

# Can we detect the long-run effects of the measles vaccine? Replicating Atwood and Pearlman (2023)

February 6, 2024

## Summary

Atwood and Pearlman (2023) studies the effect of childhood access to the 1973 measles vaccine on adult economic outcomes in Mexico, closely following the approach in Atwood (2022). This method uses geographic variation in measles incidence, but measles is so contagious that we'd expect all states to have the same long-run average infection rate ( $\sim 95\%$ ). So it is not clear how to interpret the treatment variation, or why the vaccine would have larger benefits in states with higher reported incidence.

Section 3 repeats the main regressions, and finds positive effects for income, employment, and education. In Section 4, I show that the main results are not statistically significant when clustering standard errors by state-of-birth (similar to Atwood (2022)). The authors cluster by state-of-birth interacted with year-of-birth, which is a restrictive assumption. I use their clustering for the remaining results.

I use an event study in Section 5 to test whether the findings represent pre-existing trends. I also follow Roodman (2018) and fit a piecewise regression in Section 6, to test for trend breaks when cohorts are fully unexposed and fully exposed to the vaccine. I find trends that are inconsistent with a treatment effect of the vaccine.

I test for mean reversion in Section 7 using pre-vaccine state characteristics, and find that the results are not very robust. The effect on income disappears, and the other effects are much smaller.

As with my replication of Atwood's paper on the US vaccine (Wiebe, 2024), I conclude that using treatment variation from average reported measles incidence does not allow us to estimate the long-term effects of the measles vaccine.

# 1 Identifying variation: reported measles incidence

Since the vaccine was introduced nationally in 1973, there is no cross-sectional variation in treatment. To perform a difference-in-differences analysis, Atwood and Pearlman (2023) (henceforth AP) interact a time series variable (number of years exposed to the vaccine) with a cross-sectional variable (pre-vaccine average reported measles rate):  $Exp_t \times M_s^{pre}$ . This creates a continuous treatment variable with state-year variation, to estimate whether the vaccine had differential effects by pre-vaccine average reported measles incidence.

The main problem with applying the disease burden method to measles is that there is no clear source of geographic variation in actual measles infections. Measles is so contagious ( $R_0 = 12-18$ ) that virtually everyone becomes infected by adulthood.<sup>1</sup> So any measured geographic variation in long-run average reported measles cases likely reflects differences in reporting capacity or other factors<sup>2</sup>, and not differences in the actual incidence of the disease. There can be *short-term* variation in measles infections, based on differential timing of epidemic cycles. But since the disease burden method is based on a long-run average, this short-term variation is averaged out. If we controlled for these factors (by interacting with the time series variable  $Exp_t$ ), there would be no identifying variation left over. In contrast, hookworm and malaria are determined by geographic factors, leading to geographic variation in disease incidence (Bleakley, 2007, 2010). Because measles is different from other diseases, we should be skeptical that papers studying the effect of the measles vaccine using the disease burden method are estimating a meaningful quantity.

---

<sup>1</sup>Chuard et al. (2022) uses an epidemiological SIR model based on the assumption that the percentage ever-infected was  $\sim 100\%$  in the pre-vaccine era (Figure A3).

<sup>2</sup>States with worse initial health levels might have more severe outbreaks, and severe measles cases are more likely to be reported.

## 2 Data

Since AP is not yet published, I collected data to replicate parts of their analysis. I use data on economic outcomes (income, employment, and education) from the Mexican census via IPUMS, for years 1990, 1995, 2000, 2010, and 2015. AP use census data only from 1995 and 2000. While AP run separate regressions for each census year, I pool the census years. I also use data from the 1970 census for control variables. Note that AP also report results using the ENE, ENEU, and ENOE labor market surveys.

Following AP, I restrict the sample to men aged 18-65 who are born in Mexico. I also adjust the income variable for inflation using CPI from December, 2004.

I hand-collected disease data from 1965-1978 from PDFs of the annual epidemiology bulletins. Following AP, I collected data on measles and syphilis. AP state in footnote 32 that pertussis and tetanus are not included in the bulletins, but they are, so I collected those variables as well.

### 3 Main results

I replicate the main results on income and education from Table 4 and 5 in AP. Following Atwood (2022), I use the per capita measles rate (instead of per 100,000 population) to make the coefficients more readable. I find results with the same sign that are similar in magnitude, but off by a factor of 1000.<sup>3</sup>

Table 1: Replicating main regressions

|                                  | Log income<br>(1)    | Employed<br>(2)       | Education<br>(3)    |
|----------------------------------|----------------------|-----------------------|---------------------|
| $M_{pre} \times \text{Exposure}$ | 4.260***<br>(0.5160) | 0.5910***<br>(0.1481) | 30.19***<br>(3.350) |
| Observations                     | 6,936,812            | 10,633,483            | 10,529,025          |
| R <sup>2</sup>                   | 0.16                 | 0.10                  | 0.26                |
| Married fixed effects            | ✓                    | ✓                     | ✓                   |
| Urban fixed effects              | ✓                    | ✓                     | ✓                   |
| Birthyear fixed effects          | ✓                    | ✓                     | ✓                   |
| Birthplace fixed effects         | ✓                    | ✓                     | ✓                   |
| Survey year fixed effects        | ✓                    | ✓                     | ✓                   |

*Note:*  $M_{pre} \times \text{Exposure}$  is the state-level pre-vaccine reported measles rate interacted with exposure to the vaccine. Standard errors (in parentheses) are clustered by state-of-birth interacted with year-of-birth. Survey weights are used. Significance levels: \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

<sup>3</sup>One possible explanation is that AP used the measles rate per 1000 population. For example, I get an estimate of 0.59 for the employment rate, while AP report 0.00068 in Table 5, Column 2. Similarly, my coefficient for log income is 4.26, while AP report 0.0039.

## 4 Clustering standard errors

Following Atwood (2022), AP cluster standard errors at the state-of-birth by year-of-birth level, which is a restrictive assumption, since it does not allow the errors to be correlated across cohorts within the same state. They write: “Standard errors are clustered at the level of exposure to the measles vaccine, which is state and year of birth.” (p.14) But since the vaccine was rolled out nationally in 1973 and targeted children up to age five, it is not true that the level of exposure was state-of-birth by year-of-birth. Moreover, the original disease-burden studies (Bleakley, 2007, 2010) clustered by state-of-birth.

Below I rerun the main results clustering by state-of-birth. As with Atwood (2022), none of the results are statistically significant.

Table 2: Main results: clustering by state-of-birth

|                                  | Log income<br>(1) | Employed<br>(2)    | Education<br>(3) |
|----------------------------------|-------------------|--------------------|------------------|
| $M_{pre} \times \text{Exposure}$ | 4.260<br>(2.978)  | 0.5910<br>(0.6405) | 30.19<br>(21.65) |
| Observations                     | 6,936,812         | 10,633,483         | 10,529,025       |
| R <sup>2</sup>                   | 0.16              | 0.10               | 0.26             |
| Married fixed effects            | ✓                 | ✓                  | ✓                |
| Urban fixed effects              | ✓                 | ✓                  | ✓                |
| Birthyear fixed effects          | ✓                 | ✓                  | ✓                |
| Birthplace fixed effects         | ✓                 | ✓                  | ✓                |
| Survey year fixed effects        | ✓                 | ✓                  | ✓                |

*Note:*  $M_{pre} \times \text{Exposure}$  is the state-level pre-vaccine reported measles rate interacted with exposure to the vaccine. Standard errors (in parentheses) are clustered by state-of-birth. Survey weights are used. Significance levels: \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

## 5 Testing for convergence with an event study

Since it is questionable that average reported measles rates measure actual disease incidence, it is possible that any treatment effects found by this method simply capture convergence or divergence between High and Low measles states.

I run an event study by interacting the pre-vaccine average reported measles rate  $M_s^{pre}$  with year-of-birth indicators. AP define the vaccine exposure variable to be 0 in birth years 1957 and below, rising by 1 per year from 1958 until 1973, when it takes a value of 16 for that and all following cohorts. Hence, I omit the 1957 cohort as the reference year. Following Bleakley (2007), if the vaccine has an effect, we should observe the following pattern in the point estimates (for positive effects): flat until 1957 (as no one benefits from the vaccine), increasing steadily from 1957 to 1973 (as cohorts become more exposed to the vaccine), and flat from 1973 onwards (as everyone benefits from the vaccine). In the graphs below, I indicate the 1957-1973 period with red lines.

Following AP, I include fixed effects for survey year, state-of-birth, year-of-birth, marital status, and urban status. Standard errors are clustered by state-of-birth interacted with year-of-birth.

The results are plotted in Figures 1 - 3. The income results look like a rising trend over the full sample, both before and after the vaccine. For employment, we see a slight upward trend, but none of the coefficients are statistically different from zero. Education appears to have a trend break around 1957, but the coefficients continue to increase after 1973 when all cohorts have equal access to the vaccine (and there is no change in treatment dose).

Figure 1: Event study: log income




Figure 2: Event study: employed







Figure 3: Event study: education



## 6 Testing for a trend break

Following Roodman (2018), Fig. 6, I formally test for a trend break using a piecewise regression with kink points at 1958 and 1973, corresponding to the 1958-1973 cohorts with increasing exposure to the vaccine (the last year  $Exp_t = 0$  is 1957, and the first year  $Exp_t = 16$  is 1973). This creates three line segments corresponding to three time windows, representing cohorts with no, partial, and full access to the vaccine.

Specifically, this piecewise regression replaces  $Exp_t \times M_s^{pre}$  with three terms:

$$M_s^{pre} \times t, M_s^{pre} \times \max(0, t - 1957), \text{ and } M_s^{pre} \times \max(0, t - 1973) \quad (1)$$

where  $t$  is birth year. The coefficients on the second and third terms represent the change in slope at the kink points, so we can interpret the p-values directly as tests of a trend break. I restrict the sample to 1941-1989. As before, standard errors are clustered by state-of-birth interacted with year-of-birth.

I overlay the piecewise regression on top of the event study point estimates, scaling the piecewise fit to be centered vertically at the average of the point estimates (by adding the average difference between the piecewise fit and the point estimates). I also present the p-values for the trend break tests. The results are plotted in Figures 4-6.

The results do not support a treatment effect of the vaccine. All three variables are increasing throughout the sample, even after the vaccine is implemented and the coefficients should level off (because the treatment dose stops increasing). The 1973-1989 spline for employment is flat, as predicted by an effect of the vaccine, but the trend break is not statistically significant, and the coefficients continue to increase after 1989, indicating a general trend.

Figure 4: Piecewise regression: log income




Figure 5: Piecewise regression: employed





Figure 6: Piecewise regression: education



## 7 Controlling for state characteristics

In contrast to Atwood (2022) (Table 4, Panel C), AP do not perform a test for mean reversion, where high and low measles states converge in spite of the vaccine. Following Bleakley (2007), I interact pre-treatment state characteristics with the continuous vaccine exposure variable. I use 1970 census data to calculate state-level averages for income, employment, and education. (Recall that the vaccine was implemented in 1973.)

First, I test for a cross-sectional correlation between the 1970 state characteristics and pre-vaccine average reported measles incidence. This corresponds to Table 3 in AP. I do not find strong correlations, but also do not get results similar to AP.<sup>4</sup>

Table 3: Correlation between measles rate and 1970 state characteristics

|              | Income<br>(1)     | Log income<br>(2)   | Literacy<br>(3)     | Employment<br>(4)    | Education<br>(5)    |
|--------------|-------------------|---------------------|---------------------|----------------------|---------------------|
| $M_{pre}$    | -11.96<br>(12.03) | -0.0007<br>(0.0015) | 0.00009<br>(0.0002) | 0.0002*<br>(0.00008) | -0.0006<br>(0.0025) |
| Observations | 32                | 32                  | 32                  | 32                   | 32                  |
| $R^2$        | 0.02              | 0.01                | 0.00                | 0.12                 | 0.00                |

*Note:*  $M_{pre}$  is the state-level pre-vaccine reported measles rate. Robust standard errors. Significance levels: \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Next I run the main regression including interactions of the state characteristics with vaccine exposure. (Note that the state characteristics are collinear with the state-of-birth fixed effect.) The results are much smaller. The effect on log income disappears entirely, the effect on employment is one-third smaller, and the effect on education is three times smaller. I conclude that reported measles incidence is driven by factors correlated with income, employment, and education.

<sup>4</sup>This is possibly due to using different variables. They report average values for income, literacy, and employment of 5107, 0.72, and 0.49; in my data, these are 8999, 0.78, and 0.86.

Table 4: Interacting with 1970 state characteristics

|                                     | Log income<br>(1)      | Employed<br>(2)        | Education<br>(3)       |
|-------------------------------------|------------------------|------------------------|------------------------|
| $M_{pre} \times \text{Exposure}$    | -0.0141<br>(0.3361)    | 0.3964***<br>(0.1203)  | 7.208***<br>(1.579)    |
| Exposure $\times$ Log income (1970) | -0.0105***<br>(0.0009) | 0.0030***<br>(0.0003)  | -0.0035<br>(0.0045)    |
| Exposure $\times$ Education (1970)  | -0.0018***<br>(0.0005) | -0.0016***<br>(0.0002) | -0.0341***<br>(0.0024) |
| Exposure $\times$ Employment (1970) | 0.1016***<br>(0.0083)  | 0.0013<br>(0.0034)     | 0.3738***<br>(0.0465)  |
| Observations                        | 6,936,812              | 10,633,483             | 10,529,025             |
| R <sup>2</sup>                      | 0.16                   | 0.10                   | 0.26                   |
| Married fixed effects               | ✓                      | ✓                      | ✓                      |
| Urban fixed effects                 | ✓                      | ✓                      | ✓                      |
| Birthyear fixed effects             | ✓                      | ✓                      | ✓                      |
| Birthplace fixed effects            | ✓                      | ✓                      | ✓                      |
| Survey year fixed effects           | ✓                      | ✓                      | ✓                      |

*Note:*  $M_{pre} \times \text{Exposure}$  is the state-level pre-vaccine reported measles rate interacted with exposure to the vaccine. Standard errors (in parentheses) are clustered by state-of-birth interacted with year-of-birth. Survey weights are used. Significance levels: \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

## 8 Redoing the disease event study

In Figures 4 and 5, AP run an event study at the state-year level to test for the effect of the vaccine on measles and other diseases that should not be affected by the vaccine, to serve as a placebo test. Following AP, I collected data on measles and syphilis; I also collected data on pertussis and tetanus, but they did not. Syphilis should not be affected by the measles vaccine, and, following Atwood (2022), pertussis and tetanus should be reduced. AP control for population in their event study, which is only available from 1967 onwards. I omit the control variable to extend the sample back to 1965. I also use calendar time on the x-axis instead of event time, since there's no staggered timing. I follow AP in omitting the year before the vaccine (1972).<sup>5</sup> Standard errors are clustered at the state level.

The measles event study is similar to the original; if the disease was eradicated, we'd expect a coefficient of -1. There was a global measles epidemic in 1976-77, so the increase in coefficients is likely explained by states with high reporting rates reporting relatively more cases during an epidemic. My event study for syphilis is roughly similar to the original, with no clear trend over time. The tetanus results show a slight decline. Pertussis also decreases after the vaccine, but this is potentially due to a pre-existing trend.

---

<sup>5</sup>Note that Atwood (2022) omits the year of the vaccine's introduction, not the year before.



Figure 7: Event study: measles




Figure 8: Event study: syphilis




Figure 9: Event study: tetanus





Figure 10: Event study: pertussis



## References

- Atwood, Alicia. 2022. “The Long-Term Effects of Measles Vaccination on Earnings and Employment.” *American Economic Journal: Economic Policy* 14 (2):34–60. URL <https://www.aeaweb.org/articles?id=10.1257/pol.20190509>.
- Atwood, Alicia and Sarah Pearlman. 2023. “The Long Term Benefits of Vaccination Campaigns: Evidence from Measles.” Working paper. URL [https://web.archive.org/web/20230502002136/https://www.haverford.edu/sites/default/files/Department/Economics/econ-colloquium-Measles\\_Mexico\\_Feb2023.pdf](https://web.archive.org/web/20230502002136/https://www.haverford.edu/sites/default/files/Department/Economics/econ-colloquium-Measles_Mexico_Feb2023.pdf).
- Bleakley, Hoyt. 2007. “Disease and Development: Evidence from Hookworm Eradication in the American South\*.” *The Quarterly Journal of Economics* 122 (1):73–117. URL <https://doi.org/10.1162/qjec.121.1.73>.
- . 2010. “Malaria Eradication in the Americas: A Retrospective Analysis of Childhood Exposure.” *American Economic Journal: Applied Economics* 2 (2):1–45. URL <https://www.aeaweb.org/articles?id=10.1257/app.2.2.1>.
- Chuard, Caroline, Hannes Schwandt, Alexander D Becker, and Masahiko Haraguchi. 2022. “Economic vs. Epidemiological Approaches to Measuring the Human Capital Impacts of Infectious Disease Elimination.” Working Paper 30202, National Bureau of Economic Research. URL <http://www.nber.org/papers/w30202>.
- Roodman, David. 2018. “The Impacts of Hookworm Eradication in the American South. A replication study of Bleakley (The Quarterly Journal of Economics, 2007).” *International Journal for Re-Views in Empirical Economics (IREE)* 2 (2018-3):1–45. URL <http://hdl.handle.net/10419/190496>.
- Wiebe, Michael. 2024. “The Long-Term Effects of Measles Vaccination on Earnings and Employment: Comment.” Working paper. URL [https://michaelwiebe.com/assets/atwood/atwood\\_comment](https://michaelwiebe.com/assets/atwood/atwood_comment).

# A Appendix


## A.1 Distribution of age and birthyear

Following AP, I restrict the data to males with ages between 18-65. The census data reports age, but not birth year, so I infer birth year as census year  $-$  age. We see that the age distribution has substantial bunching at years ending in 0, and some at years ending in 5. There is also a large mass at age 18.

Examining the distribution of birthyear in Figure A2, we see corresponding bunching. There is also bunching from combining the censuses. For example, there is a large mass in 1992 relative to 1993. This corresponds to birthyear 1992 containing 18 year-olds in the 2010 census and 23 year-olds in the 2015 census, while birthyear 1993 contains only 22 year-olds from the 2015 census.


Figure A3 shows that there are spikes in average age, including in the treatment year, 1973. It is unclear whether this pattern can cause bias.

Figure A1: Distribution of age




Note: using survey weights.

Figure A2: Distribution of birthyear



Note: using survey weights.

Figure A3: Average age by birthyear




Note: weighted average using survey weights.



## B Outcome variables in high- and low-measles states

I plot averages of the economic outcomes by above- and below-median prevaccine average reported measles rates against birth year. In contrast to Atwood (2022), here the high-measles states have lower outcomes, and catch up over time.<sup>6</sup>

Figure A4: High vs. Low measles: log income




Note: Monthly wage income, inflation-adjusted using December, 2004 CPI. Zero values are dropped. Average calculated using survey weights. The vertical lines in 1957 and 1973 denote the cohorts with increasing exposure to the vaccine.

<sup>6</sup>For log income and employment, we see spikes in, for example, birth years 1973, 1978, 1983, and 1993. (Recall that the census years are 1990, 1995, 2000, 2010, and 2015.) Figure A3 shows that these are years with spikes in average age. Since younger people are poorer and less likely to be employed, this generates spikes in the average by birth-year.

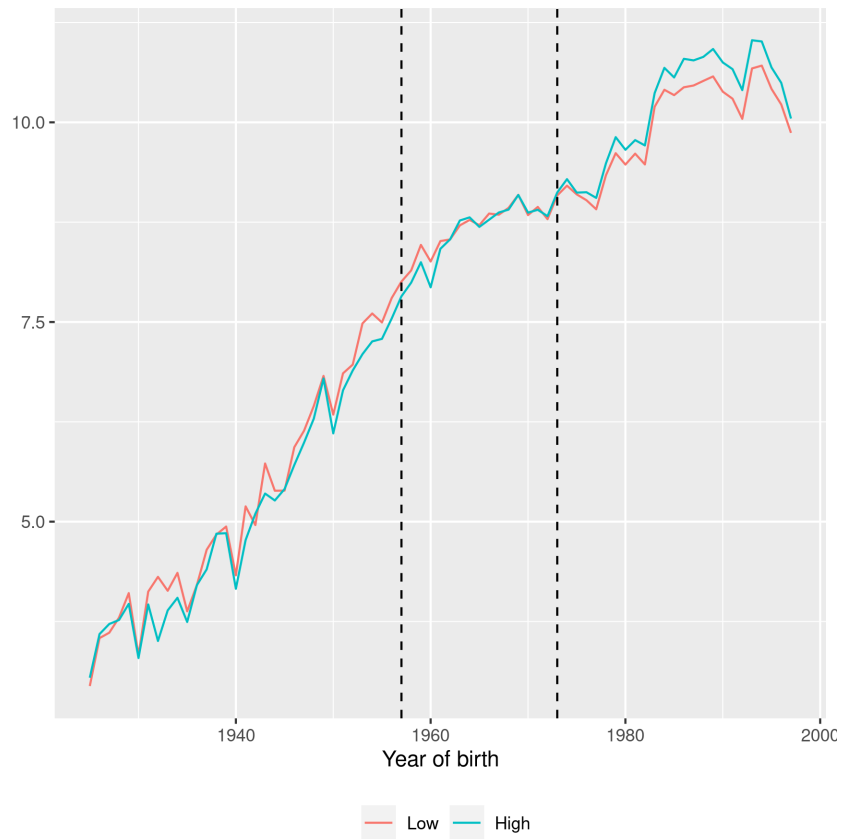
For example, the age composition is (20,30,35) in 1980, (19,29,34) in 1981, (18,28,33) in 1982, and (27,32) in 1983. The average age in these years is 28.2, 27.1, 25.6, and 29.6. This corresponds to the decline in income over 1980-1982, and the jump in 1983.

Figure A5: High vs. Low measles: employed



Note: Average calculated using survey weights. The vertical lines in 1957 and 1973 denote the cohorts with increasing exposure to the vaccine.

Figure A6: High vs. Low measles: education

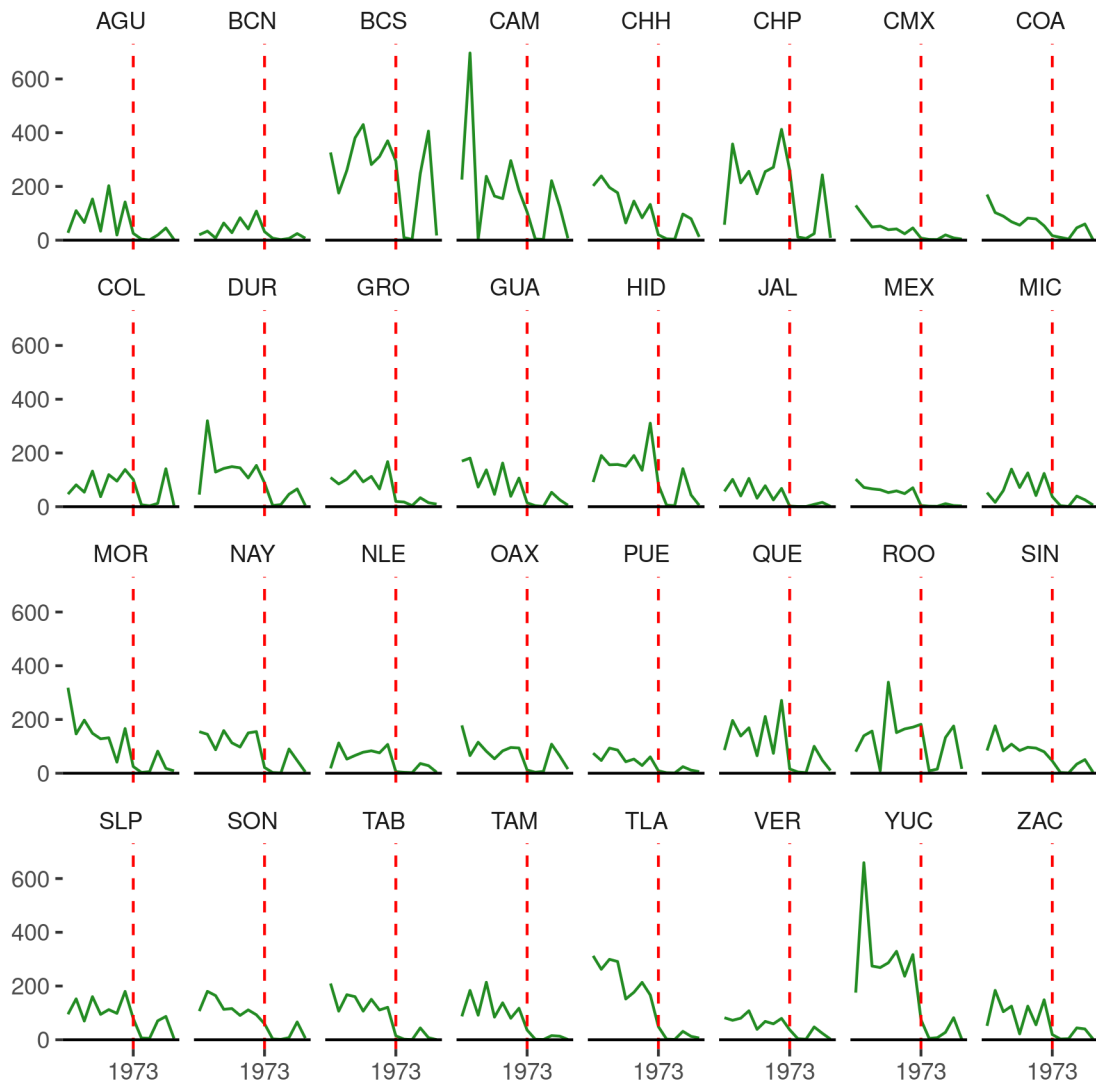


Note: Average calculated using survey weights. The vertical lines in 1957 and 1973 denote the cohorts with increasing exposure to the vaccine.

## B.1 Measles trends by state

I plot reported measles rates by state. As in Figure 2 from Chuard et al. (2022), we see high variation in reported measles incidence across states and over time.

Figure A7: Reported measles rate, by state (1965-1978)



Note: Reported measles rate per 100,000 from 1965-1978 for all Mexican states. The vaccine was introduced in 1973. A measles epidemic occurred in 1976-77.