap p-values. ***p<0.01, **p<0.05, *p<0.1

The formality estimate is small and negative, -0.12 with a standard error of 0.20, and does not reproduce the ENOE panel estimate of 1.28. Two features of the survey bound what that comparison can establish. The first is the thinness of the contrast: 373 cells and a single before-after difference for this outcome, against 5,091 cells over sixteen quarters in the panel. The second, and the more important, is that the population is not the same. ENILEMS observes only individuals who completed high school, at ages 18–20, among whom roughly 34% of the employed hold formal jobs, against roughly 8% among the employed 15–18 year-olds who drive the main estimates. These are different points of the early-career trajectory rather than the same margin measured twice. ENILEMS is therefore well suited to corroborating the participation response, which operates on a margin where its own population is still moving, and is not positioned to adjudicate the composition of employment among cohorts observed three years earlier; that margin is addressed directly in ENOE, where the population-based decomposition of Table 3 shows the formalization response in detail. I read the ENILEMS formality estimate as marking a scope condition on the main result rather than as independent evidence either for or against it.

One heterogeneity pattern is worth noting. Among exposed cohorts, states in which a larger share of students combined work with high school show a lower subsequent formality rate. This is the direction the benefit-substitution channel of Section 2 would predict for the study-and-work group, though the estimate is only marginally significant and I do not build on it.

In sum, ENILEMS delivers what a corroborating survey can: independent confirmation of the participation response under a cleaner exposure assignment, together with a scope condition on the formality result, which the main design locates among cohorts still at high-school ages rather than among graduates already observed at 18–20.

4.6 Mechanisms

Taken together, the evidence of this section makes it possible to distinguish among the three channels outlined in Section 2. The schooling channel predicts higher enrollment and delayed labor market entry; instead, enrollment does not respond (Table 6) and participation rises. Benefit substitution predicts a shift toward informal work, concentrated among the currently enrolled and dissipating at school exit; instead, informal employment falls as a share of the cohort population (Table 3), the participation and formality responses appear among youth both in and out of school, and the panel places the formality response, if anything, slightly more among those who have left school (1.15 against 1.00; Table A.7). The only trace of substitution is the marginally significant pattern among graduates who combined work and study, which is too fragile to build on. The salience and administrative-frictions channel predicts formal entry concentrated in wage employment, no response of self-employment, effects that scale with actual eligibility, and a response that does not require the coverage itself to be new to the youth; this is the configuration the data deliver: formal wage employment absorbs the entire formal-employment response while self-employment is unchanged on either margin (Table 3, Panel C), the coefficients rise under predicted exposure (Table A.6), and the formality gradient is as large for youths in households with a formal head (the group most likely to have had prior coverage through a parent and the group for whom the reform reduced to the assignment of a number) as for those with an informal head (0.96 against 0.92; Table A.7). The pattern is not confined to any one subgroup: participation and formality respond similarly for men and women and under an employees-only definition of informality (Figure 7), and ENILEMS corroborates the participation response under an independent, sharper exposure measure. On this reading, the policy operated less through the relative price of informal work than by lowering the administrative and informational costs of formal hiring at labor market entry.

5 Conclusion

Extending social insurance to students without requiring a formal job did not push young workers toward informal work. In the states where informality was highest, and where the subsidy concern is therefore sharpest, cohorts exposed to Mexico's 2016 school-linked expansion pulled further ahead of slightly older unexposed cohorts in formal employment than they did elsewhere, while informal employment fell. The response ran through wage employment, with no movement in self-employment on either margin, and it appears among youths both in and out of school.

Point estimates imply that a 1-percentage-point increase in initial informality is associated with a 0.8–1.3-percentage-point wider exposed-minus-unexposed cohort gap in the formality rate. Moving from a state at the mean of pre-policy informality to one a standard deviation higher is associated with a cohort formality gap narrower by 11.4–16.8 points at the baseline estimates, roughly a third to a half of the average gap of -32.8 points.[16]

The population-based decomposition (Table 3) shows how this occurs: the participation response flows entirely into employment, formal employment absorbs more than the full increase, informal employment declines, with no detectable change in unemployment. The evidence points to formal sector entry rather than a shift toward informality. I find no effects on school enrollment and only very small changes in income. Effects are similar for women and men, for students and non-students, and robust to excluding the self-employed. Participation estimates range from 0.15 to 0.59 depending on how flexibly cohort profiles are allowed to vary with initial informality, remaining positive and significant at conventional levels. Formality estimates remain large (between 0.78 and 1.21 under the same variations) and significant at the 1 percent level in every specification.

Taken together, the evidence is most consistent with a salience and administrative-frictions channel operating at the school-to-work transition: benefits become concrete, and a permanent social security number lowers hiring frictions. Consistent with a channel that works through formal hiring by firms, the effects are concentrated in subordinate wage employment, with no response of self-employment, and they are as large for youths already covered through a formal parent as for those who were not. I find no evidence that delinked benefits produced a shift toward informality where the concern is sharpest:

[16] The population-weighted standard deviation of pre-policy informality is 10.6 percentage points, which gives a scaled range of 9.2–13.6 percentage points.

informal employment declines, relative to less informal states, precisely where exposure was strongest. A weakly significant heterogeneity pattern among graduates who worked during high school is consistent with the substitution margin operating for that group. The results may be read as a lower bound on the schooling-linked impact under the assumption that public college eligibility generates effects in the same direction as high-school eligibility, since 39% of the comparison cohorts are observed attending school after the policy and this attenuates the cohort gap in every state; scaling by measured enrollment and public-school shares, the implied effects per enrolled-eligible student are at least 1.7 times the intention-to-treat magnitudes.

Four limitations bound the interpretation. First, the design identifies a differential and not a level: effects common to all states are absorbed. Second, the cohort contrast cannot separate net formalization from reallocation between cohorts within a state; if exposed cohorts displaced their older neighbors from formal wage jobs, the gap would widen without a net gain. The placebo cohorts show no such gradient among untreated cohort pairs, but the design does not rule the channel out directly. Third, cohorts are tracked by state of residence rather than state of schooling, so selective migration remains a residual concern, though it should equally affect the placebo cohorts. Fourth, the estimates pertain to the early-career margin over a pre-pandemic window, and the participation margin is materially more sensitive to the identifying assumption than the formality margin.

For policy design, the lesson is about delivery rather than delinking. The delinking debate asks whether benefits should be conditioned on formal employment; the Mexican experience suggests the more consequential question is how a benefit delivered outside a formal contract reaches its recipients. Here the entitlement had existed since 1998 and schools had already listed their students; what changed in 2016 was that coverage was attached to each student as a person, inside the contributory system, with the social security number that formal hiring requires, before their first job. Delivered that way, a benefit that could have subsidized informality worked instead as an on-ramp to formal employment. The same logic applies wherever social protection travels through a channel, a registry, an identifier, a payment account, that also lowers the cost of a formal transaction; whether it extends to benefits delivered through parallel systems, which leave the recipient's standing in the formal system unchanged, is the question the *Seguro Popular* literature has debated. A full welfare evaluation would integrate the labor market effects with the health benefits and longer-run outcomes; the college component and post-2019 dynamics, including cohorts enrolled after batch registration ended, are natural next steps. More broadly, advancing formalization will require understanding both workers' early-career decisions and firms' hiring margins in high-informality environments.

References

Akay, A. and Khamis, M. (2012). The Persistence of Informality: Evidence from Panel Data. *Informal Employment in Emerging and Transition Economies (Research in Labor Economics, Volume 34). Emerald Group Publishing Limited*, pages 229–255.

Akosa Antwi, Y., Moriya, A. S., and Simon, K. (2013). Effects of Federal Policy to Insure Young Adults: Evidence from the 2010 Affordable Care Act's Dependent-Coverage Mandate. *American Economic Journal: Economic Policy*, 5(4):1–28.

Alonso-Ortiz, J. and Leal, J. (2018). Cross-Subsidies, and the Elasticity of Informality to Social Expenditures: The Case of Mexico's Seguro Popular. *Review of Income and Wealth*, 64(2):482–512.

Alvarez, J. and Pizzinelli, C. (2021). COVID-19 and the Informality-driven Recovery: The Case of Colombia's Labor Market. *IMF Working Paper*.

Alzúa, M. L., Cruces, G., and Ripani, L. (2013). Welfare Programs and Labor Supply in Developing Countries: Experimental Evidence from Latin America. *Journal of Population Economics*, 26(4):1255–1284.

Arias, J., Azuara, O., Bernal, P., Heckman, J. J., and Villarreal, C. (2010). Policies to Promote Growth and Economic Efficiency in Mexico. Technical report, National Bureau of Economic Research.

Arulampalam, W. (2001). Is Unemployment Really Scarring? Effects of Unemployment Experiences on Wages. *The Economic Journal*, 111(475):F585–F606.

Ashworth, J., Hotz, V. J., Maurel, A., and Ransom, T. (2021). Changes across Cohorts in Wage Returns to Schooling and Early Work Experiences. *Journal of Labor Economics*, 39(4):931–964.

Atkin, D. (2016). Endogenous Skill Acquisition and Export Manufacturing in Mexico. *American Economic Review*, 106(8):2046–85.

Azuara, O. and Marinescu, I. (2013). Informality and the Expansion of Social Protection Programs: Evidence from Mexico. *Journal of Health Economics*, 32(5):938–950.

Bachas, P., Gadenne, L., and Jensen, A. (2023). Informality, Consumption Taxes, and Redistribution. *The Review of Economic Studies*, 91(5):2604–2634.

Bell, D. N. and Blanchflower, D. G. (2011). Young People and the Great Recession. *Oxford Review of Economic Policy*, 27(2):241–267.

Berniell, I., Gasparini, L., Marchionni, M., and Viollaz, M. (2023). Lucky Women in Unlucky Cohorts: Gender Differences in the Effects of Initial Labor Market Conditions in Latin America. *Journal of Development Economics*, 161:103042.

Bleakley, H. (2010). Malaria Eradication in the Americas: A Retrospective Analysis of Childhood Exposure. *American Economic Journal: Applied Economics*, 2(2):1–45.

Bosch, M. and Campos-Vazquez, R. M. (2014). The Trade-Offs of Welfare Policies in Labor Markets with Informal Jobs: The Case of the "Seguro Popular" Program in Mexico. *American Economic Journal: Economic Policy*, 6(4):71–99.

Breza, E. and Kaur, S. (2025). Labor Markets in Developing Countries. *Annual Review of Economics*, 17(Volume 17, 2025):747–776.

Brockmeyer, A., Gerard, F., Ulyssea, G., and Wu, L. (2025). Formal Labor Market Dynamics in Developing Countries. *Working Paper*.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90(3):414–427.

Campos-Vázquez, R. M. and Knox, M. A. (2013). Social Protection Programs and Employment: The Case of Mexico's Seguro Popular Program. *Economía Mexicana. Nueva Época*, 22(2):403–448.

Cruces, G., Ham, A., and Viollaz, M. (2012). Scarring Effects of Youth Unemployment and Informality: Evidence from Argentina and Brazil. *Working Paper*.

Donovan, K., Lu, W. J., and Schoellman, T. (2023). Labor Market Dynamics and Development. *The Quarterly Journal of Economics*, 138(4):2287–2325.

Duflo, E. (2001). Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment. *American Economic Review*, 91(4):795–813.

Esquivel, G. and Ordaz-Díaz, J. L. (2009). *¿Es Correcto Vincular la Política Social a la Informalidad en México?: Una Prueba Simple de las Premisas de esta Hipótesis*. CEPAL.

Gerard, F., Naritomi, J., and Silva, J. (2021). Cash Transfers and Formal Labor Markets: Evidence from Brazil. *CEPR Discussion Paper*.

Guner, N., Kulikova, Y. A., and Valladares-Esteban, A. (2025). Does the Added Worker Effect Matter? *Review of Economic Dynamics*, 56:101271.

Hotz, V. J., Xu, L. C., Tienda, M., and Ahituv, A. (2002). Are There Returns to the Wages of Young Men from Working While in School? *The Review of Economics and Statistics*, 84(2):221–236.

ILO (2018). Women and Men in the Informal Economy: A Statistical Picture. Technical Report 46, International Labour Organization.

Jensen, A. (2022). Employment Structure and the Rise of the Modern Tax System. *American Economic Review*, 112(1):213–34.

Le Barbanchon, T., Ubfal, D., and Araya, F. (2023). The Effects of Working While in School: Evidence from Employment Lotteries. *American Economic Journal: Applied Economics*, 15(1):383–410.

Levy, S. (2007). ¿Pueden los Programas Sociales Disminuir la Productividad y el Crecimiento Económico? Una Hipótesis para México. *El Trimestre Económico*, 74(295):491–540.

Levy, S. (2008). *Good Intentions, Bad Outcomes: Social Policy, Informality, and Economic Growth in Mexico*. Rowman & Littlefield.

Leyva, G. and Urrutia, C. (2023). Informal Labor Markets in Times of Pandemic. *Review of Economic Dynamics*, 47:158–185.

MacKinnon, J. G., Ørregaard Nielsen, M., and Webb, M. D. (2023). Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics*, 232(2):272–299.

Meghir, C. and Phillips, D. (2008). Labour Supply and Taxation. Technical report, IFS Working Papers.

Moffitt, R. A. (2002). Welfare Programs and Labor Supply. *Handbook of Public Economics*, 4:2393–2430.

Osuna-Gómez, D. (2023). The Impact of the COVID-19 Pandemic on Post-Great Recession Entrants: Evidence from Mexico. *Labour Economics*, 81:102337.

Rambachan, A. and Roth, J. (2023). A more credible approach to parallel trends. *The Review of Economic Studies*, 90(5):2555–2591.

Roodman, D., Nielsen, M. Ø., MacKinnon, J. G., and Webb, M. D. (2019). Fast and wild: Bootstrap inference in stata using boottest. *The Stata Journal*, 19(1):4–60.

Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3):305–22.

Ruhm, C. J. (1997). Is High School Employment Consumption or Investment? *Journal of Labor Economics*, 15(4):735–776.

Samaniego de la Parra, B. and Fernández Bujanda, L. (2024). Increasing the Cost of Informal Employment: Evidence from Mexico. *American Economic Journal: Applied Economics*, 16(1):377–411.

Schwandt, H. and von Wachter, T. (2019). Unlucky Cohorts: Estimating the Long-Term Effects of Entering the Labor Market in a Recession in Large Cross-Sectional Data Sets. *Journal of Labor Economics*, 37(S1):S161–S198.

Seira, E., Meza, I., González-Pier, E., and Alcaraz Prous, E. (2023). Did Mexico's Seguro Popular Universal Health Coverage Programme Really Reduce Formal Jobs? Technical Report 46, WIEGO.

Solon, G., Haider, S. J., and Wooldridge, J. M. (2015). What are we weighting for? *The Journal of Human Resources*, 50(2):301–316.

Ulyssea, G. (2020). Informality: Causes and Consequences for Development. *Annual Review of Economics*, 12(1):525–546.

Ulyssea, G., Bobba, M., Gadenne, L., and Harari, M. (2025). Informality. *VoxDevLit*, 6(2).

A Additional Material

Figure A.1: Cumulative students registered with a permanent social security number under the 2016 procedure

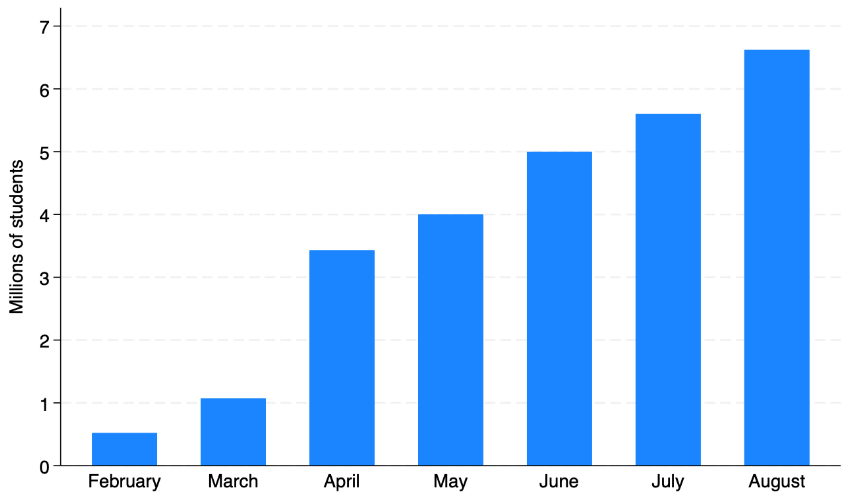

Source: IMSS. *Informe de Labores y Programa de Actividades 2016–2017*.

Figure A.2: Regions of Mexico

Source: *Regiones de México.* Northwest region includes Baja California, Baja California Sur, Chihuahua, Durango, Sinaloa and Sonora. Northeast region includes Coahuila, Nuevo León and Tamaulipas. West region includes Colima, Jalisco, Michoacán and Nayarit. East region includes Hidalgo, Puebla, Tlaxcala and Veracruz. Central north region includes Aguascalientes, Guanajuato, Querétaro, San Luís Potosí and Zacatecas. Central south region includes Mexico City, Mexico State and Morelos. South west region includes Chiapas, Guerrero, and Oaxaca. Southeast region includes Campeche, Quintana Roo, Tabasco and Yucatán.

Figure A.3: Long differences trends (difference between 15–18 and 21–24 cohorts by state)

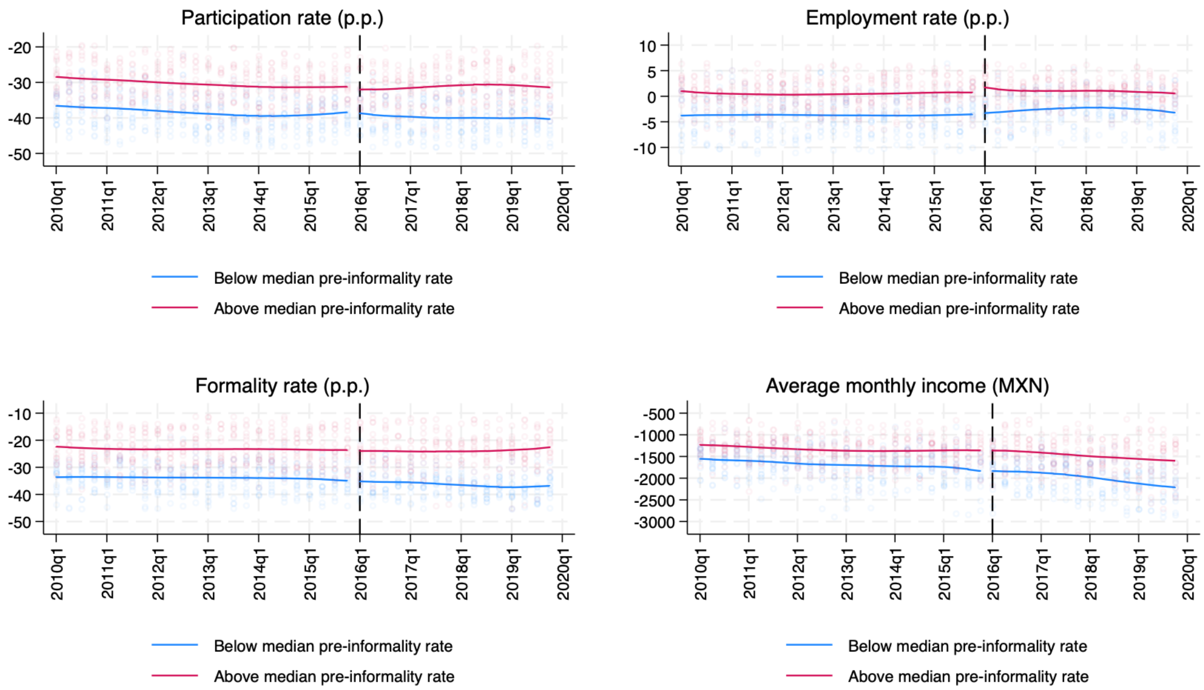

Note: Quarterly long differences are computed for each state as the outcome gap between ages 15–18 (treated) and 21–24 (control). States are split at the median pre-policy informality (2013–2015), and lines plot LOWESS-smoothed trends pre/post-2016. Table A.5 reports the corresponding formal check: baseline long-difference estimates conditioning on each state's cross-cohort gaps three years before the policy.
Source: Author's calculations with data from ENOE.

Figure A.4: Alternative definition of pre-policy informality rate by state (%), only employees
2013–2015 average

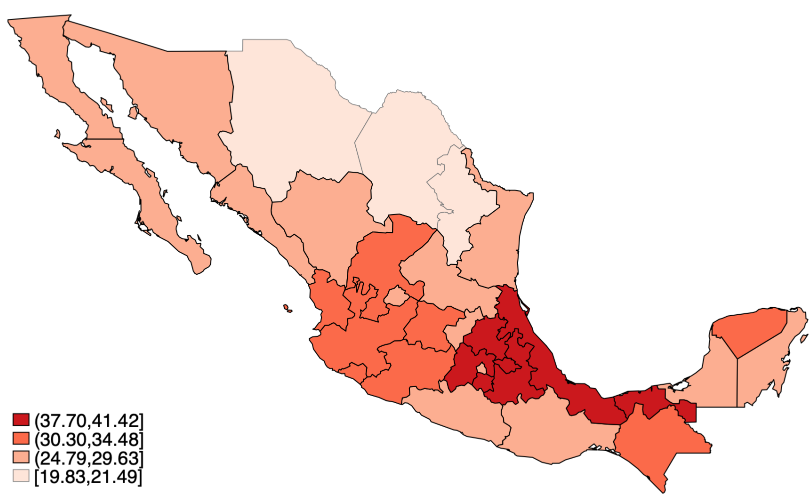

Source: Author's calculation using data from ENOE.

Figure A.5: Public share of upper-secondary enrollment, 2015–2016 (%)

Source: SEP. *Principales Cifras del Sistema Educativo Nacional 2015–2016.*

Figure A.6: Orthogonality of exposure ingredients to pre-policy informality

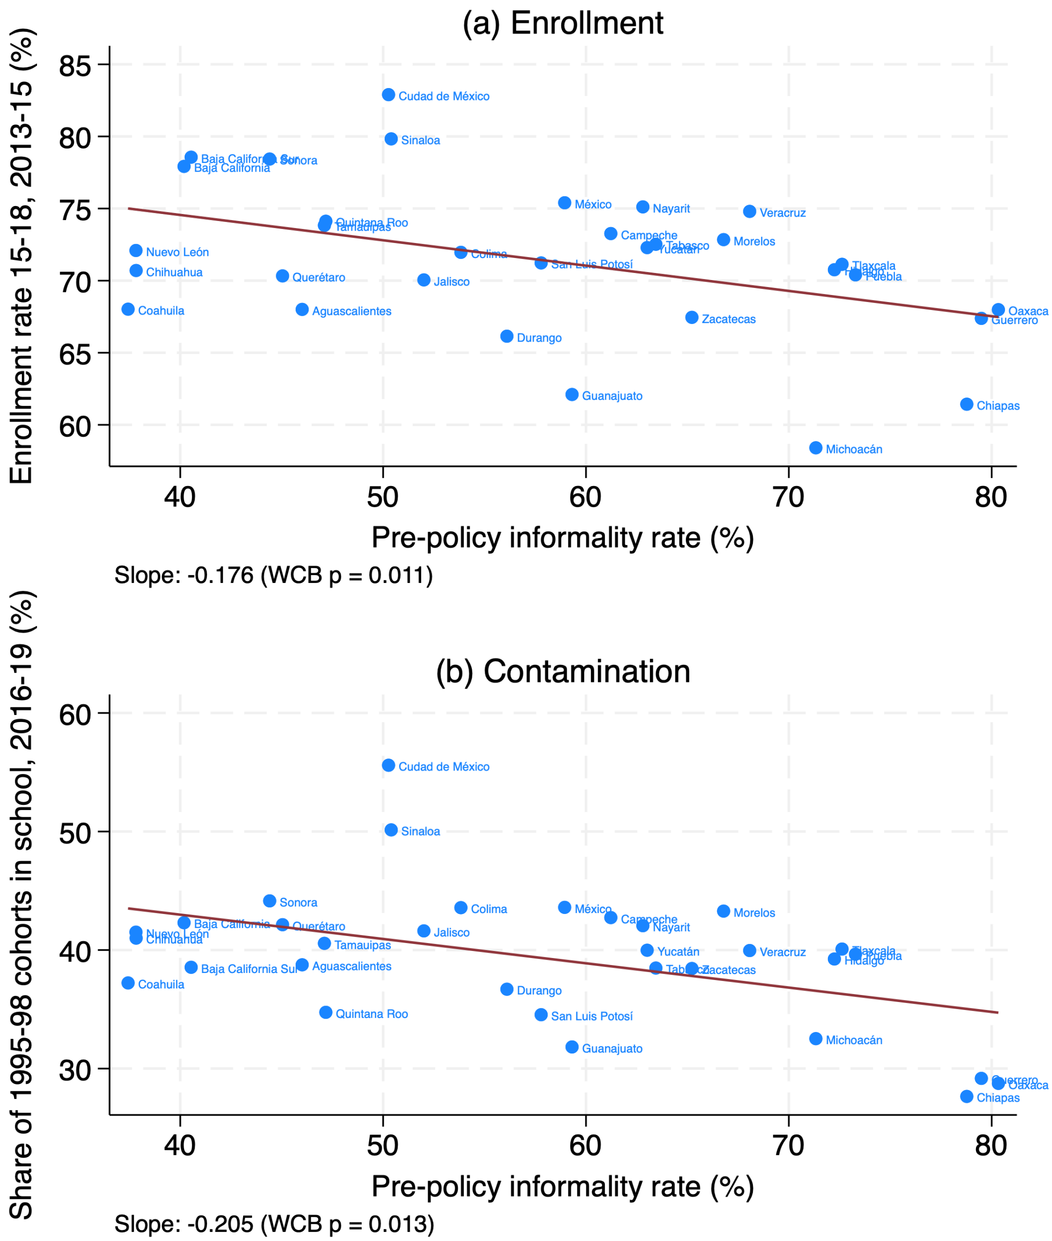

Note: Each point is a state. Panel (a) plots the pre-policy school-attendance rate of 15–18 year-olds (ENOE, 2013–2015) against the pre-policy informality rate; panel (b) plots the share of the 1995–1998 cohorts observed attending school during 2016–2019. Lines are unweighted linear fits; slopes and wild cluster bootstrap p-values are reported in Table A.6, Panel B.
Source: Author's calculations using data from ENOE.

Table A.1: Descriptive Statistics

	Long Differences (2019)		Panel Estimates (2016–2019)	
	(1) Exposed cohorts	(2) Unexposed cohorts	(3) Exposed cohorts	(4) Unexposed cohorts
Participation rate	0.23	0.61	0.21	0.53
Employment rate	0.93	0.93	0.95	0.92
Formality rate	0.08	0.41	0.07	0.35
Monthly income	3,501	5,384	2,906	4,602
School enrollment	0.77	0.31	0.79	0.40
Pre-policy informality rate	0.59	0.58	0.59	0.58
Log GDP per capita in 2013	11.68	11.72	11.67	11.70
Educational attainment in 2015	9.10	9.18	9.10	9.16
Age	16.05	22.02	15.60	20.48
Men	0.50	0.50	0.49	0.50
Formal household head	0.48	0.50	0.47	0.49
Northwest region	0.13	0.13	0.13	0.14
Northeast region	0.10	0.10	0.09	0.10
West region	0.12	0.12	0.12	0.12
East region	0.15	0.15	0.15	0.15
Central north region	0.12	0.12	0.12	0.12
Central south region	0.21	0.22	0.21	0.22
Southwest region	0.11	0.09	0.11	0.10
Southeast region	0.06	0.06	0.06	0.06
Observations	128	128	3,046	2,048

Note: Mean values for each variable are reported, weighted by survey weights. Columns (1) and (2) correspond to the long differences estimates reported in Table 2. Columns (3) and (4) correspond to the panel estimates reported in Table 4. For the long differences estimates, exposed cohorts only include fully exposed cohorts; for the panel estimates exposed cohorts include partially exposed cohorts as defined by the exposure measure plotted in figure 1. Rates (participation, employment, formality, school enrollment, pre-policy informality) are displayed as proportions (0–1) in this table. In all regression specifications the pre-policy informality rate and labor market outcome rates are scaled as percentage points (0–100); regression coefficients are therefore interpreted as the effect of a 1-percentage-point change in the respective variable. Observations are cohort-state-quarter cells; columns (3) and (4) count cells with a non-missing participation rate over 2016–2019 and sum to the 5,094 observations of Table 4, column (1). The corresponding counts are 3,043 and 2,048 for the employment and formality rates and 2,998 and 2,048 for monthly income.

Table A.2: Robustness of Inference and Leverage

	Long differences		Panel	
	Participation	Formality	Participation	Formality
Coefficient	0.59	0.87	0.44	1.28
Panel A: standard errors / p-values				
Cell bootstrap	(0.07)	(0.11)	—	—
Cluster-robust, state (CRV1)	(0.09)	(0.16)	(0.05)	(0.06)
Block bootstrap, states	(0.14)	(0.21)	—	—
Cluster jackknife (CV3)	(0.12)	(0.21)	(0.07)	(0.06)
Wild cluster bootstrap p-value	$<$0.001	0.003	$<$0.001	$<$0.001
Wild bootstrap 95% CI	[0.40, 0.81]	[0.44, 1.24]	[0.30, 0.56]	[1.16, 1.45]
State-collapsed (32 obs.)	0.59 (0.10)	0.87 (0.18)	—	—
Panel B: leverage				
Leave-one-state-out range	[0.56, 0.67]	[0.79, 0.98]	[0.40, 0.48]	[1.24, 1.31]
Excluding Chiapas, Guerrero and Oaxaca	0.61	0.88	0.46	1.29
Panel C: weighting of cells				
$\sqrt{\text{population}}$ (baseline)	0.59	0.87	0.44	1.28
Population	0.59	0.91	0.45	1.29
Unweighted	0.60	0.85	0.43	1.27

Note: All rows refer to the coefficient on pre-policy informality (long differences, Table 2 Panel B specification) or on pre-policy informality interacted with exposure (panel, Table 4 specification). Point estimates are identical across inference methods by construction. Wild cluster bootstrap uses Webb weights with $9,999$ replications, clustering by state (Roodman et al., 2019); CV3 jackknife standard errors follow MacKinnon et al. (2023). The state-collapsed specification averages the four 2019 quarters within each state (population-weighted) and drops quarter fixed effects. Panel C reports coefficients under alternative regression weights; all remain significant at the 1 percent level by wild cluster bootstrap.

Table A.3: Panel Estimates Robustness to State Linear Trends, 2016–2019

	(1) Participation rate	(2) Employment rate	(3) Formality rate	(4) Monthly income
Pre-policy informality	0.45***	0.15***	1.29***	24.52***
rate × Exposure	(0.06)	(0.03)	(0.06)	(8.29)
Log GDP per	2.05	0.44	3.73**	-64.72
capita in 2013 × Exposure	(1.05)	(1.08)	(1.05)	(131.27)
Educational attainment	-0.32	1.85**	1.92	-58.32
in 2015 × Exposure	(0.71)	(0.58)	(1.29)	(101.38)
Cohort FE	✓	✓	✓	✓
State FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
State linear time trend	✓	✓	✓	✓
Mean dependent variable	33.31	93.84	17.99	4,299.21
Observations	5,094	5,091	5,091	5,046

Note: Standard errors clustered by state in parentheses. The square root of cell sizes are used as weights for the observations. Cell sizes are identical across specifications, except for the income one as I restrict to individuals reporting strictly positive income.
***p<0.01, **p<0.05, *p<0.1

Table A.4: Robustness to Informality-Specific Cohort Trends and Age Profiles

	(1) Baseline 1995–2004	(2) Baseline 1991–2004	(3) + linear trend	(4) + quadratic trend	(5) + age profile
Panel A: participation rate					
Informality × Exposure	0.44***	0.39***	0.15*	0.22***	0.15**
	(0.05)	(0.04)	(0.07)	(0.07)	(0.06)
Informality × cohort trend			0.03***	0.04***	
			(0.01)	(0.01)	
Informality × cohort trend^2				−0.002**	
				(0.001)	
Informality × age					−0.03***
					(0.01)
WCB p-value (exposure)	<0.001	<0.001	0.052	0.006	0.039
Observations	5,094	7,142	7,142	7,142	7,142
Panel B: formality rate					
Informality × Exposure	1.28***	1.21***	1.09***	0.78***	0.93***
	(0.06)	(0.06)	(0.08)	(0.07)	(0.07)
Informality × cohort trend			0.02	−0.04**	
			(0.01)	(0.02)	
Informality × cohort trend^2				0.007***	
				(0.001)	
Informality × age					−0.04**
					(0.01)
WCB p-value (exposure)	<0.001	<0.001	<0.001	<0.001	<0.001
Observations	5,091	7,139	7,139	7,139	7,139

Note: All columns estimate the Table 4 panel specification (cohort, state, and quarter fixed effects; GDP and education controls interacted with exposure; square-root-of-cell-population weights; 2016–2019). Column (1) uses the Table 4 cohort range (1995–2004); columns (2)–(5) extend the sample to cohorts 1991–2004, adding four untreated cohorts that help separate smooth cohort profiles from the exposure step. Columns (3)–(5) add interactions of pre-policy informality with a linear cohort trend, a quadratic cohort trend, and age (year minus cohort), respectively, so the exposure coefficient is identified from the discrete exposure ramp net of smooth informality-specific profiles. Because exposure ramps from zero to one over four cohorts, flexible trends absorb part of the ramp by construction; these estimates are therefore conservative bounds. Standard errors clustered by state in parentheses; wild cluster bootstrap p-values for the exposure coefficient reported (Webb weights, 9,999 replications). ***p<0.01, **p<0.05, *p<0.1

Table A.5: Long Differences Conditioning on Pre-Policy Cohort Gaps, 2019

	(1) Participation rate	(2) Employment rate	(3) Formality rate	(4) Monthly income
Panel A				
Pre-policy informality rate	0.40***	0.10***	0.52***	28.09***
	(0.06)	(0.02)	(0.07)	(4.14)
	[<0.001]	[0.001]	[<0.001]	[<0.001]
Cohort gap in 2013	0.24	0.31***	0.46***	0.10
(same outcome)	(0.15)	(0.08)	(0.11)	(0.08)
Observations	128	128	128	128
Panel B				
Pre-policy informality rate	0.63***	0.12*	0.58**	11.02
	(0.10)	(0.05)	(0.18)	(9.98)
	[<0.001]	[0.072]	[0.017]	[0.280]
Cohort gap in 2013	-0.08	0.24**	0.42***	-0.01
(same outcome)	(0.12)	(0.09)	(0.11)	(0.08)
Log GDP per	5.98***	1.47*	0.63	-137.07
capita in 2013	(0.88)	(0.86)	(1.25)	(128.31)
Educational attainment	-3.16***	-0.24	1.70	-333.14**
in 2015	(1.07)	(0.69)	(1.55)	(142.36)
Regional controls	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Observations	128	128	128	128

Note: Each column reports the Table 2 long-difference specification augmented with the state's cross-cohort gap in the same outcome measured in 2013, three years before the start of the policy. The 2013 gap is the difference between the 15–18 and 21–24 age groups by state and quarter (the construction underlying figure A.3), matched to each 2019 cell by state and quarter of the year. Standard errors clustered by state in parentheses; wild cluster bootstrap p-values for the pre-policy informality coefficient in brackets (Webb weights, 9,999 replications). Cells are weighted by the square root of the cell population (the square root of the positive-income cell population in column (4), which restricts to individuals reporting strictly positive income).
***p<0.01, **p<0.05, *p<0.1

Table A.6: Predicted Exposure and Exposure Ingredients

Panel A: panel estimates under predicted exposure						
	Formality rate			Participation rate		
	(1) Baseline	(2) Enrollment	(3) Enroll. $\times$ public	(4) Baseline	(5) Enrollment	(6) Enroll. $\times$ public
Informality $\times$ exposure measure	1.28***	1.50***	1.59***	0.44***	0.51***	0.56***
	(0.06)	(0.07)	(0.12)	(0.05)	(0.05)	(0.08)
	[<0.001]	[<0.001]	[<0.001]	[<0.001]	[<0.001]	[<0.001]
Observations	5,091	5,091	5,091	5,094	5,094	5,094

Panel B: orthogonality of exposure ingredients to pre-policy informality		
	Enrollment rate 15–18, 2013–15	Contamination 1995–98 cohorts
Pre-policy informality rate	-0.18**	-0.21**
	(0.06)	(0.07)
	[0.011]	[0.013]
R-squared	0.19	0.22
Observations	32	32

Note: Panel A re-estimates the Table 4 panel specification replacing cohort exposure with predicted exposure: cohort exposure scaled by the state's pre-policy school-attendance rate of 15–18 year-olds (columns (2) and (5); ENOE 2013–2015, mean 71.5 percent) and additionally by the public share of upper-secondary enrollment (columns (3) and (6); SEP, *Principales Cifras del Sistema Educativo Nacional 2015–2016*, mean 81 percent). Controls are interacted with the same exposure measure as the treatment variable in each column. Panel B regresses each exposure ingredient on the pre-policy informality rate across the 32 states; contamination is the share of the 1995–1998 cohorts observed attending school during 2016–2019. Standard errors clustered by state in parentheses (robust in Panel B); wild cluster bootstrap p-values in brackets (Webb weights, 9,999 replications). The implied exposure differential between designated exposed and unexposed cohorts is at most 0.58 (enrollment $\times$ public share), so effects per enrolled-eligible student are at least 1.7 times the intention-to-treat estimates, and larger once the comparison group's partial treatment is counted. ***p<0.01, **p<0.05, *p<0.1

Table A.7: Heterogeneity and Robustness Estimates

	Participation rate	Employment rate	Formality rate	Monthly income
Panel A: Long Differences (2019)				
Full sample	0.59***	0.15*	0.87***	11.18
	(0.09)	(0.06)	(0.16)	(10.00)
	[<0.001]	[0.053]	[0.003]	[0.278]
Not in school	0.37**	0.10	0.83***	18.86
	(0.13)	(0.07)	(0.17)	(12.54)
	[0.039]	[0.255]	[0.007]	[0.155]
In school	0.56***	0.20*	0.95***	29.16
	(0.08)	(0.09)	(0.19)	(15.33)
	[<0.001]	[0.076]	[0.001]	[0.115]
Men	0.90***	0.12	0.97***	−3.34
	(0.16)	(0.08)	(0.15)	(13.24)
	[<0.001]	[0.227]	[<0.001]	[0.824]
Women	0.27**	0.15	0.70**	39.15**
	(0.10)	(0.08)	(0.21)	(13.01)
	[0.024]	[0.119]	[0.029]	[0.018]
Employees only	0.43*	0.28	1.10*	−1.96
	(0.21)	(0.14)	(0.44)	(18.13)
	[0.094]	[0.221]	[0.088]	[0.922]
Panel B: Panel Estimates (2016–2019)				
Full sample	0.44***	0.15***	1.28***	21.33**
	(0.05)	(0.03)	(0.06)	(7.05)
	[<0.001]	[0.001]	[<0.001]	[0.021]
Not in school	0.39***	0.18***	1.15***	18.87**
	(0.08)	(0.04)	(0.08)	(6.90)
	[0.005]	[<0.001]	[<0.001]	[0.036]
In school	0.51***	0.13***	1.00***	28.84***
	(0.06)	(0.03)	(0.12)	(5.75)
	[<0.001]	[0.007]	[<0.001]	[0.002]
Men	0.47***	0.12***	1.27***	18.12*
	(0.07)	(0.03)	(0.07)	(7.46)
	[<0.001]	[0.002]	[<0.001]	[0.051]
Women	0.40***	0.22***	1.30***	30.10***
	(0.05)	(0.05)	(0.07)	(6.97)
	[<0.001]	[0.002]	[<0.001]	[<0.001]
Informal hh head	0.36***	0.07**	0.92***	24.85***
	(0.06)	(0.03)	(0.06)	(6.42)
	[<0.001]	[0.022]	[<0.001]	[0.001]
Formal hh head	0.59***	0.19***	0.96***	1.27
	(0.06)	(0.04)	(0.06)	(8.39)
	[<0.001]	[<0.001]	[<0.001]	[0.897]
Employees only	0.46***	0.21***	1.27***	21.73**
	(0.10)	(0.06)	(0.36)	(8.86)
	[<0.001]	[0.005]	[0.008]	[0.024]

Note: Panel A reports the Table 2 long-difference specification; Panel B reports the Table 4 panel specification. Both are re-estimated on each subsample described in the row label (the household-head split is estimated only in the panel specification, hence no corresponding Panel A rows). In the "Employees only" rows both the outcome samples and the pre-policy informality regressor are restricted to subordinate employees (figure A.4), so magnitudes are not directly comparable to the other rows. "Full sample" rows reproduce the corresponding coefficient from Tables 2 and 4. Standard errors clustered by state in parentheses; wild-cluster-bootstrap p-values for the pre-policy informality coefficient in brackets (Webb weights, 9,999 replications), matching the convention used throughout the paper (e.g. Table 2). Stars are based on the wild-bootstrap p-values.
***p<0.01, **p<0.05, *p<0.1

B Operational Details of the School-Linked Social Insurance Expansion

This appendix documents the legal basis, the 2016 procedural change, the financing, and the scale of the program, and notes the implications for measurement.

B.1 Legal basis and scope of the benefit

Students of public upper-secondary and higher-education institutions were incorporated into the compulsory regime of social insurance, for the in-kind benefits of the Sickness and Maternity Insurance only, by a presidential agreement of June 1987, abrogated and replaced by the decree of September 1998 that remains in force.[17] Three features of the entitlement matter for the interpretation of the estimates.

- **Benefits are in kind only.** Coverage comprises medical, surgical, pharmaceutical and hospital care, obstetric care, and preventive programs. It carries no cash benefits, no retirement or housing contributions, and no work-risk coverage, so it does not replicate the full bundle attached to a formal job.

- **Coverage is individual.** The affiliation covers the student alone and extends to no dependents.

- **Eligibility is residual.** The decree covers students who do not hold "the same or similar protection" from IMSS or any other social security institution, and the 2015 rules require each student to declare under oath that this is so before the school may enroll them.[18] Coverage lapses when the student leaves the institution, and it also lapses if the student takes a formal job, since employment-based affiliation takes precedence in IMSS records. The social security number, by contrast, is retained for life.[19]

B.2 What changed in 2016

On October 28, 2015 the IMSS Technical Council approved the rules governing student incorporation, published in the *Diario Oficial de la Federación* on December 16, 2015 and

[17] *Diario Oficial de la Federación*, June 10, 1987, and September 14, 1998.

[18] Acuerdo ACDO.SA1.HCT.281015/246.P.DIR, Anexo Único, Reglas 1 and 6.1, *Diario Oficial de la Federación*, December 16, 2015.

[19] IMSS, *Seguro de Salud para Estudiantes*: "Lo primero que debes hacer es obtener tu Número de Seguridad Social (NSS), que es para toda la vida."

in force the following day.[20] The rules changed the identifier attached to the student, the party responsible for enrollment, and the channel through which enrollment travels.

- **From a school-scoped number to a personal one.** The rules distinguish two identifiers. The *conventional* social security number is the one IMSS assigned to students under the previous practice: a number valid for the period of study which, in the language of the rules, students "may not use again for any other type of insurance within the Institute" (Rule 2, fracc. XII). The *ordinary* social security number is "unique, permanent and non-transferable" and is assigned to each applicant (Rule 2, fracc. XI). From 2016, students are enrolled under the ordinary number. The pre-2016 identifier was therefore, by construction, unusable for employment-based affiliation; the post-2016 identifier is the same number the student will use when first hired formally.

- **Schools became the obligated party.** Public institutions must register with IMSS to obtain an ordinary employer registration number in modalidad 32, enroll and de-enroll their students, and keep a roster containing each student's full name, ordinary social security number, and CURP (Rules 5.1–5.2).

- **Enrollment moved to an electronic channel with a deadline.** Institutions transmit enrollment and de-enrollment notices through IDSE (*IMSS Desde Su Empresa*), the employers' electronic filing system, signed with the legal representative's advanced electronic signature, within no more than five business days of a student acquiring or losing student status (Rules 7.1–7.3, 8.2).

- **Obtaining the number was made close to costless for the student.** A student obtains or locates the ordinary number through the IMSS web portal using only a CURP and an email address, or in person at a subdelegación (Rule 6.2). The student then supplies it to the school together with the sworn declaration of Rule 6.1.

- **The transition had a hard deadline.** Institutions were required to register and transmit affiliation movements from January 2016 (Rule 8.4), and after the first semester of 2016 IMSS would de-register institutions and their students that had not complied (Rule 8.5). Subdelegaciones were instructed to inform and assist registered institutions in implementing the rules.

[20] Acuerdo ACDO.SA1.HCT.281015/246.P.DIR and Anexo Único, *Diario Oficial de la Federación*, December 16, 2015. All rule numbers cited in this subsection refer to the Anexo Único.

The 2016 campaign described in Section 2.1 operated within this window. Two consequences of Rule 8.5 are worth noting: it explains why the rollout was compressed into a single semester, and it is consistent with the fall in year-end registered affiliations from 7.4 to 6.8 million between 2015 and 2016, which IMSS attributes to improved registration.[21]

B.3 Financing

The federal government pays the entire contribution; neither the student nor the school pays anything, and neither faces a marginal cost from an additional registration. The contribution is a fixed daily amount per registered student, set in 1987 as 1.723 percent of the daily minimum wage of Mexico City and, from 2017, of the *Unidad de Medida y Actualización*.[22] The rate has not changed since 1987. In 2019 it amounted to 1.47 pesos per day, or 524 pesos per student per year; IMSS estimates that its spending on this population exceeded the corresponding revenue by a factor of about 2.5 in that year, and that the 1989 contribution of 63 pesos would be 893 rather than 524 pesos had it been indexed to inflation.[23] The program is therefore financed centrally and its cost is invisible at the point of enrollment, which is why school-level registration decisions were driven by administrative compliance rather than by budget.

B.4 Scale and utilization

Year-end affiliations in modalidad 32 rose from 5.0 million in 2010 to 7.4 million in 2015, fell to 6.8 million in 2016 as records were regularized, and stood at 7.8 million in 2019.[24] In 2019 this population received 2.4 million medical consultations, 71 percent of them at the primary level, and generated 63 thousand hospital discharges, of which obstetric care accounted for 45 percent.[25] Access carries no exclusions or waiting periods, including for chronic conditions.

[21] IMSS, *Informe al Ejecutivo Federal y al Congreso de la Unión sobre la Situación Financiera y los Riesgos del IMSS 2019–2020*, gráfica II.22 and note 34.

[22] Consejo Técnico agreements 1041/87 and 1313/87; the change of base follows Acuerdo ACDO.SA2.HCT.250117/26.P.DJ of January 25, 2017.

[23] IMSS, *Informe 2019–2020*, section II.2.2 and gráfica II.23.

[24] IMSS, *Informe 2019–2020*, gráfica II.22. The 2019 figure includes beneficiaries of *Jóvenes Construyendo el Futuro*, who are affiliated in this modality. IMSS reports that about 98 percent of modalidad-32 affiliations correspond to students.

[25] IMSS, *Informe 2019–2020*, section II.2.2. The three leading reasons for consultation were factors influencing health status and contact with health services (20 percent), respiratory conditions (13 percent), and digestive conditions (13 percent).

B.5 Implications for measurement

Four features of the administrative design bear on the empirical strategy.

1. **Administrative counts cannot serve as the outcome.** Student affiliation and job-based affiliation are mutually exclusive in IMSS records, so a student who takes a formal job leaves modalidad 32 and appears in modalidad 10. A rise in formal employment among students mechanically lowers modalidad-32 counts. Formality is therefore measured throughout from ENOE, using INEGI's employment-based definition, and not from IMSS registers.

2. **A pre-policy administrative first stage by cohort is not constructible.** Before 2017, IMSS coded every modalidad-32 affiliation in the "age not available" category, because registration under the conventional number did not require the student's date of birth; age is reported for this modality only from 2017, a change IMSS attributes to the 2016 assignment of ordinary numbers.[26] The exposure measure of Section 3.1 is therefore built from survey data on cohorts rather than from administrative counts.

3. **The residual-eligibility rule shapes the household-head split.** Because students holding equivalent protection were required to declare as much and were not enrolled, youths already covered as dependents of a formally employed parent were largely outside the affiliation margin. What the 2016 procedure extended to them was the request, made by their school, to obtain a permanent social security number in their own name. Section 4.2 reads the similarity of the formality response across the two groups in this light.

4. **The five-day rule and the January 2016 start compressed the rollout** into a window too short to exploit as staggered adoption at ENOE's quarterly frequency, as discussed in Section 3.1.

[26] IMSS, *Glosario de datos abiertos, asegurados*, "Rango de edad".

Figure B.1: "Tienes IMSS" campaign flyer

Source: IMSS. Official Twitter account (April 18, 2016).

Figure B.2: Political event announcing the affiliation of students to IMSS in Mexico City

Source: IMSS. *El Seguro Médico para estudiantes es un escudo para prevenir enfermedades crónicas: Mikel Arriola.*

C Bounding Violations of Parallel Cohort Profiles

This appendix describes the sensitivity analysis reported at the end of Section 3.3. The exercise applies the framework of Rambachan and Roth (2023) to the cohort-space design of this paper. Its purpose is to replace a binary test of the identifying assumption with a quantitative statement about how much of that assumption the results actually require.

C.1 Motivation

Section 3.3 documents that the informality gradients of untreated cohorts lie in a narrow band but are not exactly equal, and that the drift across those cohorts runs opposite in sign to the estimated policy effect. The conventional response is to test the null that the untreated gradients are equal and, having failed to reject, proceed as though parallel cohort profiles held exactly. That response is unavailable here in any case, since the test rejects ($F(7, 31) = 10.8$ for formality and 7.4 for participation, both $p < 0.001$), and it would be unattractive even if it did not, for two reasons. First, with 32 clusters the test has limited power: failing to reject a null of no differential drift is weak evidence that none exists. Second, conditioning the analysis on the outcome of a pre-test changes the sampling distribution of the estimator, so that reported confidence intervals no longer have their nominal coverage, and the distortion is anti-conservative (Roth, 2022). The relevant question is therefore not whether the assumption holds exactly, but how far it can fail before the conclusions change.

C.2 The Design as an Event Study in Cohort Space

The fully interacted specification of Section 3.3,

$$Y_{kst} = \alpha_k + \beta_k I_{s,pre} + X_s \Gamma_k + \delta_{kt} + \nu_{skt}, \tag{1}$$

is an event study in which the birth cohort plays the role of event time. Older cohorts stand further before treatment, younger cohorts after it, and the exposure schedule of Figure 1 supplies the analogue of an event date. Joint estimation over cohorts is what makes the exercise feasible: it yields a single state-clustered covariance matrix across the $\hat{\beta}_k$, which is the object the sensitivity analysis requires.

Because equation (1) is estimated without an omitted cohort, the $\hat{\beta}_k$ are levels of a gradient rather than deviations, and for untreated cohorts they lie well away from zero. The event-study normalization is therefore

$$\theta_k \equiv \beta_k - \beta_{k_0}, \tag{2}$$

where k_0 denotes the reference cohort. Under the identifying assumption $\theta_k = 0$ for every untreated cohort, while θ_k measures the policy effect for exposed cohorts. The level β_{k_0} absorbs permanent cross-state differences in the age profile of formality and carries no information about the policy. Since (2) is linear, the covariance of $\hat{\theta}$ follows exactly as $A\hat{V}_\beta A'$, with A placing a one in the own-cohort column and a minus one in the reference column; no approximation is involved, and the reference cohort can be varied without re-estimation.

I take the 1997 cohort as the reference. The 1998 cohort would be the mechanical choice as the last cohort with zero high-school exposure, but as noted in Section 3.3 its gradient shows an uptick relative to 1995–1997 that is consistent with partial exposure through the public college channel, which reaches the youngest untreated cohorts first, and/or misalignment between age and school-year. Normalizing on a partially treated cohort both understates the effect and inflates the apparent pre-period violation. With 1997 as the reference, the 1998 cohort moves into the post-reference set, so the largest second difference among the remaining untreated cohorts falls from 0.24 to 0.11 for formality, confirming that the 1998 gradient was its source. Results under the 1998 reference are reported in Table C.1 for completeness; they are uniformly weaker, as this reasoning implies.

C.3 Restricting the Counterfactual Cohort Profile

Write the vector of normalized gradients as $\theta = \tau + \delta$, where τ is the policy effect and δ the violation of parallel cohort profiles. Since $\tau_k = 0$ for untreated cohorts, the untreated $\hat{\theta}_k$ estimate δ directly and contain all of the available information about it; the identification problem is that δ is unrestricted for exposed cohorts. In this setting δ has a concrete interpretation: it is the extent to which the cross-state informality gradient would have drifted across birth cohorts absent the policy, which is precisely the object generated by the two threats named at the start of Section 3.3: convergence and heterogeneous cohort-age profiles.

Rambachan and Roth (2023) propose restricting δ to a convex set Δ and reporting confidence sets for the parameter of interest that are valid for every $\delta \in \Delta$. I use their two leading choices. The smoothness class

$$\Delta^{SD}(M) = \{\delta : |(\delta_{k+1} - \delta_k) - (\delta_k - \delta_{k-1})| \leq M \ \ \forall k\}$$

bounds the curvature of the counterfactual cohort profile; $M = 0$ forces it to be exactly linear, so the untreated cohorts pin down its slope and it is extrapolated forward. The relative-magnitudes class

$$\Delta^{RM}(\bar{M}) = \left\{\delta : |\delta_{k+1} - \delta_k| \leq \bar{M} \cdot \max_{j \text{ untreated}} |\delta_{j+1} - \delta_j|\right\}$$

imposes no smoothness at all, requiring only that the profile not move more sharply after 1997 than it does among untreated cohorts by a factor of $\bar{M}$. The summary statistic in each case is the breakdown value, the smallest M or $\bar{M}$ at which the robust confidence set first includes zero.

These classes generalize the parametric restrictions already imposed in Table A.4. Column (3) of that table adds an informality-specific linear cohort trend and column (4) a quadratic; $\Delta^{SD}(0)$ and $\Delta^{SD}(M)$ are the nonparametric counterparts. The generalization is not merely cosmetic. Because exposure ramps over four cohorts, a trend estimated jointly over treated and untreated cohorts absorbs part of the ramp, which is why Table A.4 describes its estimates as conservative bounds. Under $\Delta^{SD}(0)$ the trend is identified from untreated cohorts alone and cannot absorb the ramp. The two therefore move in opposite directions: column (3) reports 1.09 for formality against a baseline of 1.28, whereas $\Delta^{SD}(0)$ shifts the interval to $[1.26, 1.62]$. This confirms that the attenuation in Table A.4 reflects ramp absorption rather than a genuine differential trend.

C.4 Target Parameter

The procedure delivers inference on a scalar $\theta = \ell'\tau_{post}$, so ℓ must be specified. I report as the baseline the simple average over the 2001 and 2002 cohorts, the fully exposed cohorts observed over a working-age window in 2016–2019 (ages 15–18 and 14–17 respectively). This is the direct analogue of the long-difference contrast in Table 2: it compares fully exposed cohorts with an unexposed reference cohort in the same units.

Extending the average to all fully exposed cohorts (2001–2004) is reported as an alternative. It raises the point estimate but lowers the breakdown value, and I do not use it as the baseline. The 2003 and 2004 cohorts are observed at ages 13–16 and 12–15, largely below the legal working age of fifteen, where formal employment among the employed is near zero in every state and the cross-state gradient is mechanically compressed toward zero. Their normalized gradients are correspondingly large (1.50 and 1.52) and continue to rise after exposure has already saturated at the 2001 cohort, which a constant intention-to-treat effect would not generate. Restricting the target to 2001–2002 avoids attributing this compression to the policy.

C.5 Inference

The framework requires an estimate of the covariance matrix of $\hat{\theta}$ and cannot accommodate the wild cluster bootstrap used for the main results. With 32 clusters, CRV1 standard errors are optimistic. I therefore construct a cluster-jackknife (CV3) covariance directly,

$$\hat{V}^{CV3} = \frac{G-1}{G} \sum_{g=1}^{G} \left(\hat{\beta}^{(-g)} - \bar{\beta}\right) \left(\hat{\beta}^{(-g)} - \bar{\beta}\right)',$$

re-estimating equation (1) $G = 32$ times, each omitting one state, and transforming to $\hat{V}_{\theta}^{CV3}$ as above (MacKinnon et al., 2023). This is the covariance analogue of the resampling the wild bootstrap performs, and it is consistent with the CV3 standard errors already reported in Table A.2. Standard errors are 10 to 56 percent larger than under CRV1 depending on the cohort. Results under both are reported.

Confidence sets under Δ^{SD} use the fixed-length construction of Rambachan and Roth (2023); those under Δ^{RM} use their conditional least-favorable hybrid, since the relative-magnitudes class is a union of polyhedra rather than a single convex set. The hybrid does not reduce to the conventional interval at $\bar{M} = 0$, because it allocates part of the size budget to a first-stage test and conditions on the untreated-cohort moments. This accounts for the small discrepancy between the $\bar{M} = 0$ row of Table C.1 and the conventional interval reported in the memorandum line.

C.6 Results

Table C.1 reports the robust confidence sets and Figure C.1 plots them. Three results stand out.

First, the formality estimate is robust in the relevant range. The breakdown value is $\bar{M}^* = 1.15$, so the estimate remains significantly positive even if the counterfactual cohort profile departs from parallel trends after 1997 by somewhat more than the largest departure visible among untreated cohorts. At the natural focal value $\bar{M} = 1$, which permits a violation exactly as large as the worst observed, the interval is $[0.13, 2.38]$ and excludes zero.

Table C.1: Sensitivity to Violations of Parallel Cohort Profiles

	(1) Participation rate	(2) Formality rate
Panel A: relative magnitudes, $\Delta^{RM}(\bar{M})$		
$\bar{M} = 0.0$	[0.24, 0.51]	[1.15, 1.36]
$\bar{M} = 0.5$	[−0.03, 0.99]	[0.72, 1.79]
$\bar{M} = 1.0$	[−0.56, 1.47]	[0.13, 2.38]
$\bar{M} = 1.5$	[−1.04, 1.95]	[−0.40, 2.92]
$\bar{M} = 2.0$	[−1.53, 2.49]	[−0.94, 3.45]
Breakdown $\bar{M}^*$	0.43	1.15
Panel B: smoothness, $\Delta^{SD}(M)$		
$M = 0$ (linear extrapolation)	[0.34, 0.75]	[1.26, 1.62]
Breakdown M^*	0.010	0.073
Largest untreated second difference	0.166	0.112
Ratio	0.06	0.65
Panel C: alternative specifications, breakdown $\bar{M}^*$		
CRV1 covariance	0.51	1.38
Target: cohorts 2001–2004	0.41	1.02
Reference cohort 1998	0.48	0.78
Conventional 95% CI	[0.24, 0.55]	[1.12, 1.38]
Pre-reference cohorts / post-reference cohorts	6 / 7	6 / 7

Note: Robust 95% confidence sets for the exposure differential, constructed following Rambachan and Roth (2023) from the cohort-specific gradients of the fully interacted specification described in Section 3.3 and Appendix C. Gradients are normalized on the 1997 cohort. The target parameter is the average over the 2001 and 2002 cohorts except where noted. Panel A permits the counterfactual informality-specific cohort profile to move after 1997 by at most $\bar{M}$ times its largest movement among untreated cohorts; Panel B bounds its curvature by M in the units of the gradient. Breakdown values are the smallest $\bar{M}$ or M at which the robust set includes zero. Panel A uses the conditional least-favorable hybrid and Panel B the fixed-length construction. The covariance matrix is the cluster jackknife (CV3) over 32 states except in the first row of Panel C. The conventional interval is the state-clustered normal interval on the same target parameter.

Second, the participation estimate is considerably more fragile, with $\bar{M}^* = 0.43$. This is not a new finding but an independent confirmation: Table A.4 already shows the participation coefficient attenuating to 0.15–0.22 under flexible trends while formality holds between 0.78 and 1.21 and remains significant at the one percent level throughout. Two robustness strategies with different assumptions agree on which of the two margins carries the paper's conclusions.

Third, relaxing the assumption in the direction the data actually indicate strengthens rather than weakens the result. Under $\Delta^{SD}(0)$ the formality interval shifts from $[1.12, 1.38]$ to $[1.26, 1.62]$, because the drift across untreated cohorts is negative while the estimated effect is positive. Extrapolating that drift forward and removing it makes the effect larger. This is the formal version of the argument in Section 3.3 that smooth cohort trending would push exposed cohorts' gradients down rather than sharply up.

Figure C.1: Robust Confidence Sets under $\Delta^{RM}(\bar{M})$

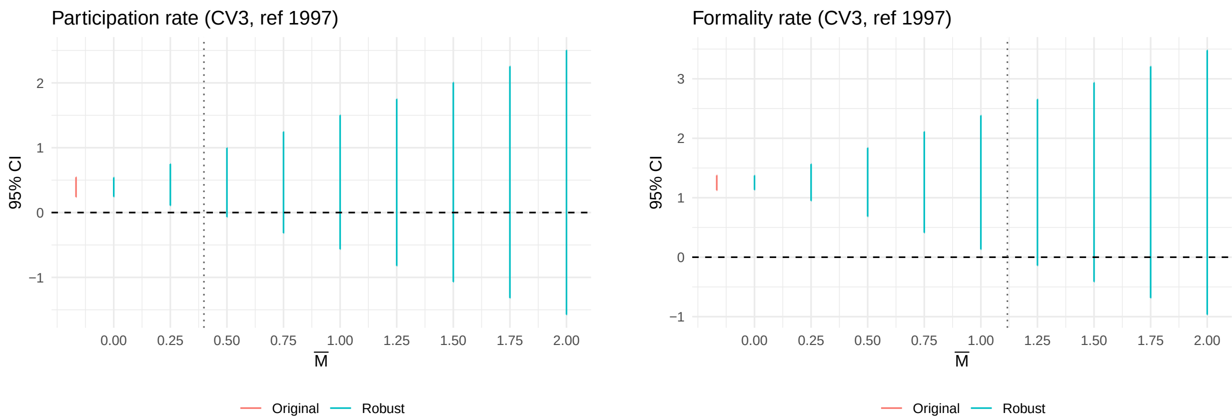

(a) Participation rate

(b) Formality rate

Note: Each panel plots 95% robust confidence sets for the exposure differential as a function of $\bar{M}$, the permitted ratio of the post-reference violation of parallel cohort profiles to the largest violation observed among untreated cohorts. The leftmost interval labelled "Original" is the conventional state-clustered interval. The dotted vertical line marks the breakdown value. Covariance is the cluster jackknife (CV3) over 32 states; gradients are normalized on the 1997 cohort; the target parameter is the average over the 2001 and 2002 cohorts. Source: Author's calculations using data from ENOE.

C.7 Limitations

Three caveats qualify the exercise. First, the smoothness class is weakly informative at this horizon. With seven post-reference cohorts, $\Delta^{SD}(M)$ permits the counterfactual profile to deviate from its linear extrapolation by an amount growing as $M \cdot k(k-1)/2$, so any curvature comparable to that observed among untreated cohorts overwhelms an effect of the estimated magnitude. The breakdown values in Panel B are therefore best read relative to the observed curvature: 0.07 against a largest untreated second difference of 0.11 for formality, or roughly two thirds, and 0.06 of the corresponding figure for participation. The $M = 0$ row remains informative for the reason given above, but I do not read the rest of Panel B as evidence either way.

Second, the reference cohort is improved rather than clean. As footnote 10 notes, 39 percent of the 1995–1998 cohorts are observed attending school during 2016–2019, so the 1997 cohort carries some college exposure of its own. Referencing on 1997 reduces contamination in the reference relative to 1998 but does not eliminate it. The direction is unchanged: any remaining contamination attenuates the estimate, so the reported intervals remain conservative.

Finally, the exercise disciplines the counterfactual cohort profile and nothing else. It has no bearing on the measurement of exposure, on the interpretation of the intention-to-treat differential, or on the demand-shock concerns addressed in Section 3.2.2. It answers one question narrowly and should be read as such.

D Contamination of the Comparison Group

Suppose college eligibility moves outcomes in the same direction as high-school eligibility and scales with informality in the same way, and let τ_s denote the schooling-linked effect on the exposed cohorts.[27] Net of the baseline age profile, the observed cohort gap is $\tau_s(1-c_s)$: the effect on the exposed, less the share of it that the partially treated comparison group also receives. Under the premise that motivates the design, $\tau_s = \beta I_{s,pre}$, so the gap is $\beta I_{s,pre}(1 - c_s)$ and the gradient the design estimates is

$$\frac{\partial}{\partial I_{s,pre}}\Big[\beta I_{s,pre}(1 - c_s)\Big] = \beta\Big[(1 - c_s) + 0.0021\, I_{s,pre}\Big] = \beta\Big[1 - \big(c_s - 0.0021\, I_{s,pre}\big)\Big],$$

where the fitted relation from Table A.6 is $c_s = a - 0.0021 I_{s,pre}$; the intercept drops out of the derivative. The two terms in the bracket are the two forces just described: the level of contamination enters as $-c_s$, its gradient as $+0.0021\, I_{s,pre}$. The estimate is attenuated, the bracket below one, whenever $c_s > 0.0021\, I_{s,pre}$.

That condition holds in every state by a wide margin. The least contaminated state has $c_s = 0.28$, while the threshold reaches its largest value, $0.0021 \times 80 = 0.17$, at the most informal state in the sample. Contamination therefore attenuates the estimated gradient throughout, and the estimates may be read as a lower bound of the combined high-school-plus-college schooling-linked impact. This lower-bound reading is conditional on

[27] Throughout this derivation the contamination share c_s is expressed as a proportion and the informality rate $I_{s,pre}$ in percentage points, so that β keeps its usual interpretation as the effect of a one-percentage-point increase in initial informality; on that scaling the Table A.6 slope of -0.21 percentage points of contamination per percentage point of informality becomes $\partial c_s/\partial I_{s,pre} = -0.0021$.

how the college effect varies across states. The derivation above assumes it scales with initial informality exactly as the high-school effect does, so that what the comparison group absorbs, $c_s\,\beta I_{s,pre}$, is itself larger where informality is higher. Suppose instead that college eligibility raises formality by a constant τ^C, and its gradient is $\beta + 0.0021\,\tau^C$. The attenuation term is gone: a contamination that does not scale with informality shifts each state's gap by a fixed amount rather than flattening its slope, so only the uneven distribution of contamination across states survives. That residual works upward, but it is small. Because contamination falls by about 8 percentage points across the full range of initial informality, a college effect of τ^C raises the amount differenced out by $0.08\,\tau^C$ more in the least informal state than in the most informal one; spread over that 40-point range, the estimated gradient is inflated by roughly $0.002\,\tau^C$. Even a college effect as large as 5 percentage points would inflate the gradient by about 0.01, roughly one percent of the long-difference formality coefficient of 0.87 in Table 2.

Finally, these ingredients translate the intention-to-treat estimates into magnitudes per enrolled-eligible student. The exposure differential between designated exposed and unexposed cohorts is at most 0.58 (the product of the enrollment rate and the public share) and smaller once the comparison group's partial treatment is counted. Dividing by this differential implies effects per enrolled-eligible student of at least 1.7 times the baseline estimates, and larger under fuller accounting of contamination.