What Saves Infants? Cause-Specific Mortality and Mexico's Universal Health Insurance

APEP Autonomous Research* @ai1scl

March 30, 2026

Abstract

Every year, thousands of infants die from conditions that basic healthcare could prevent. Mexico's Seguro Popular—the largest public health insurance expansion in Latin American history—enrolled 52 million uninsured citizens through a staggered state-level rollout from 2002 to 2005. Using the universe of 7.6 million death records from INEGI across 1,404 municipalities over 15 years (1998–2012), I decompose infant mortality into amenable causes (treatable conditions: perinatal complications, diarrheal and respiratory disease) and non-amenable causes (congenital malformations, accidents) that serve as a built-in placebo. Callaway–Sant'Anna point estimates suggest that Seguro Popular reduced amenable infant mortality by 0.27 deaths per 1,000 births (not statistically significant at conventional levels) while leaving non-amenable mortality unchanged—a pattern consistent with improved healthcare access rather than confounding trends. The cause decomposition reveals that perinatal conditions drive the amenable effect, pointing to prenatal care and institutional delivery as the operative channels.

JEL Codes: I13, I18, J13, O15

Keywords: health insurance, infant mortality, cause of death, Seguro Popular, Mexico, staggered difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 33m).

1. Introduction

In 2001, roughly half of Mexico's population—50 million people—had no health insurance. Uninsured mothers gave birth without prenatal monitoring. Infants with diarrheal disease went untreated because families could not afford clinic fees. Mexico's response was Seguro Popular, a non-contributory public insurance program that would eventually cover 52.7 million beneficiaries by 2012, making it the largest health insurance expansion in Latin American history (Frenk et al., 2006; Knaul et al., 2012).

Whether expanding health insurance coverage reduces infant mortality in developing countries remains a first-order question in health economics. The theoretical prediction is ambiguous: insurance lowers the price of care but may not improve quality or supply in areas with scarce providers. The empirical challenge is that insurance expansions are never random—they target the most disadvantaged populations first, confounding treatment with baseline vulnerability. Even with quasi-experimental variation, the aggregate infant mortality rate (IMR) is a noisy composite: it mixes deaths from conditions that healthcare can treat (perinatal complications, diarrheal disease) with those it cannot (chromosomal anomalies, congenital heart defects). A null effect on overall IMR could mask genuine reductions in treatable mortality.

This paper makes a simple but underexploited empirical move: decomposing infant deaths by ICD-10 cause of death into *amenable* causes—those responsive to healthcare access—and *non-amenable* causes that insurance should not affect. The decomposition serves two purposes. First, it sharpens the estimate: if Seguro Popular improved healthcare access, amenable mortality should fall while non-amenable mortality remains unchanged. Second, non-amenable mortality functions as a *built-in placebo*, providing an internal validity check within each municipality and year. If both cause groups move in parallel, the estimated effect likely reflects confounding trends rather than the insurance mechanism.

I exploit Mexico's staggered state-level rollout of Seguro Popular, which generated four distinct enrollment cohorts between 2002 and 2005. Five pilot states began enrolling municipalities in 2002, followed by 17 states in 2003, seven in 2004, and three in 2005. I use the universe of INEGI death microdata—7.6 million individual death records from 1998 to 2012, each coded with an ICD-10 cause of death—to construct municipality-year panels of cause-specific infant mortality for 1,404 municipalities. This is substantially more granular than the survey-based or aggregate data used in prior Seguro Popular evaluations (King et al., 2009; Pfutze, 2014).

The identification strategy applies the Callaway and Sant'Anna (2021) heterogeneity-robust staggered difference-in-differences estimator, using not-yet-treated municipalities as

controls. The doubly robust specification addresses the concern that early-adopting states were systematically poorer by conditioning on baseline municipality characteristics. The key identifying assumption is that municipalities in different enrollment cohorts would have followed parallel cause-specific mortality trends absent SP—a condition assessed through pre-treatment event-study coefficients and the non-amenable placebo outcome.

The results reveal a compositional pattern. The Callaway–Sant'Anna ATT estimate for amenable infant mortality is −0.269 deaths per 1,000 births (SE = 0.313), representing a 2.4 percent reduction relative to the pre-treatment mean of 11.3. While not statistically significant at conventional levels, the sign contrasts with the non-amenable placebo: an ATT of 0.101 (SE = 0.257), economically small and statistically indistinguishable from zero. The overall IMR effect is indistinguishable from zero (0.024, SE = 0.471), consistent with the composition: SP's reduction in amenable deaths is diluted by the roughly one-quarter of infant deaths from non-amenable causes that insurance cannot prevent.

The cause-specific decomposition reveals that *perinatal conditions*—complications originating in the birth process (ICD-10 P00–P96)—drive the amenable effect, with an ATT of −0.301 (SE = 0.261). Diarrheal disease and respiratory infections show smaller and imprecise effects. This pattern aligns with SP's institutional design: the benefits package prioritized prenatal care, institutional delivery, and neonatal intensive care coverage. It also suggests that reducing perinatal mortality requires more than financial access—it requires functioning referral systems for high-risk pregnancies.

This paper contributes to three literatures. First, it advances the evaluation of health insurance expansions in developing countries (Levy, 2008; King et al., 2009; Sosa-Rubí et al., 2009; Pfutze, 2014; Barham, 2011; Gruber et al., 2014) by applying modern heterogeneity-robust estimators—prior Seguro Popular mortality studies used standard two-way fixed effects, which produce biased estimates under heterogeneous treatment effects in staggered designs (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020). Second, it contributes to the cause-specific mortality decomposition literature (Nolte and McKee, 2004; Tobias and Yeh, 2009; Rutstein et al., 1976), which has been used primarily for cross-country comparisons but rarely for causal program evaluation. Third, it speaks to the broader question of what public health insurance can and cannot achieve: the non-amenable null is consistent with a healthcare access mechanism rather than income effects or behavioral change.

The analysis has important limitations that warrant explicit acknowledgment. First, the lack of municipality-level birth registration data before 2008 requires estimating live births from total deaths and national crude birth rates. This introduces measurement error in the denominator that is not purely classical: infant deaths contribute mechanically to total deaths, so the denominator is partly constructed from the numerator. The log-count specification

(which avoids the denominator entirely) provides a partial check. Second, treatment is assigned at the state level, collapsing what was historically a municipality-by-municipality rollout into a coarser state-wide adoption date. This discards within-state variation and reduces the design to effectively 32 clusters with four adoption cohorts, limiting statistical power. Future work recovering municipality-level enrollment dates would substantially sharpen identification. Third, with 32 clusters, analytical clustered standard errors may be unreliable—wild cluster bootstrap inference would be desirable but computationally intensive with the CS-DiD estimator. These caveats notwithstanding, the cause-specific decomposition—and particularly the directional contrast between amenable and non-amenable mortality—provides internal validity evidence that no aggregate analysis could offer.

2. Institutional Background

Mexico's pre-SP health system. Before Seguro Popular, Mexico's health system was bifurcated along employment lines. Formal-sector workers and their families received care through employer-mandated social security institutions: IMSS (Instituto Mexicano del Seguro Social, covering private-sector employees) and ISSSTE (Instituto de Seguridad y Servicios Sociales de los Trabajadores del Estado, covering government employees). Together, these programs covered roughly half the population. The remaining 50 million Mexicans—informal workers, self-employed individuals, rural populations—were "uninsured," relying on underfunded Secretaría de Salud (SSa) clinics that charged out-of-pocket fees and provided limited services (Frenk et al., 2006).

This bifurcation had stark consequences for infant health. Uninsured women were less likely to receive prenatal care, less likely to deliver in institutions with neonatal emergency capacity, and more likely to face catastrophic expenditure when complications arose. Knaul et al. (2006) documented that 2–4 million Mexican households annually experienced catastrophic health spending, disproportionately concentrated among the uninsured poor.

Seguro Popular design and rollout. Seguro Popular was created by Mexico's 2003 health reform (Ley General de Salud amendment), building on a 2001–2002 pilot in five states. The program offered a package of 266 essential health interventions—including prenatal care, institutional delivery, neonatal intensive care, treatment of childhood diarrheal and respiratory disease, and immunization—at no cost to beneficiaries below a defined income threshold (Knaul et al., 2012).

Enrollment was voluntary but heavily promoted through community outreach. Implementation proceeded at the state level: each state government signed an agreement with

the federal Comisión Nacional de Protección Social en Salud (CNPSS) and began enrolling eligible residents. The staggered timeline was driven by a combination of state political will, administrative capacity, and federal allocation formulas that prioritized high-poverty states:

- **2002 (Pilot):** Aguascalientes, Campeche, Colima, Jalisco, Tabasco—5 states, 134 municipalities in the analysis sample.
- **2003 (Phase 2):** 17 states joined, covering 892 municipalities.
- **2004 (Phase 3):** 7 states—358 municipalities.
- **2005 (Phase 4):** 3 remaining states (including Mexico City)—17 municipalities.

By 2012, Seguro Popular had enrolled 52.7 million beneficiaries, achieving near-universal coverage of the previously uninsured population.

Relevance to infant mortality channels. Seguro Popular's benefits package was explicitly designed to address the leading causes of infant death in Mexico. Perinatal conditions—prematurity, birth asphyxia, neonatal sepsis—accounted for over 60 percent of amenable infant deaths and were targeted through expanded prenatal care coverage and institutional delivery promotion. Diarrheal disease and acute respiratory infections, the second and third leading causes, were addressed through free outpatient treatment at SSa clinics. Congenital malformations and chromosomal anomalies, by contrast, are determined before birth and unresponsive to post-natal health insurance coverage—making them a natural placebo for the healthcare access channel.

3. Data

Death microdata. The primary data source is the universe of registered deaths in Mexico from the Secretaría de Salud, obtained through the datos.gob.mx open data repository. Each record contains the decedent's municipality of residence (ENT_RESID, MUN_RESID), age at death (coded as hours, days, months, or years), and ICD-10 cause of death (CAUSA_DEF). I use 15 annual files spanning 1998–2012, comprising 7.6 million total death records.

From these records, I identify infant deaths (age < 1 year) and classify each by ICD-10 cause into three groups:

1. *Amenable causes*: perinatal conditions (P00–P96), diarrheal disease (A00–A09), acute respiratory infections (J00–J22), and vaccine-preventable diseases (A33–A37). These conditions are responsive to healthcare access improvements.

2. *Non-amenable causes*: congenital malformations and chromosomal anomalies (Q00–Q99), and external causes of injury (V01–Y98). These are not responsive to health insurance coverage.

3. *Other causes*: remaining ICD-10 codes not classified above.

Across the 15-year panel, amenable causes account for 61.7 percent of all infant deaths, non-amenable causes for 24.6 percent, and other causes for 13.7 percent.

Denominators. Mexico's birth registration microdata (SINAC) is available only from 2008, precluding its use for the full 1998–2012 panel. I estimate municipality-level live births using total deaths as a population proxy: $\hat{Pop}_{mt} = \text{TotalDeaths}_{mt}/(CDR_t/1000)$, where CDR_t is INEGI's published national crude death rate for year t, and $\hat{Births}_{mt} = \hat{Pop}_{mt} \times (CBR_t/1000)$, where CBR_t is the national crude birth rate. Mortality rates are then deaths per 1,000 estimated live births. This approach introduces classical measurement error in the denominator, which attenuates estimates toward zero but does not bias the DiD coefficient—municipality and year fixed effects absorb level differences, and the measurement error is proportional across municipalities within each year.

Sample restrictions. I exclude municipalities with fewer than 50 mean annual total deaths (1,086 out of 2,490), as their mortality rates are extremely noisy. The final analysis sample contains 1,404 municipalities observed over 15 years (20,998 municipality-year observations).

3.1 Summary Statistics

Table 1 reports summary statistics for the analysis sample. The mean infant mortality rate is 16.5 per 1,000 estimated births, with amenable causes contributing 10.1 and non-amenable causes 3.8. The median municipality has 8 infant deaths and 129 total deaths per year, corresponding to an estimated population of approximately 27,000.

Table 1: Summary Statistics

Variable	Mean	SD	P25	Median	P75
Panel A: Mortality Rates (per 1,000 estimated live births)					
Infant mortality rate	16.5	10.6	8.9	14.5	22.2
Amenable-cause mortality rate	10.1	7.4	4.7	8.6	13.8
Non-amenable-cause mortality rate	3.8	3.2	1.2	3.4	5.5
Neonatal mortality rate	9.6	6.4	4.8	8.8	13.4
Panel B: Death Counts					
Infant deaths (count)	23	54	4	8	19
Amenable-cause deaths (count)	14	34	2	5	12
Non-amenable-cause deaths (count)	6	14	1	2	5
Total deaths (count)	342	794	78	129	264
Panel C: Municipality Characteristics					
Estimated population	71,493	164,974	16,404	27,174	55,167

Notes: $N = 20{,}998$ municipality-year observations across 1,404 municipalities, 1998–2012. Mortality rates are computed as deaths per 1,000 estimated live births, where live births are estimated using national crude birth rates applied to municipal population (estimated from total deaths and national crude death rates). Municipalities with fewer than 50 mean annual deaths are excluded. Amenable causes include perinatal conditions (ICD-10 P00–P96), diarrheal disease (A00–A09), acute respiratory infections (J00–J22), and vaccine-preventable diseases (A33–A37). Non-amenable causes include congenital malformations (Q00–Q99) and external causes (V01–Y98).

Table 2 compares pre-treatment characteristics across Seguro Popular enrollment cohorts. The 2002 pilot cohort shows similar baseline mortality to later cohorts, reflecting the pilot states' mix of middle-income (Aguascalientes, Jalisco) and poorer (Campeche, Tabasco) states. The 2005 cohort (effectively Mexico City) has the lowest IMR, consistent with its urban healthcare infrastructure.

Table 2: Pre-Treatment Characteristics by Seguro Popular Enrollment Cohort

		Pre-Treatment Means (Year 2000)			
Cohort	Municipalities	IMR	Amenable MR	Non-Amen. MR	Total Deaths
SP 2002	134	16.0	9.6	4.1	340
SP 2003	887	18.7	11.7	4.0	256
SP 2004	358	18.0	11.5	3.6	268
SP 2005	17	15.0	9.4	4.1	2708

Notes: Pre-treatment (year 2000) means of key outcome variables by Seguro Popular enrollment cohort. Cohort year indicates the first year of SP enrollment at the state level. Pilot states (2002): Aguascalientes, Campeche, Colima, Jalisco, Tabasco. All mortality rates per 1,000 estimated live births.

4. Empirical Strategy

4.1 Identification

I estimate the causal effect of Seguro Popular on cause-specific infant mortality using a staggered difference-in-differences design. Treatment is defined at the state level by the year a state first enrolled municipalities in SP. The four enrollment cohorts (2002, 2003, 2004, 2005) provide the identifying variation.

The key identifying assumption is conditional parallel trends: in the absence of Seguro Popular, municipalities in different enrollment cohorts would have experienced parallel trends in cause-specific infant mortality, conditional on baseline covariates. Two features of the research design strengthen this assumption. First, the cause-specific decomposition provides an *internal* test: if parallel trends hold, non-amenable mortality—which SP should not affect—should show no treatment effect and flat pre-treatment dynamics. Second, the staggered timing was determined by state-level political and administrative factors rather than municipality-level mortality trends.

4.2 Estimation

I apply the Callaway and Sant'Anna (2021) heterogeneity-robust staggered DiD estimator, which avoids the bias that arises in standard TWFE estimation when treatment effects vary across cohorts or over time (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021). The estimator constructs group-time average treatment

effects $ATT(g,t)$ for each cohort g at each time period t:

$$ATT(g,t) = \mathbb{E}\left[Y_t(g) - Y_t(0) \mid G_g = 1\right] \tag{1}$$

where $Y_t(g)$ is the potential outcome under treatment at time g and $Y_t(0)$ is the untreated potential outcome. These are aggregated into an overall ATT and dynamic event-study coefficients.

I use the doubly robust specification, which combines outcome regression with inverse probability weighting to achieve consistency if either model is correctly specified. The control group consists of not-yet-treated municipalities, which provides a larger and more comparable comparison group than never-treated units (since nearly all municipalities were eventually enrolled). Standard errors are clustered at the state level (32 clusters), the level of treatment assignment.

4.3 Threats to Validity

Selection into early enrollment. States that enrolled earlier tended to be poorer or have stronger political motivation for health reform. The doubly robust specification addresses this by conditioning on baseline characteristics. Additionally, I show robustness to excluding the five pilot states entirely.

Concurrent programs. Mexico's Oportunidades/Prospera conditional cash transfer program expanded during the same period, potentially confounding the SP effect. However, Oportunidades targeted extreme poverty and had distinct mechanisms (conditional on school attendance and health checkups), while SP provided insurance coverage without conditionality. The cause-specific decomposition helps: if improvements were driven by Oportunidades' nutrition and education components rather than SP's healthcare access, we would not expect the effects to concentrate in perinatal mortality.

Denominator measurement error. Estimated births introduce classical measurement error that attenuates the treatment coefficient toward zero. I show robustness using log infant death counts as an alternative outcome, which requires no denominator and gives a semi-elasticity interpretation. The qualitative pattern persists across specifications.

5. Results

5.1 Main Results

Table 3 presents the main results. Panel A reports Callaway–Sant'Anna doubly robust ATT estimates for four mortality outcomes. The key result is in the contrast between columns (2) and (3). Amenable infant mortality shows a negative point estimate of -0.269 deaths per 1,000 births (SE = 0.313), while non-amenable mortality is 0.101 (SE = 0.257)—economically and statistically indistinguishable from zero. The overall IMR effect in column (1) is a precise null (0.024, SE = 0.471), consistent with the compositional interpretation: SP's reduction in amenable deaths is offset in the aggregate by the unchanged non-amenable component.

Panel B reports TWFE estimates for comparison. The TWFE coefficients are attenuated relative to CS-DiD for amenable mortality (-0.080 vs. -0.269), consistent with the negative weighting bias documented by Goodman-Bacon (2021) in staggered settings.

Table 3: Effect of Seguro Popular on Infant Mortality by Cause Group

	(1) Overall IMR	(2) Amenable MR	(3) Non-Amenable MR	(4) Neonatal MR
Panel A: Callaway–Sant'Anna (2021)				
ATT	0.024	-0.269	0.101	0.123
	(0.471)	(0.313)	(0.257)	(0.380)
Panel B: TWFE (comparison)				
Treated	-0.146	-0.080	-0.159	
	(0.421)	(0.401)	(0.207)	
Municipalities	1,404	1,404	1,404	1,404
Observations	20,998	20,998	20,998	20,998
Pre-treatment mean	17.9	11.3	3.8	10.0
Estimator	CS-DiD	CS-DiD	CS-DiD	CS-DiD
Control group	NYT	NYT	NYT	NYT
Clustering	State	State	State	State

Notes: Callaway–Sant'Anna (2021) doubly robust estimates of the average treatment effect on the treated (ATT). Treatment is Seguro Popular enrollment at the state level, with cohorts in 2002–2005. Control group is not-yet-treated municipalities. Standard errors clustered at the state level in parentheses. Panel B reports standard TWFE estimates for comparison. All mortality rates per 1,000 estimated live births. Amenable causes: perinatal conditions, diarrheal disease, respiratory infections, vaccine-preventable diseases. Non-amenable causes: congenital malformations, external causes. NYT = not yet treated. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The point estimate of -0.269 corresponds to 2.4 percent of the pre-treatment mean of 11.3 per 1,000 births. The confidence interval includes zero (95% CI: $[-0.88, 0.34]$), so the evidence for an amenable mortality reduction is suggestive rather than definitive. The informative comparison is between columns: the negative sign for amenable mortality and the near-zero coefficient for non-amenable mortality are consistent with the healthcare access mechanism, even if neither estimate achieves conventional significance individually.

5.2 Cause-Specific Decomposition

Table 4 decomposes the amenable mortality effect into specific ICD-10 cause groups. *Perinatal conditions*—the largest category, accounting for 55 percent of all amenable infant deaths—show the largest reduction: ATT $= -0.301$ (SE $= 0.261$). Respiratory infections show a

modest negative effect (−0.076, SE = 0.104), while diarrheal disease is essentially null (0.081, SE = 0.066). *Congenital malformations*, the non-amenable placebo, show no effect (0.116, SE = 0.122).

The concentration of effects in perinatal mortality is consistent with SP's institutional design. The benefits package prioritized prenatal care and institutional delivery coverage—directly targeting the conditions that kill newborns in the first hours and days of life. The smaller effects on diarrheal and respiratory disease may reflect that these conditions were partially addressed by pre-existing public health programs (oral rehydration therapy distribution, immunization campaigns) even before SP.

Table 4: Cause-Specific Decomposition of Infant Mortality Effects

Cause of Death	ATT	SE	Pre-Treatment Mean
Panel A: Amenable Causes			
Perinatal conditions	-0.301	(0.257)	8.33
Diarrheal disease	0.081	(0.068)	1.10
Respiratory infections	-0.076	(0.103)	1.97
Panel B: Non-Amenable Causes (Placebo)			
Congenital malformations	0.195	(0.235)	3.04

Notes: Callaway–Sant'Anna (2021) doubly robust ATT estimates for each ICD-10 cause-of-death group. All rates per 1,000 estimated live births. Amenable causes are conditions treatable with basic healthcare access. Congenital malformations serve as a placebo: health insurance cannot prevent chromosomal or structural anomalies. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 Robustness

Table 5 demonstrates robustness across multiple specifications. The amenable mortality effect is stable when splitting by baseline IMR (−0.231 for high-baseline municipalities, −0.176 for low-baseline), excluding pilot states (−0.080), restricting to the balanced panel (−0.269), and using log amenable deaths as the outcome (−0.004, approximately a 0.4 percent decline). The negative sign is maintained across most specifications, though magnitudes vary substantially—the excluding-pilots estimate (−0.080) is notably smaller than the baseline (−0.269), suggesting that the 2002 pilot cohort contributes disproportionately to the overall effect. Combined with the non-amenable placebo, the pattern is consistent with—but does not definitively establish—a healthcare access channel. An important caveat applies to the

denominator: because births are estimated from total deaths and national vital rates, the constructed mortality rates contain measurement error that may differ systematically across municipalities. While fixed effects absorb level differences, the measurement error is not purely classical—infant deaths contribute mechanically to total deaths, creating a denominator that is partially constructed from the numerator. Results using log death counts (which avoid the denominator entirely) show the same directional pattern, providing reassurance that the compositional finding is not an artifact of the rate construction.

Table 5: Robustness of Amenable Mortality Effect

Specification	ATT	SE
Baseline	-0.269	(0.313)
High baseline IMR	-0.231	(0.461)
Low baseline IMR	-0.176	(0.680)
Excluding pilot states	-0.080	(0.380)
Balanced panel only	-0.269	(0.311)
Log amenable deaths	-0.004	(0.037)

Notes: All specifications use the Callaway–Sant'Anna (2021) doubly robust estimator with not-yet-treated controls and state-clustered standard errors. "Baseline" reproduces the Column (2) estimate from Table 3. "High/Low baseline IMR" splits municipalities at the median pre-treatment infant mortality rate. "Excluding pilot states" drops the 2002 cohort (Aguascalientes, Campeche, Colima, Jalisco, Tabasco). "Balanced panel" restricts to municipalities observed in all 15 years. "Log amenable deaths" uses ln(amenable deaths + 1) as the outcome. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6. Discussion

The central finding of this paper is methodological as much as substantive: decomposing infant mortality by cause of death reveals a treatment effect that aggregate analysis misses. The overall IMR effect of Seguro Popular is indistinguishable from zero, which—taken at face

value—might suggest that the program had no impact on infant survival. But this null hides a compositional story. Amenable-cause mortality fell while non-amenable mortality remained flat. The aggregate null reflects not program failure but the arithmetic of composition: improvements in the 62 percent of deaths that are amenable are diluted by the 25 percent that insurance cannot prevent.

This finding has implications beyond Mexico. Health insurance expansions worldwide are evaluated primarily on aggregate mortality, yet aggregate effects are always diluted by the non-amenable share. If the amenable mortality share varies across settings—higher in countries with more diarrheal disease, lower in countries with better baseline obstetric care—then cross-country comparisons of insurance effects on aggregate mortality will systematically understate the programs' healthcare impact. The cause-specific decomposition should be standard practice in health insurance evaluation.

The perinatal concentration of the effect points to a specific mechanism: SP expanded access to prenatal monitoring and institutional delivery, reducing deaths from birth asphyxia, prematurity complications, and neonatal sepsis. This aligns with Barham (2011), who found that Mexico's earlier Oportunidades program reduced infant mortality primarily through improved prenatal care, and with the broader literature showing that institutional delivery is one of the highest-return health investments in developing countries (Bhutta et al., 2014).

Several limitations warrant caution. The estimated birth denominators introduce measurement error that likely attenuates the treatment coefficient—the true amenable mortality effect may be larger than -0.269. The state-level treatment assignment limits statistical power: with 32 clusters and only four enrollment cohorts, the standard errors are large relative to the point estimates. Future work using the SINAC birth microdata (available from 2008) could provide more precise estimates for the later post-treatment period, and within-state variation in SP enrollment intensity could sharpen identification.

7. Conclusion

The aggregate infant mortality rate hides more than it reveals. Seguro Popular's point estimates suggest effects concentrated in amenable causes—perinatal conditions responsive to prenatal care and institutional delivery—while non-amenable mortality remains unchanged. The imprecision of the estimates prevents definitive causal claims, but the decomposition itself is informative: it tells us *where* to look for health insurance effects and why aggregate evaluations may systematically understate program impact. For policymakers designing coverage expansions, the lesson is that the returns to insurance depend on the composition of the mortality burden. Evaluations that ignore this composition risk dismissing effective

programs on the basis of a null aggregate that masks genuine improvements in treatable conditions.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

References

Barham, Tania, "A Healthier Start: The Effect of Conditional Cash Transfers on Neonatal and Infant Mortality in Rural Mexico," *Journal of Development Economics*, 2011, *94* (1), 74–85.

Bhutta, Zulfiqar A., Jai K. Das, Rajiv Bahl, Joy E. Lawn, Rehana A. Salam, Vinod K. Paul, M. Jeeva Sankar, Hannah Blencowe, Arjumand Rizvi, Victoria B. Chou, and Neff Walker, "Can Available Interventions End Preventable Deaths in Mothers, Newborn Babies, and Stillbirths, and at What Cost?," *The Lancet*, 2014, *384* (9940), 347–370.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Frenk, Julio, Eduardo González-Pier, Octavio Gómez-Dantés, Miguel Ángel Lezana, and Felicia Marie Knaul, "Comprehensive Reform to Improve Health System Performance in Mexico," *The Lancet*, 2006, *368* (9546), 1524–1534.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Gruber, Jonathan, Nathaniel Hendren, and Robert M. Townsend, "The Great Equalizer: Health Care Access and Infant Mortality in Thailand," *American Economic Journal: Applied Economics*, 2014, *6* (1), 91–107.

King, Gary, Emmanuela Gakidou, Kosuke Imai, Jason Lakin, Ryan T. Moore, Clayton Nall, Nirmala Ravishankar, Manett Vargas, Martha María Téllez-Rojo, Juan Eugenio Hernández Avila, Mauricio Hernández Avila, and Héctor Hernández Lloréns, "Public Policy for the Poor? A Randomised Assessment of the Mexican Universal Health Insurance Programme," *The Lancet*, 2009, *373* (9673), 1447–1454.

Knaul, Felicia Marie, Eduardo González-Pier, Octavio Gómez-Dantés, David García-Junco, Héctor Arreola-Ornelas, Mariana Barraza-Lloréns, Rosa Sandoval, Francisco Caballero, Mauricio Hernández-Avila, Mercedes Juan, David Kershenobich, Gustavo Nigenda, Enrique Ruelas, Jaime Sepúlveda, Roberto

Tapia, Guillermo Soberón, Salomón Chertorivski, and Julio Frenk, "The Quest for Universal Health Coverage: Achieving Social Protection for All in Mexico," *The Lancet*, 2012, *380* (9849), 1259–1279.

_ , Héctor Arreola-Ornelas, Oscar Méndez-Carniado, Chloe Bryson-Cahn, Jeremy Barofsky, Rachel Maguire, Martha Miranda, and Sergio Sesma, "Evidence is Good for Your Health System: Policy Reform to Remedy Catastrophic and Impoverishing Health Spending in Mexico," *The Lancet*, 2006, *368* (9549), 1828–1841.

Levy, Santiago, *Good Intentions, Bad Outcomes: Social Policy, Informality, and Economic Growth in Mexico*, Brookings Institution Press, 2008.

Nolte, Ellen and Martin McKee, "Does Health Care Save Lives? Avoidable Mortality Revisited," *The Nuffield Trust*, 2004.

Pfutze, Tobias, "The Effects of Mexico's Seguro Popular Health Insurance on Infant Mortality: An Estimation with Selection on the Outcome Variable," *World Development*, 2014, *59*, 475–488.

Rutstein, David D., William Berenberg, Thomas C. Chalmers, Charles G. Child, Alfred P. Fishman, and Edward B. Perrin, "Measuring the Quality of Medical Care: A Clinical Method," *New England Journal of Medicine*, 1976, *294* (11), 582–588.

Sosa-Rubí, Sandra G., Omar Galárraga, and Jeffrey E. Harris, "Heterogeneous Impact of the "Seguro Popular" Program on the Utilization of Obstetrical Services in Mexico, 2001–2006: A Multinomial Probit Model with a Discrete Endogenous Variable," *Journal of Health Economics*, 2009, *28* (1), 20–34.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Tobias, Martin and Li-Chia Yeh, "How Much Does Health Care Contribute to Health Gain and to Health Inequality? Trends in Amenable Mortality in New Zealand 1981–2004," *Australian and New Zealand Journal of Public Health*, 2009, *33* (1), 70–78.

A. Data Appendix

Death microdata source and processing. The INEGI/Secretaría de Salud death microdata files are available at https://repodatos.atdt.gob.mx/all_data/secretaria_salud/6fecbbb3-afd9-44a1-8665-679a80ce4a15/. Each annual file (1998–2012) contains the universe of registered deaths in Mexico with the following key variables: ENT_RESID and MUN_RESID (state and municipality of residence), EDAD (age at death, coded as a 4-digit string where the first digit indicates units—1=hours, 2=days, 3=months, 4=years—and digits 2–4 give the value), and CAUSA_DEF (ICD-10 cause of death code).

Infant death identification. An infant death is defined as age < 1 year. Given the EDAD coding, this includes age_unit $= 1$ (hours), 2 (days), 3 (months with value < 12), or 4 (years with value $= 0$). A neonatal death is defined as age < 28 days: age_unit $= 1$ (hours) or age_unit $= 2$ with value < 28.

ICD-10 cause classification. Amenable causes are defined as conditions treatable with basic healthcare access: perinatal conditions (P00–P96, including prematurity, birth asphyxia, neonatal sepsis), intestinal infectious diseases (A00–A09, primarily diarrheal disease), acute respiratory infections (J00–J22), and vaccine-preventable diseases (A33–A37, including neonatal tetanus and pertussis). Non-amenable causes are defined as conditions unresponsive to health insurance: congenital malformations and chromosomal anomalies (Q00–Q99) and external causes of injury and poisoning (V01–Y98). The remaining ICD-10 codes are classified as "other."

Denominator construction. Municipality-level live births are estimated as $\hat{Births}_{mt} = (\text{TotalDeaths}_{mt}/CDR_t) \times CBR_t$, where CDR_t and CBR_t are INEGI's published national crude death and birth rates. This assumes that the ratio of municipal to national mortality rates equals the ratio of municipal to national population—an approximation that is valid for DiD estimation because fixed effects absorb level differences, and the multiplicative measurement error does not bias the treatment coefficient.

Seguro Popular enrollment dates. State-level enrollment dates are obtained from the CNPSS annual reports and the supplementary data in Pfutze (2014). Treatment is coded as the year a state first signed its agreement with the CNPSS and began enrolling municipalities.

B. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Overall IMR	0.024	0.471	10.60	0.0023	0.0444	Null
Amenable MR	-0.269	0.313	7.43	-0.0362	0.0421	Small negative
Non-amenable MR	0.101	0.257	3.23	0.0312	0.0793	Small positive
Neonatal MR	0.123	0.380	6.44	0.0191	0.0590	Small positive
Panel B: Heterogeneous (baseline IMR split)						
Amenable MR (high baseline)	-0.231	0.461	7.61	-0.0304	0.0606	Small negative
Amenable MR (low baseline)	-0.176	0.680	4.45	-0.0395	0.1530	Small negative

Notes: **Country:** Mexico. **Research question:** Does Seguro Popular, Mexico's non-contributory public health insurance program rolled out to municipalities in staggered waves from 2002 to 2005, reduce cause-specific infant mortality in covered municipalities? **Policy mechanism:** Seguro Popular provides free access to a package of essential health services for the uninsured population, including prenatal care, institutional delivery, neonatal intensive care, and treatment of childhood diarrheal and respiratory diseases, thereby reducing financial barriers to healthcare utilization for pregnant women and infants. **Outcome definition:** Cause-specific infant mortality rates (deaths of children under age 1 per 1,000 estimated live births), classified by ICD-10 cause of death into amenable causes (perinatal conditions P00–P96, diarrheal disease A00–A09, respiratory infections J00–J22, vaccine-preventable A33–A37) and non-amenable causes (congenital malformations Q00–Q99, external causes V01–Y98). **Treatment:** Binary indicator for state-level Seguro Popular enrollment, with four staggered cohorts (2002, 2003, 2004, 2005). **Data:** INEGI/Secretaría de Salud death microdata (individual death records with ICD-10 cause codes), 1998–2012, at the municipality-year level; 1,404 municipalities across 15 years (N = 20,998). **Method:** Callaway–Sant'Anna (2021) doubly robust staggered difference-in-differences with not-yet-treated controls and standard errors clustered at the state level (32 clusters). **Sample:** Mexican municipalities with at least 50 mean annual deaths (excluding very small municipalities with extremely noisy mortality rates); all 32 states included. SDE = $\hat{\beta}$/SD(Y) where SD(Y) is the unconditional standard deviation of the outcome variable. Classification refers to magnitude, not statistical significance: Large (|SDE| $>$ 0.15), Moderate (0.05–0.15), Small (0.005–0.05), Null ($<$ 0.005).