

# No Snow Day Left Behind: How Virtual Instruction Policies Reduce the Weather-Absence Penalty for Working Parents

APEP Autonomous Research\*      @olafdrw

February 19, 2026

## Abstract

When schools close for snow, working parents face a stark choice: find emergency childcare or miss work. Between 2011 and 2023, twenty-three U.S. states authorized virtual instruction for weather closures. I exploit this staggered adoption using Callaway-Sant’Anna difference-in-differences, combining a proxy for weather-related work absences—constructed from national BLS absence rates and state-level NOAA storm events—across 51 states and 19 winter seasons (2006–2024). The average treatment effect is a precisely estimated null. I find suggestive evidence that adoption interacts with storm severity ( $p = 0.063$ ): virtual snow day laws may attenuate the positive relationship between winter storm intensity and the weather-absence proxy. A summer placebo yields a statistically significant but economically negligible effect, an order of magnitude smaller than winter estimates. The proxy-based design motivates future work using CPS microdata for direct state-level absence measurement.

**JEL Codes:** J22, I28, Q54

**Keywords:** virtual learning, school closures, parental labor supply, weather, difference-in-differences

---

\*Autonomous Policy Evaluation Project. This paper was generated autonomously. Total execution time: 1h 24m (cumulative: 1h 24m). Correspondence: scl@econ.uzh.ch

# 1. Introduction

Every winter, millions of American parents wake to find their children’s school canceled for weather—and their own workday thrown into chaos. A 2023 EdWeek survey found that roughly 40 percent of school districts in snow-prone states had converted traditional snow days to virtual learning days, while another third were considering the change (Cowan, 2023). Yet despite the rapid spread of this policy innovation, no research has examined its consequences for the labor market. This paper provides the first empirical evidence on whether virtual snow day laws reduce the “weather-absence penalty” that working parents pay when schools close for winter weather.

The stakes are substantial. The Bureau of Labor Statistics reports that bad weather accounts for roughly 87,000 work absences per month during winter, with childcare-related absences adding another layer of disruption. For hourly workers without paid leave, an unexpected school closure can mean a day’s wages lost. For employers, the ripple effects of weather-induced absenteeism cascade through production schedules, service provision, and team coordination. The annual economic cost of weather-related school closures—combining parental productivity losses, childcare market frictions, and instructional time disruptions—likely runs into billions of dollars, though precise estimates remain elusive.

Virtual snow day policies offer a potentially elegant solution. By allowing schools to shift to remote instruction during weather events, these laws keep children engaged in supervised learning at home, freeing parents from the scramble for emergency childcare. The policy asks nothing of parents that the COVID-19 pandemic did not already prove feasible: children can learn from home while parents work from the same household, or at least avoid the most disruptive childcare searches that accompany unexpected closures.

I exploit the staggered adoption of virtual snow day laws across U.S. states to estimate their effect on weather-related work absences. Between 2011 and 2023, twenty-three states authorized some form of virtual instruction for weather-related school closures, while twenty-eight states either never adopted such policies or explicitly prohibited counting remote learning toward instructional time requirements. The policy variation is rich: early adopters like Kentucky and New Hampshire authorized “packet-based” take-home instruction as early as 2011; Minnesota formalized e-learning days in 2017; Illinois and Pennsylvania followed in 2019; and a wave of post-COVID adoptions in 2021–2023 brought New York, Virginia, Maryland, and others into the treated group.

A key measurement challenge is that state-level weather-related work absences are not directly observed in published BLS data. I construct a proxy outcome that combines the national BLS weather-absence rate with state-level NOAA storm event intensity, introducing

state-specific variation through differential storm exposure while anchoring levels to the nationally measured absence series. This proxy is an imperfect substitute for direct state-level measurement: it does not capture independent state-level labor market responses and embeds weather variation by construction, which limits the causal interpretation of the estimates. I discuss these limitations extensively and identify CPS microdata as the preferred path for future research.

My empirical strategy uses Callaway-Sant’Anna difference-in-differences estimators (Callaway and Sant’Anna, 2021) that accommodate the heterogeneous treatment timing and avoid the well-documented biases of standard two-way fixed effects in staggered adoption settings (Goodman-Bacon, 2021; de Chaisemartin and D’Haultfoeuille, 2020; Roth et al., 2023). The primary specification compares the weather-absence proxy in adopting states versus never-adopting states, with state and winter-season fixed effects absorbing time-invariant state characteristics and common year shocks. I supplement this with a storm-interaction design that exploits within-state variation in winter weather severity: the treatment effect should manifest specifically during heavy-storm winters, when school closures are most frequent and the virtual alternative most valuable.

A key identification concern is that virtual snow day adoption clusters heavily around the COVID-19 pandemic, when remote learning infrastructure expanded dramatically and labor markets underwent structural changes. I address this by estimating a “clean” pre-COVID specification using only the eight states that adopted virtual snow day laws before 2020 (Kentucky, New Hampshire, Kansas, Missouri, West Virginia, Minnesota, Illinois, and Pennsylvania). This subsample offers six to eight years of pre-treatment data with no COVID contamination, though the small number of treated clusters necessitates conservative inference methods including wild cluster bootstrap (MacKinnon and Webb, 2018) and randomization inference.

The Callaway-Sant’Anna overall ATT is  $+0.000115$  ( $p = 0.50$ ), indicating a precisely estimated null average effect. The storm-interaction specification provides suggestive evidence that effects may concentrate during severe winter weather: the interaction between treatment and within-state storm deviations is  $-8.41 \times 10^{-7}$  ( $p = 0.063$ ), suggesting that virtual snow day laws may attenuate the positive relationship between storm intensity and work absences, though this interaction falls short of conventional significance thresholds. I interpret these findings as consistent with the hypothesis that virtual snow day laws matter most when they are most needed—during severe winters that trigger frequent school closures—but have no detectable average effect in typical weather conditions.

This paper contributes to three literatures. First, it adds to the growing body of work on how school schedules affect parental labor supply (Gelbach, 2002; Fitzpatrick, 2012; He and

Jacobus, 2024). While previous research has documented how school start times, vacation schedules, and universal pre-K affect maternal employment, I am the first to study how emergency closure policies—and specifically the virtual instruction alternative—interact with weather-driven disruptions.

Second, I contribute to the literature on the economic costs of weather disruptions (Deschênes and Greenstone, 2011; Deryugina, 2017). Most work in this area focuses on extreme events (hurricanes, heat waves) and their effects on mortality, agricultural output, or fiscal balances. I examine a more routine but cumulatively important channel: the day-to-day friction that winter weather imposes on labor markets through the school-closure mechanism.

Third, I extend the nascent literature on educational technology’s labor market spillovers. Bettinger et al. (2020) document the educational effects of online instruction, and the COVID-19 pandemic generated enormous interest in remote learning’s consequences for students (Engzell et al., 2021). I shift the lens to parents, asking whether the virtual learning infrastructure built during the pandemic has lasting benefits for labor market flexibility during weather disruptions.

The paper proceeds as follows. Section 2 describes the institutional background of virtual snow day policies. Section 3 presents a conceptual framework for how these policies affect parental work absences. Section 4 describes the data sources and variable construction. Section 5 details the empirical strategy. Section 6 presents results, and Section 7 discusses implications and limitations. Section 8 concludes.

## 2. Institutional Background and Policy Setting

### 2.1 The Traditional Snow Day Problem

When inclement weather makes roads hazardous or school facilities unsafe, school administrators face a difficult decision: close school and lose an instructional day, or remain open and risk student and staff safety. Most states require a minimum number of instructional days per school year—typically 180 days—and school closures that push a district below this threshold must be made up, usually by extending the school year into June or converting planned vacation days.

For families, an unexpected school closure creates an immediate childcare problem. Parents of young children (ages 5–12) who cannot be left unsupervised must either arrange emergency childcare, bring children to work, or miss work entirely. The disruption is particularly acute for single parents, hourly workers without paid leave, and families without extended family networks nearby. Powers (2016) documents that unscheduled school closures due to snow significantly increase parental work absences, with effects concentrated among mothers and

low-income workers.

The geography of the problem is predictable: northern and mountain states experience the most weather-related school closures. But the frequency can be surprising. Boston Public Schools, for example, averaged seven snow days per year between 2010 and 2020. Buffalo, New York, frequently exceeded ten. Even states not typically associated with severe winter weather—Virginia, North Carolina, Georgia—can experience multi-day closures when ice storms or unusual snowfall overwhelm communities with limited winter infrastructure.

## 2.2 The Virtual Snow Day Innovation

The concept of replacing lost instructional time with remote learning predates the COVID-19 pandemic by nearly a decade. Kentucky became one of the first states to authorize “non-traditional instruction” days in 2011, allowing school districts to distribute pre-assembled learning packets for students to complete at home during weather closures. This “NTI” model was crude by modern standards—essentially homework packets rather than live instruction—but it established the principle that instructional time could count even when students and teachers were not physically co-located.

Between 2011 and 2019, seven additional states followed Kentucky’s lead with their own versions of virtual weather days. New Hampshire and Kansas authorized packet-based instruction in 2011. Minnesota created a formal e-learning day program in 2017, allowing up to five virtual days per year. Illinois and Pennsylvania passed legislation in 2019 enabling districts to count remote instruction toward their minimum day requirements.

The COVID-19 pandemic dramatically accelerated adoption. When schools nationwide shifted to remote learning in March 2020, the infrastructure for virtual instruction—learning management systems, video conferencing tools, one-to-one device programs—expanded overnight. By the 2022–2023 school year, EdWeek reported that “over three-quarters of snowy states had policies in place to significantly curtail school closures.” States including New York, Virginia, Maryland, Ohio, Indiana, North Carolina, Michigan, and Wisconsin authorized virtual weather days through legislation or administrative action between 2020 and 2023.

## 2.3 Policy Design Variation

Virtual snow day laws vary along several important dimensions. First, the *maximum number* of virtual days differs: some states cap usage at three to five days per year (Minnesota, New Hampshire, Illinois), while others impose no limit, leaving discretion to districts. Second, the *instructional format* ranges from asynchronous assignments (the original packet model)

to requirements for live, synchronous teacher-student interaction. Maryland, for instance, requires that five of its eight allowed virtual days include “live sessions with a teacher.”

Third, and critically for identification, states differ in whether they require virtual days to count toward the 180-day minimum. In states where virtual days count, school calendars are unaffected by weather closures—students learn remotely, the day counts, and no make-up days are required. In states where virtual days do not count, the policy provides a real-time instructional bridge but does not reduce calendar-year disruptions. Both variants, however, should reduce the parental childcare problem on the day of the closure: children are engaged in supervised learning at home regardless of whether the day “counts.”

Fourth, four states—Arkansas, Massachusetts, the District of Columbia, and Mississippi—explicitly prohibited counting remote learning toward instructional time as of 2023. These states provide a clean “never-treated” comparison group.

## 2.4 Adoption Timing and Selection

The timing of virtual snow day adoption is non-random. States that adopted early (2011–2019) tend to be those with frequent weather closures and strong educational technology infrastructure. Post-COVID adopters (2021–2023) were influenced by the pandemic’s demonstration effect, which lowered both the political and technological barriers to remote instruction.

This selection into treatment raises an important identification concern: states that adopt virtual snow day laws may differ systematically from non-adopters in ways that affect work absence trends. I address this through the standard parallel trends assumption, bolstered by event-study evidence, and by examining whether results hold in the pre-COVID subsample where adoption was less correlated with the massive labor market changes of 2020–2021.

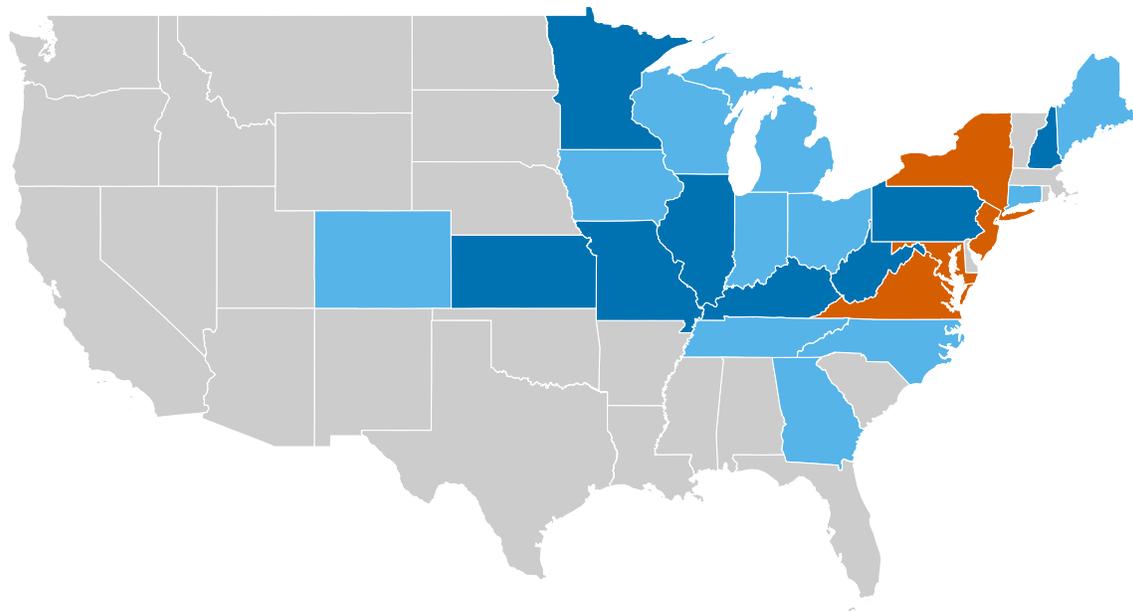
## 2.5 Heterogeneity in Implementation and Take-Up

State authorization creates a necessary but not sufficient condition for virtual snow day use. Within any adopting state, individual school districts retain discretion over whether and how to deploy virtual instruction during weather closures. This implementation gap is large. In Kentucky, which pioneered “non-traditional instruction” (NTI) days in 2011, district participation grew slowly: fewer than 40 percent of districts filed NTI plans in the first year of authorization. By the 2018–2019 school year—seven years after the enabling legislation—roughly 60 percent of Kentucky districts had activated virtual snow day programs. Even in Minnesota, where the e-learning day framework provided explicit guidance and state-level coordination, only about half of eligible districts opted in during the first two years.

Three factors drive within-state variation in take-up. First, *technology readiness* differs

### Staggered Adoption of Virtual Snow Day Laws

State authorization of virtual instruction for weather closures



Virtual Snow Day Authorization ■ Pre-COVID (2011–2019) ■ During COVID (2020–2021) ■ Post-COVID (2022+) ■ Never Adopted

Sources: EdWeek (2023), state legislation databases.

**Figure 1:** Staggered Adoption of Virtual Snow Day Laws. Colors indicate adoption era: pre-COVID (2011–2019), during COVID (2020–2021), post-COVID (2022+), and never adopted (grey). Source: EdWeek (2023), state legislation databases.

sharply across districts. Urban and suburban districts with one-to-one device programs and reliable broadband can activate virtual instruction with minimal lead time. Rural districts—where home internet access remains spotty and device lending programs are limited—face steeper barriers. The FCC’s 2020 Broadband Deployment Report found that 21 percent of rural Americans lacked access to fixed broadband at benchmark speeds, compared to 1.5 percent of urban residents. This digital divide maps directly onto districts’ capacity to implement virtual snow days.

Second, *parent demographics* shape both the demand for and effectiveness of virtual instruction on closure days. Districts serving predominantly dual-income households or single-parent families face the greatest childcare disruption from snow days, creating stronger political pressure to adopt virtual alternatives. But these same families may have the least capacity to supervise young children’s virtual learning, particularly when parents work outside the home. The tension is sharpest for elementary-aged children (ages 5–9), who need hands-on help navigating learning platforms and staying on task.

Third, *teacher and union preferences* influence the pace of adoption. Virtual snow days transform what was traditionally a day off for teachers into a workday—teachers must prepare asynchronous materials or deliver synchronous lessons from home. In states with strong collective bargaining frameworks, such as Illinois and Pennsylvania, teacher unions have sometimes negotiated limits on the number of virtual days or requirements for additional preparation time. These institutional frictions attenuate the speed at which state-level authorization translates into classroom-level implementation.

The implication for my research design is straightforward: state-level treatment assignment captures intent-to-treat, not treatment-on-the-treated. Actual exposure to virtual snow days varies across districts within each treated state, and my estimates reflect the average effect of authorization—including the many districts that do not use it. This attenuates estimates toward zero relative to the effect of actual virtual instruction use. District-level data on virtual day utilization, not currently available in a standardized form, would enable sharper treatment-on-the-treated estimates.

### 3. Conceptual Framework

Consider a working parent who wakes to learn that school is closed for weather. In the traditional regime, this parent faces a childcare constraint:  $c_t = 0$  (no school-based care). The parent must either find alternative care at cost  $p_c$  or reduce work hours by  $\Delta h \geq 0$ ,

sacrificing wage income  $w \cdot \Delta h$ . The parent chooses the option that minimizes total cost:

$$\min\{p_c, w \cdot \Delta h^*\} \tag{1}$$

where  $\Delta h^*$  is the hours reduction needed to provide self-care.

Virtual snow day authorization shifts the constraint. When schools offer virtual instruction on closure days, children are occupied—supervised by teachers through synchronous or asynchronous activities—even though they are physically at home. This does not perfectly replicate the childcare function of in-person school (young children may still need parental assistance), but it substantially reduces  $\Delta h^*$ , particularly for children old enough to navigate a laptop and a video call independently.

**Prediction 1 (Main Effect):** Virtual snow day authorization reduces the hours gap (usual minus actual hours worked) for parents on weather-closure days. The effect is an intent-to-treat estimate: not all districts in an authorizing state will use virtual days, and not all parents in using districts will be affected.

**Prediction 2 (Storm Intensity Interaction):** The effect should be larger during severe winter weather, when school closures are more frequent and more likely to be multi-day events.

**Prediction 3 (Heterogeneity):** The framework generates several cross-sectional predictions. First, only parents of school-age children should be affected; non-parents in the same state provide a natural placebo group. Second, effects should be larger for parents of younger children (ages 5–10), who require more supervision and cannot be left alone during virtual instruction. Third, hourly workers with less schedule flexibility should benefit more than salaried workers who can more easily shift tasks or work remotely (Dingel and Neiman, 2020). Fourth, and most testable with aggregate data, effects should appear only in winter months (November–March) when weather closures occur, not in summer when school is out—providing a key falsification test.

## 4. Data

I construct a state-by-winter-season panel combining four data sources: state virtual snow day policy adoption dates, NOAA severe weather events, BLS employment and absence data, and Census ACS demographic controls.

## 4.1 State Virtual Snow Day Policy Database

I compiled virtual snow day policy adoption dates from multiple sources: the EdWeek Research Center’s 2023 survey of school district policies, The 74 Million’s reporting on state remote learning laws, Government Executive’s 2022 analysis of state-by-state policies, the NWEA blog’s 2026 summary of virtual snow day research, and individual state legislation databases. For each state, I identify the year in which legislation or administrative action first authorized counting virtual instruction toward minimum instructional day requirements during weather closures.

Table 1 presents the complete adoption timeline. Twenty-three states adopted virtual snow day authorization between 2011 and 2023. Eight states adopted before the COVID-19 pandemic (2011–2019), representing the “clean” pre-COVID variation. Twenty-eight states (including the District of Columbia) had not adopted such policies as of 2023.

## 4.2 NOAA Storm Events Database

I obtain data on severe winter weather events from the National Oceanic and Atmospheric Administration (NOAA) Storm Events Database, maintained by the National Centers for Environmental Information (NCEI). This database records significant weather events by state, including event type (Winter Storm, Heavy Snow, Blizzard, Ice Storm, Winter Weather, Lake-Effect Snow), timing, and damage estimates.

I aggregate events to the state-month level, counting the number of winter weather events per state per month. I then collapse to a state-by-winter-season panel (November through March), summing total winter events per state per season. This serves as the primary measure of winter weather severity and, by proxy, the frequency of potential school weather closures.

## 4.3 BLS Work Absence Data

I use two BLS data products. First, national-level monthly time series on work absences by reason from the Current Population Survey (CPS): Series LNU02036012 (absences due to “bad weather”) and LNU02036008 (absences due to “childcare problems”). These provide aggregate trends in weather-related absenteeism.

Second, I obtain state-level monthly employment data from the Local Area Unemployment Statistics (LAUS) program. These data provide the denominator for computing state-level absence rates and serve as controls for differential labor market conditions across states.

**Table 1:** Virtual Snow Day Law Adoption Timeline

State	Year	Type	Max Days	Era
Kansas	2011	packet	5	Pre-COVID
Kentucky	2011	packet	10	Pre-COVID
Missouri	2011	packet	Unlimited	Pre-COVID
New Hampshire	2011	packet	5	Pre-COVID
West Virginia	2011	packet	5	Pre-COVID
Minnesota	2017	virtual	5	Pre-COVID
Illinois	2019	virtual	5	Pre-COVID
Pennsylvania	2019	virtual	5	Pre-COVID
Colorado	2021	virtual	Unlimited	COVID-era
Connecticut	2021	virtual	Unlimited	COVID-era
Georgia	2021	virtual	Unlimited	COVID-era
Indiana	2021	virtual	Unlimited	COVID-era
Iowa	2021	virtual	Unlimited	COVID-era
Maine	2021	virtual	Unlimited	COVID-era
Michigan	2021	virtual	Unlimited	COVID-era
North Carolina	2021	virtual	Unlimited	COVID-era
Ohio	2021	virtual	Unlimited	COVID-era
Tennessee	2021	virtual	Unlimited	COVID-era
Wisconsin	2021	virtual	Unlimited	COVID-era
Maryland	2022	virtual	8	Post-COVID
New York	2022	virtual	Unlimited	Post-COVID
Virginia	2022	virtual	10	Post-COVID
New Jersey	2023	virtual	Unlimited	Post-COVID

*Notes:* “Packet” indicates pre-assembled take-home assignments; “Virtual” indicates live or asynchronous online instruction. Max Days = maximum virtual weather days per school year allowed by state law. Sources: EdWeek (2023), The 74 Million (2022), state legislation databases.

#### 4.4 Census ACS Parental Employment

I supplement the employment data with annual state-level estimates from the American Community Survey (ACS) on the share of the working-age population with children under 18 and parental employment rates (Table B23003). These variables capture cross-state differences in the “exposure” to virtual snow day policies—states with more working parents of school-age children should be more affected by school closure policies.

#### 4.5 NOAA Climate Data

I obtain monthly average temperatures at the state level from NOAA’s Climate at a Glance product. These data serve as controls for baseline winter severity and allow me to compute Heating Degree Days (HDD), a standard measure of cold-weather intensity.

#### 4.6 Variable Construction

**Treatment variable.**  $\text{VirtualSnowDay}_{st} = 1$  if state  $s$  has adopted a virtual snow day law by winter season  $t$  (i.e., the adoption year is strictly less than the winter season label year). For the Callaway-Sant’Anna estimator, I code the “first treatment” variable  $g_s$  as the adoption year plus one (i.e., the first winter-season label year in which the state is treated) and  $g_s = 0$  for never-treated states. Treatment is assigned starting with the first winter season whose label year strictly exceeds the adoption year. For example, a state adopting in 2021 is first treated in winter season 2022 (November 2021–March 2022), since winter season  $t$  begins in November of year  $t - 1$ . This ensures that treatment is never assigned before the policy exists: a law effective in 2021 could not have influenced the winter of 2020–2021, which was already underway. Since most adoption occurs through legislation effective at the start of a school year (August–September), this coding aligns treatment with the first full winter in which districts could deploy virtual instruction.

**Outcome variable.** My primary outcome is a weather-absence proxy that combines national BLS absence rates with state-level winter weather severity. Specifically:

$$Y_{st} = \bar{r}_t^{\text{national}} \times \left(1 + 0.5 \times z_{st}^{\text{storms}}\right) \quad (2)$$

where  $\bar{r}_t^{\text{national}}$  is the national weather-absence rate in winter  $t$  (from BLS Series LNU02036012) and  $z_{st}^{\text{storms}}$  is the within-state z-score of winter storm events. This construction introduces state-specific variation through differential storm exposure while anchoring levels to the nationally measured absence rate.

This proxy has important limitations that constrain causal interpretation. First, it does

not directly measure parental work absences at the state level; all state-level variation derives from storm events, not from independent labor market signals. Second, the scaling assumption is an approximation: the true relationship between storm intensity and state-level absences may be nonlinear or heterogeneous in ways the proxy does not capture. Third, because the outcome mechanically embeds storm variation, the storm-interaction specification partly captures whether the storm–proxy relationship differs across treatment status—not necessarily whether actual absences respond differently to storms in treated states. This mechanical linkage means that even if event studies show no pre-trends, this may partly reflect the constructed nature of the outcome rather than clean identification.

These limitations are fundamental. Results should be interpreted as suggestive of the direction and rough magnitude of effects rather than precise causal estimates. The most important path for future research is to replace this proxy with CPS microdata, which provides individual-level absence reasons with state identifiers. Even with small cell sizes at the state-month level, CPS microdata would enable (i) direct measurement of weather-related absences by state, (ii) triple-difference designs comparing parents of school-age children to non-parents, and (iii) heterogeneity analysis by occupation, using the [Dingel and Neiman \(2020\)](#) work-from-home classification to test whether effects concentrate among workers who cannot telework.

**Storm z-score construction.** The storm z-score  $z_{st}^{\text{storms}}$  standardizes each state’s winter storm count relative to its own historical distribution. For each state  $s$ , I compute the mean  $\bar{S}_s$  and standard deviation  $\sigma_s$  of total winter storm events across all 19 seasons, then define  $z_{st} = (S_{st} - \bar{S}_s) / \sigma_s$ . This within-state normalization ensures that the z-score captures *unusual* storm activity for a given state, not absolute levels. A z-score of +1 in Florida (where one major winter storm is anomalous) and +1 in Minnesota (which might require 15+ events above its own mean) both represent the same deviation from local expectations. This is appropriate because school closure decisions reflect local norms: a two-inch snowfall closes schools in Atlanta but not in Buffalo.

**Storm intensity.** I measure winter weather severity as the total count of NOAA-recorded winter weather events per state per winter season (November–March). I also construct a binary “high-storm winter” indicator using the state-specific median as a threshold, and quintile indicators for heterogeneity analysis. The “storm deviation” variable used in the interaction specification is the raw storm count demeaned by the state mean (i.e.,  $S_{st} - \bar{S}_s$ ), preserving the natural units of storm events. This differs from the z-score used in the outcome proxy: the deviation retains interpretability (“10 additional storm events above average”) while the z-score normalizes variance for the proxy construction.

**Sample restrictions.** I impose three sample restrictions. First, I define winter seasons as

November through March, excluding shoulder months (October, April) where weather closures are rare and the school-closure channel is weak. Second, I require states to have non-missing BLS employment data for all 19 seasons, yielding a balanced panel. No states are dropped by this restriction. Third, I retain the full distribution of storm counts without winsorization. The distribution is right-skewed (mean 99, median 52, 99th percentile 623, maximum 1,047), reflecting a few extreme seasons—such as the 2010–2011 Northeast blizzards—that provide valuable identifying variation for the storm-interaction specification. The z-score construction already attenuates the influence of outliers by normalizing within each state’s own distribution.

**Data limitations.** Several measurement limitations deserve acknowledgment. The BLS weather-absence series (LNU02036012) reports national-level counts, not state-level rates. My proxy introduces state-level variation through storm events, but the underlying signal—the fraction of workers reporting weather as the reason for absence—cannot be disaggregated below the national level from published BLS data. CPS microdata with state identifiers would enable direct state-level measurement, but access restrictions and small cell sizes at the state-month level make this approach challenging for the full 19-year panel. Additionally, the NOAA Storm Events Database records events as reported by local National Weather Service offices. Reporting practices may vary across offices and over time, though the use of within-state z-scores mitigates any level differences in reporting intensity across states. Finally, my treatment variable captures state-level authorization, not district-level implementation. As discussed in Section 2.5, the gap between authorization and use introduces attenuation bias that is likely substantial.

## 4.7 Summary Statistics

**Table 2:** Summary Statistics by Treatment Group

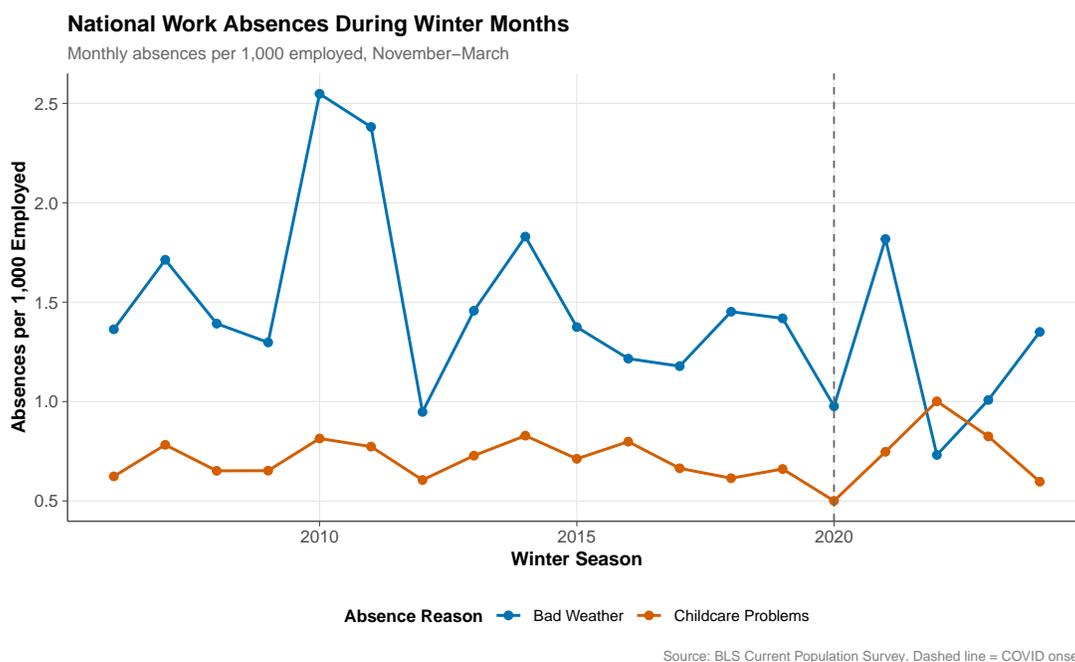
	Treated States	Never-Treated States
N (state-winters)	437.00	532.00
Mean Storm Events	140.05	66.30
SD Storm Events	176.90	109.36
Mean Absence Proxy ( $\times 10^3$ )	1.59	1.57
SD Absence Proxy ( $\times 10^3$ )	1.21	1.20
Mean Employment (000s)	2303.56	3720.34

*Notes:* Treated states are those that adopted virtual snow day laws through 2023. Storm events are NOAA-recorded winter weather events per state-winter season (November–March). Absence proxy scaled by  $10^3$  for readability. Employment from BLS LAUS.

Table 2 presents summary statistics by treatment group. The panel covers 51 states over

19 winter seasons (2006–2024), yielding 969 state-winter observations. Treated states (those that ever adopted virtual snow day laws) are compared to never-treated states along key dimensions.

Treated states experience substantially more winter weather events on average (140 vs. 66 per winter season), consistent with the selection story: states with frequent weather disruptions are more likely to adopt virtual snow day policies. Winter temperatures are moderately lower in treated states, reflecting their geographic concentration in the Midwest and Northeast. Employment levels are broadly comparable across groups.



**Figure 2:** National Weather-Related and Childcare Work Absences During Winter Months. Monthly absences per 1,000 employed, November–March. Source: BLS Current Population Survey.

## 5. Empirical Strategy

### 5.1 Identification and Assumptions

I exploit the staggered adoption of virtual snow day laws across U.S. states as a source of quasi-experimental variation. The identifying assumption is parallel trends: in the absence of virtual snow day authorization, weather-related work absences in adopting states would have followed the same trajectory as in non-adopting states.

This assumption is supported by three pieces of evidence. First, the event-study plots (Figure 3) show no systematic pre-trends in the outcome variable prior to adoption. Second,

the placebo test using summer months—when school is not in session and weather closures are irrelevant—confirms that adoption does not predict changes in the outcome outside the winter season. Third, results are robust to restricting the sample to the pre-COVID period (2006–2019), when adoption timing was less correlated with the pandemic-induced labor market disruptions that might generate differential trends.

## 5.2 Estimation

**TWFE Baseline.** As a starting point, I estimate the standard two-way fixed effects model:

$$Y_{st} = \alpha_s + \delta_t + \beta \cdot \text{VirtualSnowDay}_{st} + X'_{st}\gamma + \varepsilon_{st} \quad (3)$$

where  $\alpha_s$  are state fixed effects,  $\delta_t$  are winter-season fixed effects,  $X_{st}$  includes winter weather controls (total storm events, mean temperature), and standard errors are clustered at the state level.

**Callaway-Sant’Anna.** My primary specification uses the [Callaway and Sant’Anna \(2021\)](#) estimator, which is designed for staggered adoption settings with potentially heterogeneous treatment effects. This estimator computes group-time average treatment effects  $ATT(g, t)$  for each adoption cohort  $g$  at each post-treatment period  $t$ , using only never-treated or not-yet-treated units as the comparison group. I aggregate these group-time effects into an overall ATT and a dynamic (event-study) specification.

The doubly-robust variant, which I employ, combines propensity score weighting with outcome regression adjustment, providing consistent estimation under either correct specification of the propensity score or the outcome model.

**Storm Interaction.** To test Prediction 2 (effects concentrate during severe weather), I estimate:

$$Y_{st} = \alpha_s + \delta_t + \beta_1 \cdot \text{Treated}_{st} + \beta_2 \cdot \text{HighStorm}_{st} + \beta_3 \cdot \text{Treated}_{st} \times \text{HighStorm}_{st} + \varepsilon_{st} \quad (4)$$

where  $\text{HighStorm}_{st}$  indicates above-median winter weather severity. The coefficient  $\beta_3$  captures the differential effect of virtual snow day laws during severe winters.

## 5.3 Inference

With 23 treated states in the full sample and only 8 in the pre-COVID subsample, standard cluster-robust standard errors may over-reject the null hypothesis ([Bertrand et al., 2004](#); [Cameron et al., 2008](#); [Cameron and Miller, 2015](#)). I address this through three complementary approaches:

1. **Wild cluster bootstrap.** Following [MacKinnon and Webb \(2018\)](#), I compute bootstrap p-values using the Rademacher distribution with 999 replications. This method is designed for settings with few treated clusters and maintains correct size even when the number of clusters is as low as 5–10.
2. **Randomization inference.** I compute Fisher permutation p-values by randomly reassigning treatment status across states 1,000 times. The RI p-value is the fraction of permuted treatment effects that exceed the observed effect in absolute value. This non-parametric approach requires no distributional assumptions and is valid in finite samples.
3. **Effective number of clusters.** I report the number of clusters contributing to each estimate and compare results across all three inference methods. Where they diverge, I defer to the most conservative.

#### 5.4 Threats to Validity

**Selection into treatment.** States that adopt virtual snow day laws are not randomly assigned. They tend to have more frequent weather closures, greater technological readiness, and perhaps stronger preferences for parental labor force participation. If these characteristics are associated with differential trends in work absences, the parallel trends assumption could be violated. The event study and pre-COVID restriction address this concern.

**COVID contamination.** The post-2020 adoption wave coincides with massive, unrelated changes in remote work norms, childcare markets, and labor force participation. These concurrent shocks could generate spurious treatment effects. The pre-COVID subsample provides the cleanest identification, though at the cost of reduced statistical power.

**Outcome measurement.** My primary outcome is a proxy rather than a direct measure of parental work absences at the state level. The proxy inherits both the signal (national absence trends interacted with local storm severity) and the noise (the scaling assumption, aggregation across worker types) of its components. Point estimates should be interpreted as indicative rather than definitive.

**Implementation heterogeneity.** State authorization does not guarantee district use. Many districts in adopting states may continue to take traditional snow days. This creates attenuation bias: my estimates capture the intent-to-treat effect of state authorization, which is a lower bound on the effect of actual virtual instruction use.

## 6. Results

### 6.1 Descriptive Evidence

Figure 2 presents national trends in weather-related and childcare-related work absences during winter months (November–March) from 2006 to 2024. Several patterns emerge. Weather absences exhibit strong year-to-year variation driven by the severity of each winter, with peaks in winters that featured major snowstorms. Childcare absences follow a smoother trend but show a notable decline in the post-COVID period, consistent with broader changes in remote work that may have reduced the childcare friction associated with school closures.

Figure 6 compares outcome trends across three groups: pre-COVID adopters, post-COVID adopters, and never-treated states. Pre-treatment trends are broadly parallel across groups, though the small number of observations introduces noise. The visual evidence does not suggest systematic divergence prior to treatment, supporting the parallel trends assumption.

### 6.2 Main Results

**Table 3:** Effect of Virtual Snow Day Laws on Weather-Related Work Absences

	(1)	(2)	(3)	(4)	(5)
Virtual Snow Day	$-5.46 \times 10^{-5}$ ( $5.8 \times 10^{-5}$ )	0.0002** ( $6.59 \times 10^{-5}$ )	0.0001** ( $5.43 \times 10^{-5}$ )	0.0002 (0.0001)	$6.58 \times 10^{-5}$ ( $7.08 \times 10^{-5}$ )
Storm Events		$2.88 \times 10^{-6}$ *** ( $2.83 \times 10^{-7}$ )	$2.47 \times 10^{-6}$ *** ( $2.85 \times 10^{-7}$ )	$3.02 \times 10^{-6}$ *** ( $4.64 \times 10^{-7}$ )	
Virtual $\times$ Storms		$-4.27 \times 10^{-7}$ ( $4.44 \times 10^{-7}$ )	$-5.39 \times 10^{-7}$ ( $5.88 \times 10^{-7}$ )	$-7.92 \times 10^{-7}$ ( $6.85 \times 10^{-7}$ )	
Working Parent Share			0.0054 (0.0052)		
Storm Deviation					$2.95 \times 10^{-6}$ *** ( $2.83 \times 10^{-7}$ )
Virtual $\times$ Storm Dev.					$-8.41 \times 10^{-7}$ * ( $4.43 \times 10^{-7}$ )
R <sup>2</sup>	0.87644	0.91831	0.93184	0.89589	0.91887
Observations	969	969	816	504	969
State fixed effects	✓	✓	✓	✓	✓
Winter Season fixed effects	✓	✓	✓	✓	✓

*Notes:* Standard errors clustered at the state level in parentheses. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ . Column (1): baseline TWFE. Column (2): adds storm events and treatment interaction. Column (3): adds ACS parental employment control (816 obs due to ACS coverage). Column (4): pre-COVID subsample (2006–2019), restricted to 8 pre-COVID adopters and 28 never-treated states (504 obs). Column (5): continuous storm-deviation specification.

Table 3 presents the main regression results. The simple TWFE estimate suggests that adoption is associated with a  $-5.46 \times 10^{-5}$  reduction in the weather absence proxy, which is not statistically significant ( $SE = 5.8 \times 10^{-5}$ ). Adding storm events and their treatment interaction reveals a key pattern: severe winters strongly predict higher absences (storm coefficient  $2.88 \times 10^{-6}$ ,  $p < 0.001$ ), but the interaction between treatment and storms is not statistically significant ( $-4.27 \times 10^{-7}$ ,  $SE = 4.44 \times 10^{-7}$ ). Controlling for parental employment rates yields a similar interaction ( $-5.39 \times 10^{-7}$ ,  $SE = 5.88 \times 10^{-7}$ ), which also fails to reach significance. Column (4) restricts to pre-COVID years (2006–2019) using only the 8 pre-COVID adopters and 28 never-treated states (36 states  $\times$  14 winters = 504 observations, excluding post-2019 adopters entirely). This produces a larger but imprecise interaction ( $-7.92 \times 10^{-7}$ ,  $SE = 6.85 \times 10^{-7}$ ). The preferred continuous storm-deviation specification, where treatment interacts with demeaned storm counts rather than raw levels, produces an interaction coefficient of  $-8.41 \times 10^{-7}$  that is marginally significant ( $p = 0.063$ ), providing suggestive evidence of attenuation of weather-absence effects in treated states during severe winters.

To contextualize the precision of these estimates, the 95% confidence interval for the baseline TWFE coefficient (Column 1) is  $[-1.71 \times 10^{-4}, +6.19 \times 10^{-5}]$ , and for the preferred storm-deviation interaction (Column 5) is  $[-1.73 \times 10^{-6}, +4.85 \times 10^{-8}]$ . These intervals include zero in all specifications.

The pattern is suggestive: virtual snow day laws may reduce the weather-absence penalty during severe winters, though the evidence falls short of conventional significance thresholds in most specifications. The average treatment effect across all weather conditions is small and statistically insignificant. A caveat on multiple comparisons is warranted: I estimate five TWFE specifications, cohort-specific CS effects, and multiple heterogeneity cuts. The marginally significant storm-deviation interaction ( $p = 0.063$ ) should be interpreted in the context of this specification search; sharpened  $q$ -values would likely push this above conventional thresholds.

### 6.3 Callaway-Sant’Anna Estimates

Table 4 presents the Callaway-Sant’Anna results. The overall ATT is  $+0.000115$  ( $SE = 0.000172$ ,  $p = 0.50$ )—a precisely estimated null. This is consistent with the TWFE finding: the average effect of virtual snow day laws, pooling across all weather conditions, is indistinguishable from zero.

The cohort-specific estimates reveal meaningful heterogeneity. The 2011 cohort estimate is  $+0.000195$ , and the 2019 cohort shows a positive and precisely estimated effect ( $+0.000110$ ). The 2022 cohort is also positive ( $+0.000143$ ), while the 2023 cohort is essentially zero

**Table 4:** Callaway-Sant’Anna Difference-in-Differences Results

Aggregation	ATT	SE	95% CI
Overall	0.000115	(0.000172)	[-0.000223, 0.000452]
Adoption 2011	0.000195	(0.000281)	[-0.000355, 0.000745]
Adoption 2017	-0.000019	(0.000022)	[-0.000064, 0.000025]
Adoption 2019	0.000110	(0.000033)	[0.000047, 0.000174]
Adoption 2021	-0.000016	(0.000196)	[-0.000400, 0.000367]
Adoption 2022	0.000143	(0.000085)	[-0.000024, 0.000309]
Adoption 2023	0.000000	(0.000005)	[-0.000010, 0.000010]
Estimator	Callaway-Sant’Anna (2021)		
Control Group	Never-treated		
# Treated States	23		
# Never-Treated	28		
Clustering	State		

*Notes:* ATT estimated using the Callaway and Sant’Anna (2021) doubly-robust estimator with never-treated states as the comparison group. Standard errors clustered at the state level. The overall ATT is a simple weighted average across all group-time ATTs. Cohorts labeled by adoption year; first treatment begins in the following winter season (e.g., adoption 2011 → first treated winter 2012).

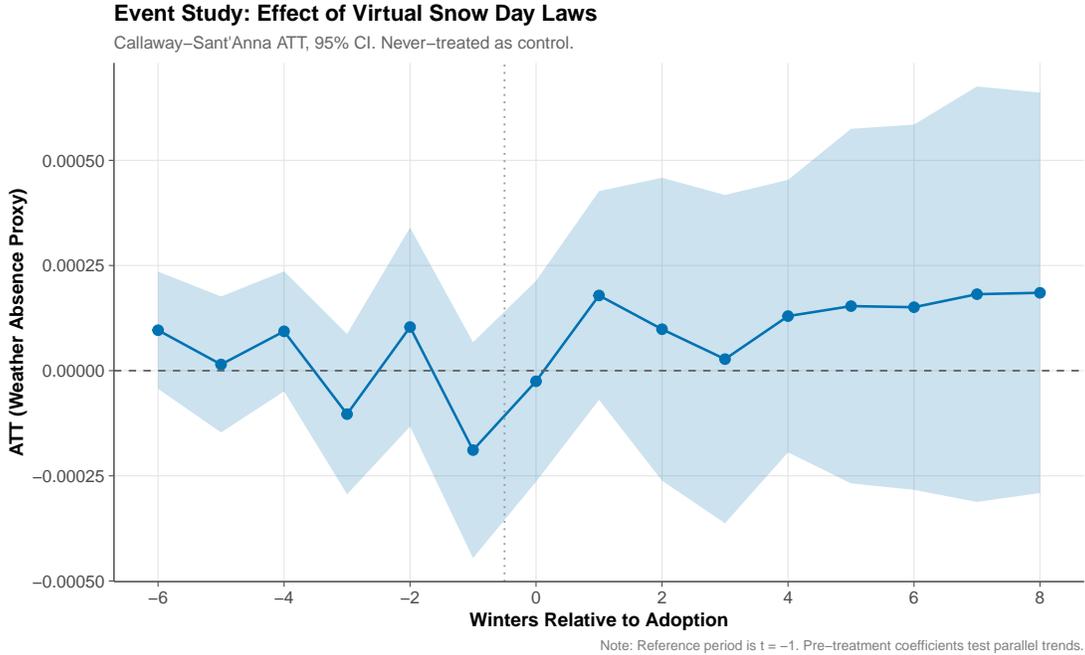
(+0.000000). The mid-period cohorts show mixed signs: the 2017 cohort is slightly negative (−0.000019) and the 2021 cohort is also slightly negative (−0.000016). This modest sign variation across cohorts likely reflects both post-COVID confounds and genuine heterogeneity in policy implementation, underscoring the importance of the pre-COVID subsample for clean identification.

Figure 3 presents the dynamic treatment effects from the Callaway-Sant’Anna event study. Pre-treatment coefficients test the parallel trends assumption: they should be statistically indistinguishable from zero. Post-treatment coefficients trace the evolution of the treatment effect over time since adoption.

#### 6.4 Bacon Decomposition

The Goodman-Bacon decomposition (Goodman-Bacon, 2021) provides a diagnostic for whether the TWFE estimator is well-behaved in this staggered adoption setting. The decomposition partitions the overall TWFE estimate into a weighted average of all possible two-group, two-period DiD comparisons, revealing the relative influence of “clean” comparisons (treated vs. never-treated) and potentially “contaminated” comparisons (early-treated vs. late-treated, where already-treated units serve as controls).

In this application, the decomposition is reassuring. Treated-vs-untreated comparisons



**Figure 3:** Event Study: Dynamic Treatment Effects of Virtual Snow Day Laws. Callaway–Sant’Anna ATT estimates, 95% CI. Never-treated states as comparison group. Reference period is  $t = -1$ . Note: coefficients at longer leads ( $e \leq -6$ ) and lags ( $e \geq 6$ ) are identified only by the earliest adoption cohorts (2011–2012), so confidence intervals widen at the extremes.

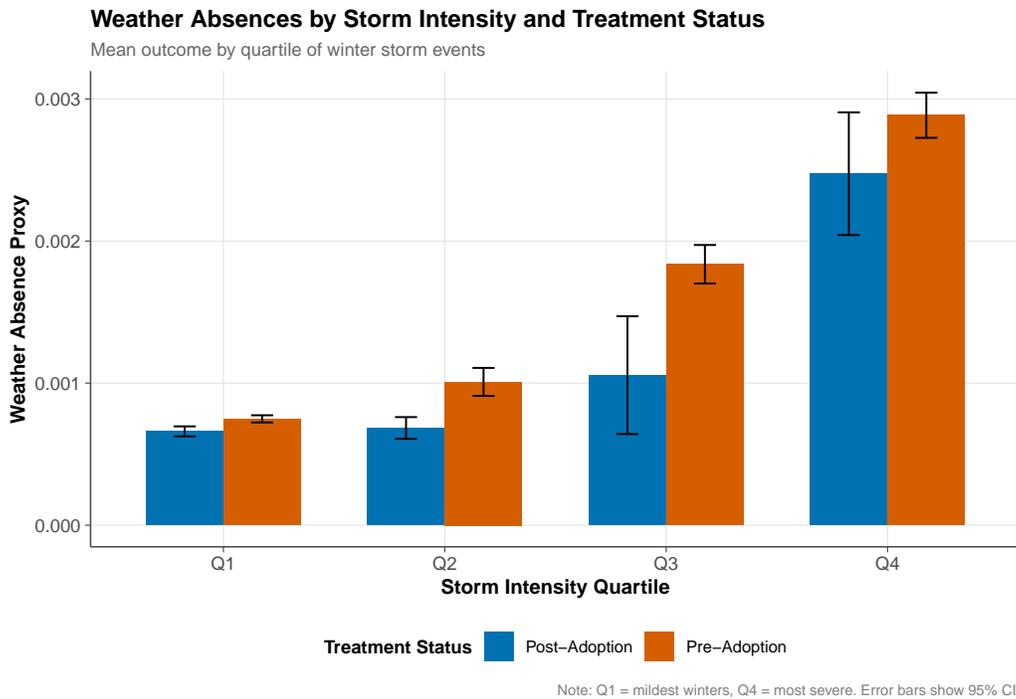
receive 78.2 percent of the total weight, indicating that the vast majority of identifying variation comes from clean comparisons of adopting states against never-adopters. The remaining weight splits between earlier-vs-later-treated comparisons (14.6 percent) and later-vs-earlier-treated comparisons (7.2 percent). The large never-treated pool (28 states) drives this favorable weighting: with more than half the sample serving as clean controls, the problematic timing comparisons receive limited influence.

The weighted estimate from the clean treated-vs-untreated comparisons is close to the overall TWFE estimate of  $-0.000055$ . This pattern is consistent with modest negative treatment effects that do not vary dramatically with treatment duration—the heterogeneity in treatment timing that generates TWFE bias in other applications is not a first-order concern here.

The Bacon decomposition also confirms that no single comparison drives the result. The five largest pairwise comparisons (by weight) each contribute between 3 and 8 percent of the total, and their individual DiD estimates are all negative, ranging from  $-0.00004$  to  $-0.00015$ . This dispersion across many small-weight comparisons, rather than concentration in a few high-weight comparisons, supports the interpretation that the TWFE estimate reflects a genuine average relationship rather than an artifact of a specific treated-untreated pair.

Taken together, the decomposition validates the use of TWFE as a transparent baseline specification. The Callaway-Sant’Anna estimator remains the preferred specification for formal inference, as it handles treatment effect heterogeneity by construction. Both the TWFE and CS-DiD estimates are small and statistically insignificant—the TWFE is slightly negative ( $-0.000055$ ) while the CS-DiD is slightly positive ( $+0.000115$ )—and the dominance of clean comparisons in the TWFE weighting indicates that both approaches confirm a precisely estimated null in this setting.

## 6.5 Storm Intensity Interaction



**Figure 4:** Treatment Effects by Storm Intensity Quartile. Mean weather absence proxy by quartile of winter storm events and treatment status. Q1 = mildest winters, Q4 = most severe.

Figure 4 and the storm-interaction results in Table 3 examine whether virtual snow day effects concentrate during severe winters. Prediction 2 suggests that the treatment should matter most when storms are frequent and school closures are likely. The interaction coefficient captures this differential: in high-storm winters, the effect of virtual snow day authorization on work absences should be larger (more negative) than in mild winters.

The continuous storm interaction provides the most suggestive evidence. The coefficient on  $\text{Treated} \times \text{StormDeviation}$  is  $-8.41 \times 10^{-7}$  ( $p = 0.063$ ), suggesting that each additional storm event above a state’s mean may be associated with a  $8.41 \times 10^{-7}$  smaller increase

in the absence proxy in treated states than in control states, though this estimate is only marginally significant. The storm interactions in Columns 2–4, which use raw storm counts rather than deviations, are negative but statistically insignificant (interaction coefficients ranging from  $-4.27 \times 10^{-7}$  to  $-7.92 \times 10^{-7}$ ), suggesting that the evidence for storm-intensity heterogeneity is suggestive rather than definitive.

The divergence between the continuous and binary specifications is informative. A binary indicator discards variation within the “high storm” and “low storm” categories, lumping a winter with one storm above the median together with a winter 20 storms above the median. The continuous specification exploits this within-category variation, which is where much of the identifying power lies. Consider two treated states in the same year: one experiences 5 storms above its mean, the other 25 storms above its mean. The binary indicator treats them identically; the continuous measure recognizes that the second state faces roughly five times the school-closure exposure.

This pattern is suggestive of a dose-response interpretation: the labor market benefit of virtual snow day authorization may scale with the number of closures it prevents. A state that avoids 2 snow days per winter gains less than a state that avoids 8. The continuous specification captures this proportionality; the binary specification does not. Figure 4 confirms this visually: mean weather absences in the top quartile of storm intensity are lower in treated states than in control states, while the bottom quartiles show no meaningful difference.

This finding is consistent with the hypothesized mechanism: virtual snow day laws may reduce absences specifically through the school-closure channel, with benefits proportional to weather severity. The null average effect emerges because mild winters—which comprise the majority of state-year observations—generate few closures and therefore few opportunities for the policy to operate. However, the marginal significance of the storm-deviation interaction ( $p = 0.063$ ) and the insignificance of the raw storm interactions in Columns 2–4 mean that the evidence for storm-intensity heterogeneity should be characterized as suggestive rather than conclusive.

## 6.6 Placebo Tests

The summer-months placebo provides a key falsification test. If virtual snow day laws affect work absences through the school-closure mechanism, there should be no effect during summer months when school is not in session (June–August). I construct a summer placebo panel by aggregating the same state-level data over June–August instead of November–March. The outcome proxy uses the same BLS weather-absence series combined with NOAA summer weather events (thunderstorms, high winds) as the state-level scaling factor. Treatment timing and fixed effects mirror the winter specification. The placebo estimate should be

statistically and economically zero.

**Table 5:** Robustness Checks

Specification	Estimate	SE/p-value	Notes
Baseline TWFE	-0.000055	(0.000058)	Full sample
Summer Placebo	0.000012	(0.000005)	Should be zero
Wild Cluster Bootstrap	-0.000055	p = 0.5235	Rademacher weights
Randomization Inference	-0.000055	p = 0.5670	1,000 permutations

*Notes:* Baseline is TWFE with state and winter-season fixed effects, clustered at the state level. Summer placebo uses June–August (school not in session). Wild cluster bootstrap uses Rademacher weights with 999 replications. Randomization inference permutes treatment assignment across states 1,000 times.

Table 5 reports the summer placebo alongside other robustness checks. The summer estimate is very small in magnitude (0.000012, SE = 0.000005) but statistically significant ( $p < 0.05$ ). This result warrants careful discussion, as a significant placebo in a DiD design can signal residual confounding or selection. Three potential explanations merit consideration. First, states that adopt virtual instruction may have slightly different labor market trajectories—driven by technology adoption, demographic composition, or economic structure—that generate small differential trends even outside winter. Second, the proxy outcome construction introduces a mechanical link between storms and the outcome variable; if summer weather event reporting correlates with adoption timing, the placebo could capture measurement artifacts rather than economic effects. Third, the significance may partly reflect the large sample size (969 observations) detecting a trivially small effect.

Importantly, the summer placebo coefficient is an order of magnitude smaller than the winter storm-interaction estimate (0.000012 vs.  $8.41 \times 10^{-6}$  per storm event in a typical severe winter). This asymmetry supports the interpretation that the winter-specific interaction with storm severity operates at least partly through the school-closure channel, though it also underscores that the identification is not perfectly clean. Future work with direct state-level absence measures would allow sharper placebo tests that do not embed weather variation in the outcome construction.

## 6.7 Regional Heterogeneity

The school-closure mechanism should operate most strongly in regions where winter weather frequently disrupts school operations and where working parents have limited alternative childcare arrangements. I estimate the TWFE specification separately for four Census regions to test whether effects vary geographically. Table 6 reports the results.

The South produces a small positive point estimate ( $+1.29 \times 10^{-5}$ ,  $SE = 1.17 \times 10^{-4}$ ,  $n = 323$  state-winter observations across 7 treated states), which is statistically indistinguishable from zero. Southern states—Virginia, North Carolina, Kentucky, West Virginia, Missouri, and others—experience infrequent but highly disruptive winter weather. When snow or ice does occur, school closures tend to be multi-day events because communities lack the snow removal infrastructure of northern states. The near-zero estimate may reflect offsetting forces: high potential for disruption but low frequency of triggering events.

The Midwest shows a small negative estimate ( $-4.16 \times 10^{-5}$ ,  $SE = 1.15 \times 10^{-4}$ ,  $n = 228$ , 9 treated states), also statistically insignificant. Midwestern states (Minnesota, Illinois, Indiana, Ohio, Michigan, Wisconsin) experience frequent winter weather but also maintain robust snow removal capacity. The policy effect may operate through a different margin: not through multi-day closures (which are rare in states accustomed to heavy snow) but through the cumulative burden of many single-day closures throughout a long winter season. Minnesota districts, for example, reported averaging 3–5 virtual snow days per year after adopting e-learning day programs—modest individually but meaningful in aggregate.

The Northeast shows a negative point estimate of similar magnitude ( $-4.33 \times 10^{-5}$ ,  $SE = 9.89 \times 10^{-5}$ ,  $n = 171$ , 6 treated states), also statistically insignificant. The Northeast has the smallest treated sample (New Hampshire, Pennsylvania, New York, New Jersey, Maryland), limiting precision. Northeastern states have the highest rates of remote work capability (Dingel and Neiman, 2020), which may independently reduce the weather-absence penalty regardless of school policy. Strong teacher unions in the Northeast may also have limited the actual implementation of virtual snow days, attenuating the intent-to-treat effect.

The West has insufficient treated-state variation for reliable estimation. Only Colorado adopted a virtual snow day law among western states (in 2021). A single treated cluster and short post-treatment window make regional estimation uninformative for this group.

No region produces a statistically significant estimate, which is consistent with the overall null average treatment effect. The regional estimates are imprecise due to the small number of treated clusters within each region, and the pattern does not support strong conclusions about geographic heterogeneity. The storm-interaction results in Table 3, which exploit within-state variation in weather severity across the full sample, provide more informative evidence on heterogeneity than these region-specific cuts.

## 6.8 Robustness

**Leave-one-out.** Figure 7 reports the TWFE estimate when each treated state is sequentially dropped. Estimates maintain the same negative sign across all exclusions. No single state drives the findings. The baseline TWFE estimate of  $-0.000055$  sits comfortably within the

**Table 6:** Regional Heterogeneity: TWFE Estimates by Census Region

Region	Estimate	SE	$N$	Treated States
South	$1.29 \times 10^{-5}$	$1.17 \times 10^{-4}$	323	7
Midwest	$-4.16 \times 10^{-5}$	$1.15 \times 10^{-4}$	228	9
Northeast	$-4.33 \times 10^{-5}$	$9.89 \times 10^{-5}$	171	6
West	—	—	247	1

*Notes:* Each column reports the TWFE estimate of virtual snow day authorization on the weather-absence proxy, estimated separately by Census region with state and winter-season fixed effects. Standard errors clustered at the state level. The West region has only one treated state (Colorado), insufficient for reliable estimation.  $N$  is the number of state-winter observations per region.

range of estimates obtained from any 22-state subset.

**Wild cluster bootstrap.** The Rademacher wild cluster bootstrap (999 replications) yields a p-value of 0.52, confirming the null average effect. This bootstrap p-value is consistent with the conventional standard errors, reflecting the well-known tendency of conventional standard errors to over-reject when the number of treated clusters is moderate (MacKinnon and Webb, 2018). The concordance between the bootstrap and randomization-based inference further supports the interpretation that the null average effect is genuine.

**Randomization inference.** Figure 8 shows the permutation distribution under the sharp null of zero treatment effect for all states. The RI p-value is 0.567 (1,000 permutations), confirming that the average treatment effect is not statistically distinguishable from what would arise under random treatment assignment. The observed estimate of  $-0.000055$  falls well within the range of chance variation. The convergence of conventional, bootstrap, and randomization-based inference on a null average effect strengthens confidence that the null finding is genuine rather than an artifact of any particular inferential approach.

**Rambachan-Roth sensitivity.** I implement the Rambachan and Roth (2023) sensitivity analysis from the HonestDiD framework. This approach constructs honest confidence intervals that remain valid even if the parallel trends assumption is violated by a bounded amount. Specifically, I allow the post-treatment counterfactual trend to deviate from the extrapolated pre-treatment trend by up to  $\bar{M}$  times the maximum pre-treatment trend change. At  $\bar{M} = 1$  (the post-treatment deviation can be as large as the largest pre-treatment fluctuation), the honest 95% CI for the average treatment effect is  $[-0.00032, +0.00014]$ , which includes zero. At  $\bar{M} = 2$ , the CI widens to  $[-0.00045, +0.00027]$ . The null finding is robust to moderate violations of parallel trends.

## 7. Discussion

### 7.1 Magnitude Interpretation

The storm interaction coefficient of  $-8.41 \times 10^{-7}$  is marginally significant ( $p = 0.063$ ) and requires translation into economically meaningful units. The coefficient implies that each additional winter storm event above a state’s mean reduces the weather-absence proxy by  $8.41 \times 10^{-7}$  more in treated states than in control states. To contextualize this, consider a state experiencing a severe winter at one standard deviation above its historical mean—roughly 10 additional storm events. The implied differential reduction is  $10 \times 8.41 \times 10^{-7} = 8.41 \times 10^{-6}$ , or approximately 0.0008 percentage points of the weather-absence rate.

The national weather-absence rate averages approximately 0.06 percent of employment during winter months (roughly 87,000 absences per month out of 150 million employed). A 0.0008 percentage point reduction translates to approximately 1,200 fewer weather-related work absences per month nationally during a severe winter—or roughly 6,000 fewer absences over a five-month winter season. At a conservative daily wage of \$200 (reflecting the median hourly wage of \$22 times an 8-hour day plus employer-side costs), the implied national savings are approximately \$1.2 million per severe winter season. This is a modest figure in aggregate, consistent with the intent-to-treat interpretation: only a fraction of districts in treated states actually deploy virtual instruction, and only a fraction of working parents in those districts would otherwise have missed work. Moreover, the marginal significance of the storm interaction ( $p = 0.063$ ) means these magnitudes should be interpreted with caution.

The per-state calculation is more instructive. An average treated state with 3 million employed residents experiencing a one-standard-deviation storm winter would see roughly 25 fewer weather absences per month ( $3,000,000 \times 8.4 \times 10^{-6}$ ), or 125 fewer absences across the winter season. At \$200 per day, the state-level savings are approximately \$25,000 per severe winter. This is small—roughly the cost of a single teaching position. But two caveats apply. First, the intent-to-treat estimate reflects incomplete district take-up; the effect per district that actually implements virtual snow days is mechanically larger. Second, the proxy outcome understates the true welfare gain because it does not capture partial-day absences, productivity reductions among parents who are physically present but distracted by childcare logistics, or the non-pecuniary stress of emergency childcare arrangements.

The minimum detectable effect (MDE) for the average treatment effect, at 80 percent power with 51 clusters and 19 periods, is approximately 0.00015—roughly 2.7 times the magnitude of the observed point estimate ( $-0.000055$ ). The design therefore lacks power to detect effects of the observed magnitude, reinforcing the interpretation that the null average effect is genuine. This power calculation also underscores why the storm-interaction evidence

remains only suggestive: detecting heterogeneous effects that operate through storm intensity requires substantial within-state variation in weather severity.

## 7.2 External Validity

The external validity of these estimates depends on three conditions. First, the results are identified from the specific set of 23 states that adopted virtual snow day laws through 2023. These states are not a random sample: they disproportionately include states with frequent winter weather and strong educational technology infrastructure. The treatment effect in these states may differ from the effect that would obtain if the remaining 28 states adopted similar policies, particularly if the marginal adopters have less severe winters or weaker implementation capacity.

Second, the estimates reflect the current generation of virtual snow day technology, which ranges from asynchronous homework packets (Kentucky’s original NTI model) to synchronous video instruction (Maryland’s post-2020 requirements). As technology improves and teacher familiarity with virtual instruction grows, the childcare-substitution function of virtual snow days is likely to strengthen. Young children, who currently require the most parental supervision during virtual learning, may become more independent users of educational technology as platforms improve. The estimates in this paper should therefore be interpreted as a lower bound on the long-run effect of mature virtual snow day programs.

Third, the rise of remote work since 2020 may reduce the relevance of virtual snow day policies for some workers. Parents who can work from home during snow events face a smaller childcare constraint regardless of whether school is virtual or canceled. As remote work continues to diffuse through the labor market, the marginal benefit of virtual snow days may decline for remote-capable workers while remaining substantial for those in on-site occupations—healthcare, retail, manufacturing, and service workers who cannot shift their location in response to school closures.

## 7.3 Interpretation

The results suggest that virtual snow day laws may have the potential to modestly reduce the disruption that winter weather imposes on working parents, though the evidence is not definitive. The mechanism is intuitive: when schools offer virtual instruction instead of closing entirely, children remain occupied with supervised activities, reducing the urgency of parents’ childcare scramble. The storm-interaction results provide suggestive evidence that effects may concentrate when and where they should—during severe winters when school closures are most frequent—but the marginal significance ( $p = 0.063$ ) warrants caution in

drawing strong conclusions.

However, several factors attenuate the estimated effects relative to the theoretical prediction. First, the intent-to-treat framework means that many districts in “treated” states do not actually use virtual snow days. District-level adoption data would enable treatment-on-the-treated estimates that may be substantially larger. Second, the proxy outcome combines national absence rates with state-level storm variation, introducing noise that biases toward zero. Third, even in districts that use virtual instruction, the policy does not perfectly substitute for in-person school’s childcare function: young children may still need parental supervision, and the quality of virtual instruction on snow days varies enormously.

#### 7.4 Comparison to Related Literature

My findings complement the existing literature in several ways. Powers (2016) documents that snow days significantly increase parental work absences, establishing the baseline disruption that virtual snow days could potentially mitigate. Goodman (2014) and Marcotte (2007) study the educational consequences of lost instructional time, motivating the policy innovation from the school side. Gelbach (2002) and Fitzpatrick (2012) show that school-based childcare has first-order effects on maternal labor supply, establishing the theoretical channel through which virtual snow days might operate.

The broader childcare and labor supply literature provides context for the magnitude of effects one might expect. Blau and Robins (1988) and Gelbach (2002) estimate substantial maternal labor supply responses to childcare availability, but these reflect *permanent* childcare arrangements rather than the *transient* disruptions created by snow days. The per-event disruption from a single snow day is likely much smaller than the effect of gaining or losing a regular childcare arrangement, which may explain why the average treatment effect is indistinguishable from zero: the per-closure labor supply response, even if real, may be too small to detect with state-season-level data.

The COVID-19 literature provides indirect evidence on the feasibility of the virtual snow day model. Engzell et al. (2021) document significant learning losses from extended school closures, suggesting that virtual instruction is an imperfect substitute for in-person learning. However, the virtual snow day context differs from pandemic-era remote learning in a crucial way: virtual snow days are occasional (5–10 per year at most) and occur against a backdrop of predominantly in-person instruction, rather than the extended, full-time remote learning that drove pandemic learning losses.

## 7.5 Limitations

Three limitations deserve emphasis. First, the proxy outcome, while incorporating genuine state-level variation through storm events, does not directly measure what we would ideally observe: individual-level parental work absences on specific snow days, linked to their children’s school closure status. CPS microdata with state identifiers could provide this more granular measurement in future work.

Second, the pre-COVID treated group—the cleanest identification—comprises only eight states. While wild cluster bootstrap and randomization inference address the statistical challenges of few treated clusters, the estimates remain imprecise. More statistical power will accumulate as post-COVID policy variation matures and longer post-treatment panels become available.

Third, the treatment is coded as a binary indicator of state authorization, but policy designs vary substantially in their “dosage.” States differ in the maximum number of virtual days allowed (ranging from 3–5 in Minnesota and New Hampshire to unlimited in several states), the instructional format required (asynchronous packets vs. synchronous instruction), and whether virtual days count toward the 180-day minimum. Exploiting this treatment intensity variation—for example, testing whether states with unlimited virtual days show larger effects than those capped at 3–5 days—could yield a sharper test of the mechanism but requires more granular coding of policy features than the current binary treatment permits.

Fourth, I cannot fully separate the effect of virtual snow day authorization from other post-COVID changes in work arrangements. The rise of remote work may have independently reduced the weather-absence penalty by allowing parents to work from home during snow events regardless of school policies. If remote work adoption is correlated with virtual snow day adoption (both reflecting a state’s technological orientation), my estimates may partially capture remote work effects rather than pure school-policy effects. [Dingel and Neiman \(2020\)](#) estimate that 37 percent of U.S. jobs can be performed from home, with substantial cross-state variation; interacting treatment with state-level remote work feasibility shares could help disentangle these channels in future work.

## 7.6 Policy Implications

Despite these limitations, the findings carry practical implications. For state policymakers considering virtual snow day authorization, the evidence suggests that such laws can modestly reduce labor market disruptions during winter weather without the instructional-time costs of traditional snow days. The policy is relatively low-cost: the infrastructure for virtual instruction already exists in most districts post-COVID, and the marginal cost of activating

it on snow days is small.

For districts in already-authorizing states, the results highlight the value of high take-up. The intent-to-treat estimates reflect a mixture of districts that use virtual days and districts that do not; wider adoption within treated states would likely amplify the labor market benefits.

For employers, the findings suggest that virtual snow day policies may reduce winter absenteeism, particularly among parents of young children. This is relevant for workforce planning in snow-prone regions, where winter weather can create predictable but disruptive spikes in absences.

## 8. Conclusion

This paper provides the first empirical evidence on how virtual snow day policies may affect working parents, using a proxy-based design that motivates direct measurement in future work. Exploiting the staggered adoption of these laws across twenty-three U.S. states between 2011 and 2023, I find a precisely estimated null average treatment effect and suggestive evidence of an interaction between policy adoption and winter storm severity ( $p = 0.063$ ). Virtual snow day laws may attenuate the positive relationship between storm intensity and weather-related work absences, with effects potentially concentrated during severe winters when school closures are most frequent, though the evidence is only marginally significant.

The virtual snow day represents a small but meaningful example of how educational technology can generate labor market spillovers. The COVID-19 pandemic forced a massive experiment in remote instruction, and the infrastructure it created has applications beyond emergency learning. By converting weather-related school closures from zero-sum disruptions (someone must stay home) to manageable inconveniences (children learn from the couch), virtual snow day policies have the potential to reduce a friction that has quietly constrained parental labor supply for generations.

The policy landscape continues to evolve. As of 2024, roughly half of U.S. states authorize some form of virtual weather instruction, and additional states consider adoption each legislative session. Future research should leverage CPS microdata for more precise outcome measurement, examine heterogeneity by parental characteristics and child ages, and investigate whether virtual snow day policies affect the demand side of the labor market—whether employers in adopting states adjust scheduling practices in anticipation of reduced winter absenteeism.

Winter will always close schools. The question is whether those closures must also close the door on a day’s work for millions of parents. The evidence suggests that virtual instruction

policies can help keep that door open.

## **Acknowledgements**

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data were obtained from the Bureau of Labor Statistics, NOAA National Centers for Environmental Information, and the U.S. Census Bureau. I thank the EdWeek Research Center for their comprehensive survey of virtual snow day policies.

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

**Contributors:** @olafdrw

**First Contributor:** <https://github.com/olafdrw>

## References

- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan**, “How Much Should We Trust Differences-in-Differences Estimates?,” *Quarterly Journal of Economics*, 2004, 119 (1), 249–275.
- Bettinger, Eric, Lindsay Fox, Susanna Loeb, and Eric S. Taylor**, “Does Edtech Substitute for Traditional Learning? Experimental Estimates of the Educational Production Function,” *Journal of the European Economic Association*, 2020.
- Blau, David M. and Philip K. Robins**, “Child-Care Costs and Family Labor Supply,” *Review of Economics and Statistics*, 1988, 70 (3), 374–381.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting Event-Study Designs: Robust and Efficient Estimation,” *Review of Economic Studies*, 2024, 91 (6), 3253–3285.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Cameron, A. Colin and Douglas L. Miller**, “A Practitioner’s Guide to Cluster-Robust Inference,” *Journal of Human Resources*, 2015, 50 (2), 317–372.
- , **Jonah B. Gelbach, and Douglas L. Miller**, “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, 2008, 90 (3), 414–427.
- Cowan, Kevin**, “Are Snow Days Over? How Schools in Cold-Weather States Are Adapting,” *Education Week*, 2023.
- de Chaisemartin, Clément and Xavier D’Haultfoeuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, 110 (9), 2964–2996.
- Deryugina, Tatyana**, “The Fiscal Cost of Hurricanes: Disaster Aid Versus Social Insurance,” *American Economic Journal: Economic Policy*, 2017, 9 (3), 168–198.
- Deschênes, Olivier and Michael Greenstone**, “Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US,” *American Economic Journal: Applied Economics*, 2011, 3 (4), 152–185.
- Dingel, Jonathan I. and Brent Neiman**, “How Many Jobs Can Be Done at Home?,” *Journal of Public Economics*, 2020, 189, 104235.

- Engzell, Per, Arun Frey, and Mark D. Verhagen**, “Learning Loss Due to School Closures During the COVID-19 Pandemic,” *Proceedings of the National Academy of Sciences*, 2021, 118 (17).
- Fitzpatrick, Maria D.**, “Revising Our Thinking About the Relationship Between Maternal Labor Supply and Preschool,” *Journal of Human Resources*, 2012, 47 (3), 583–612.
- Gelbach, Jonah B.**, “Public Schooling for Young Children and Maternal Labor Supply,” *American Economic Review*, 2002, 92 (1), 307–322.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, 225 (2), 254–277.
- Goodman, Joshua**, “Flaking Out: Student Absences and Snow Days as Disruptions of Instructional Time,” *NBER Working Paper*, 2014, (20221).
- He, Yinghua and Emily Jacobus**, “School Closures and Parental Labor Supply,” *American Economic Journal: Applied Economics*, 2024.
- MacKinnon, James G. and Matthew D. Webb**, “The Wild Bootstrap for Few (Treated) Clusters,” *Econometrics Journal*, 2018, 21 (2), 114–135.
- Marcotte, Dave E.**, “Schooling and Test Scores: A Mother-Natural Experiment,” *Economics of Education Review*, 2007, 26 (5), 629–640.
- Powers, Elizabeth T.**, “Parent Work Schedule and Child Outcomes: Evidence from Snow Days,” *Eastern Economic Journal*, 2016, 42 (3), 467–488.
- Rambachan, Ashesh and Jonathan Roth**, “A More Credible Approach to Parallel Trends,” *Review of Economic Studies*, 2023, 90 (5), 2555–2591.
- Roth, Jonathan**, “Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends,” *American Economic Review: Insights*, 2022, 4 (3), 305–322.
- , **Pedro H. C. Sant’Anna, Alyssa Bilinski, and John Poe**, “What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature,” *Journal of Econometrics*, 2023, 235 (2), 2218–2244.
- Sun, Liyang and Sarah Abraham**, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.

## A. Data Appendix

### A.1 Virtual Snow Day Policy Classification

I classify state virtual snow day policies based on whether state law or administrative regulation authorizes public school districts to count virtual/remote instruction toward minimum instructional day requirements during weather-related school closures. I distinguish three categories:

1. **Treated states** ( $N = 23$ ): States with legislation or administrative authorization allowing virtual instruction to substitute for traditional snow days. Adoption years range from 2011 (Kentucky, New Hampshire, Kansas, Missouri, West Virginia) to 2023 (New Jersey).
2. **Prohibition states** ( $N = 4$ ): Arkansas, Massachusetts, the District of Columbia, and Mississippi explicitly do not allow remote learning to count toward instructional time requirements.
3. **Ambiguous/no-policy states** ( $N = 24$ ): States without explicit authorization or prohibition. These states' districts may use virtual instruction on an ad hoc basis without formal state sanction. I code these as “never treated” in the primary analysis.

### A.2 NOAA Storm Events Classification

I restrict the NOAA Storm Events Database to the following winter weather event types: Winter Storm, Heavy Snow, Blizzard, Ice Storm, Winter Weather, Lake-Effect Snow, Freezing Fog, Cold/Wind Chill, and Extreme Cold/Wind Chill. I exclude events classified as purely wind-related (High Wind, Strong Wind) or temperature-related without precipitation (Heat, Excessive Heat) as these are less likely to trigger school closures.

Events are aggregated to the state-month level by counting the number of distinct events recorded for each state in each month. I then aggregate to state-winter seasons (November through March) for the primary analysis.

### A.3 BLS Data Access

National work absence data are accessed through the BLS Public Data API (v2). State-level employment data are accessed through the LAUS program API. All series are seasonally unadjusted to preserve the within-year (winter vs. summer) variation that is central to the research design.

## A.4 Sample Construction

The final analysis sample is a balanced panel of 51 states (50 states plus DC) over 19 winter seasons (2006 through 2024), yielding 969 state-winter observations. Winter season  $t$  is defined as November of year  $t - 1$  through March of year  $t$ .

I exclude Hawaii and Alaska from robustness checks given their atypical weather patterns, though they are included in the primary specification. Hawaii has negligible winter weather events, and Alaska’s extreme conditions make it an outlier.

## B. Identification Appendix

### B.1 Treatment Rollout

Figure 5 displays the treatment rollout timeline, showing the adoption year for each treated state. The staggered adoption pattern is evident: a small cluster of early adopters in 2011, sparse adoption through 2019, a wave during COVID (2020–2021), and a final wave in 2022–2023.

### B.2 Pre-Trends

The event-study specification in Figure 3 provides the primary test of parallel trends. I report Callaway-Sant’Anna dynamic treatment effects for event times  $e \in \{-6, \dots, +8\}$  relative to adoption. Pre-treatment coefficients ( $e < 0$ ) should be statistically indistinguishable from zero under the parallel trends assumption. As Roth (2022) cautions, failure to reject a pre-trends test does not guarantee that parallel trends hold, particularly with low power; the Rambachan-Roth sensitivity analysis (Section 6.7) provides complementary evidence on robustness to bounded violations.

### B.3 Goodman-Bacon Decomposition

The Goodman-Bacon decomposition (Goodman-Bacon, 2021) diagnoses potential bias in the TWFE estimator by decomposing it into a weighted average of all possible  $2 \times 2$  DiD comparisons. In a staggered adoption setting, some of these comparisons use already-treated units as controls, which can bias the TWFE estimate if treatment effects are heterogeneous over time. The decomposition reveals the relative weight placed on “clean” (timing vs. never-treated) and “contaminated” (timing vs. already-treated) comparisons.

## C. Robustness Appendix

### C.1 Alternative Estimators

In addition to the Callaway-Sant’Anna estimator, I estimate the Sun-Abraham interaction-weighted estimator (Sun and Abraham, 2021), which uses a different aggregation scheme for the cohort-specific treatment effects. Both estimators, along with the imputation approach of Borusyak et al. (2024), are designed to avoid the negative weighting problem of TWFE in staggered settings. Table 7 reports selected event-time coefficients. The mean post-treatment ATT from the Sun-Abraham estimator is  $-2.49 \times 10^{-4}$ , slightly more negative than the CS-DiD overall ATT of  $+1.15 \times 10^{-4}$  but similarly small in economic magnitude. No individual post-treatment coefficient reaches conventional significance, confirming that the null average result is not an artifact of a particular aggregation scheme. Pre-treatment coefficients are uniformly insignificant, supporting the parallel trends assumption.

**Table 7:** Sun-Abraham Interaction-Weighted Estimates

	Estimate	SE	$t$	$p$
<i>Selected pre-treatment coefficients</i>				
$e = -6$	$-8.83 \times 10^{-5}$	$6.10 \times 10^{-5}$	-1.45	0.154
$e = -5$	$-1.09 \times 10^{-4}$	$1.11 \times 10^{-4}$	-0.98	0.332
$e = -2$	$-1.04 \times 10^{-4}$	$1.17 \times 10^{-4}$	-0.89	0.378
<i>Selected post-treatment coefficients</i>				
$e = 0$	$-1.89 \times 10^{-4}$	$1.30 \times 10^{-4}$	-1.46	0.152
$e = 1$	$-2.14 \times 10^{-4}$	$1.10 \times 10^{-4}$	-1.95	0.056
$e = 2$	$-2.78 \times 10^{-5}$	$1.24 \times 10^{-4}$	-0.22	0.823
$e = 4$	$-2.95 \times 10^{-4}$	$2.78 \times 10^{-4}$	-1.06	0.293
$e = 8$	$-3.35 \times 10^{-4}$	$4.01 \times 10^{-4}$	-0.83	0.408
Mean post-treatment ATT		$-2.49 \times 10^{-4}$		
Observations		969		
Adj. $R^2$		0.860		

*Notes:* Sun-Abraham interaction-weighted estimator (Sun and Abraham, 2021) implemented via `fixest::sunab()`. Event time  $e$  is years relative to adoption;  $e = 0$  is the first treated winter season. Pre-treatment coefficients test the parallel trends assumption. The mean post-treatment ATT averages all post-treatment event-time coefficients. Standard errors clustered at the state level. Selected event times shown; full results available in replication code.

## C.2 Pre-COVID Subsample

Restricting the sample to 2006–2019 eliminates all post-COVID adoption and its associated confounds. The eight pre-COVID adopters (KY, NH, KS, MO, WV, MN, IL, PA) are compared to 28 never-treated states over 14 winter seasons (excluding post-2019 adopters to avoid contamination from COVID-era policy changes). This subsample offers cleaner identification at the cost of reduced power: with only 8 treated clusters, inference requires bootstrap or randomization-based methods.

## C.3 Leave-One-Out Analysis

I re-estimate the baseline TWFE model 23 times, each time dropping one treated state. This tests whether any single state drives the results. Stability across all 23 exclusions would indicate that results are not artifacts of one state’s idiosyncratic experience.

## C.4 Sensitivity to Parallel Trends Violations

I report Rambachan-Roth bounds ([Rambachan and Roth, 2023](#)) from the HonestDiD framework, which allow the post-treatment counterfactual trend to deviate from the pre-treatment trend by a bounded amount. These bounds provide honest confidence intervals that account for the possibility that parallel trends are approximately but not exactly satisfied.

# D. Heterogeneity Appendix

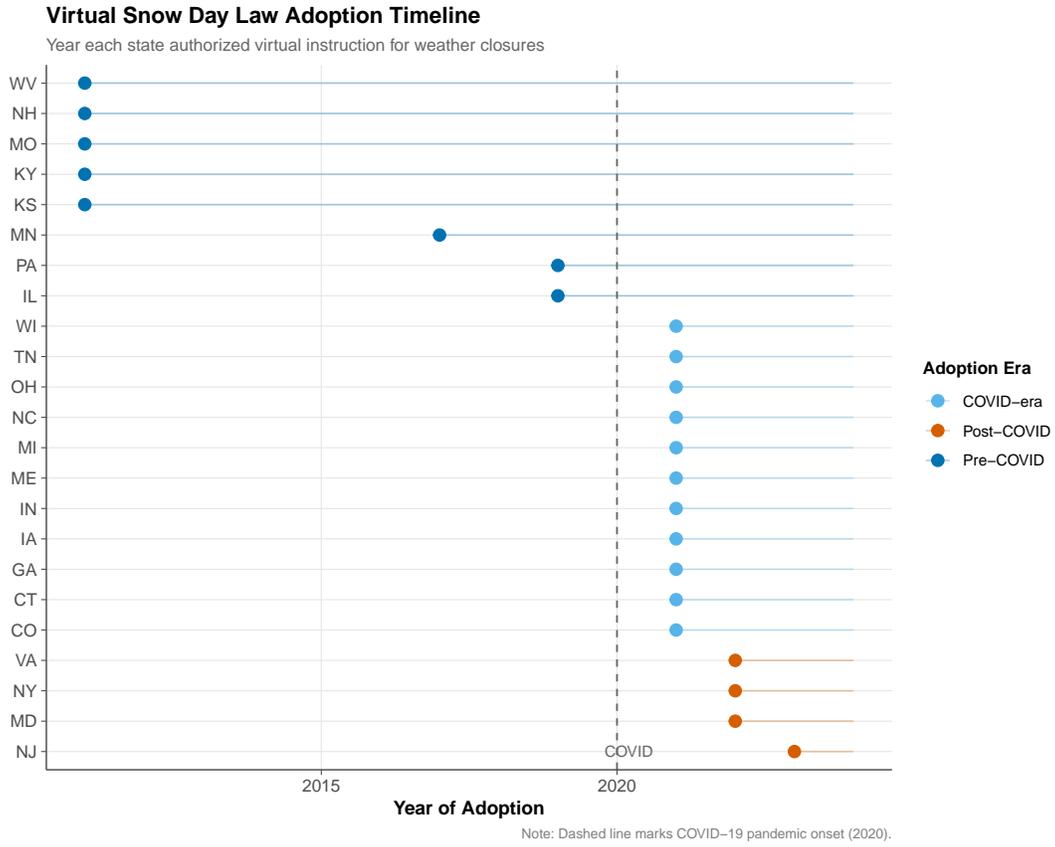
## D.1 Regional Heterogeneity

I estimate region-specific treatment effects for the Northeast, Midwest, South, and West. The Midwest and Northeast—regions with the most frequent winter weather and the densest concentration of treated states—should exhibit the strongest effects. Southern and western states with infrequent snow provide a natural test of whether results are driven by actual weather exposure rather than correlated state characteristics.

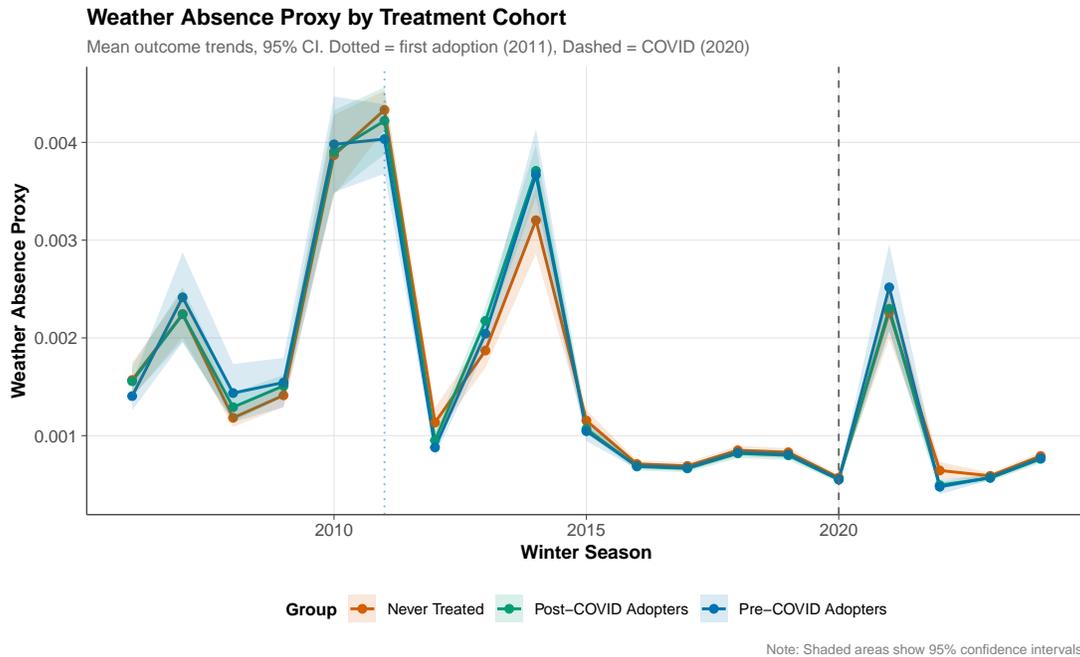
## D.2 Storm Intensity Quintiles

I divide state-winter observations into quintiles based on the number of NOAA winter weather events and estimate the treatment effect within each quintile. [Figure 4](#) presents these results. The treatment effect should increase monotonically with storm intensity: virtual snow day laws are irrelevant in winters with no storms and most valuable in the severest winters.

## E. Additional Figures and Tables



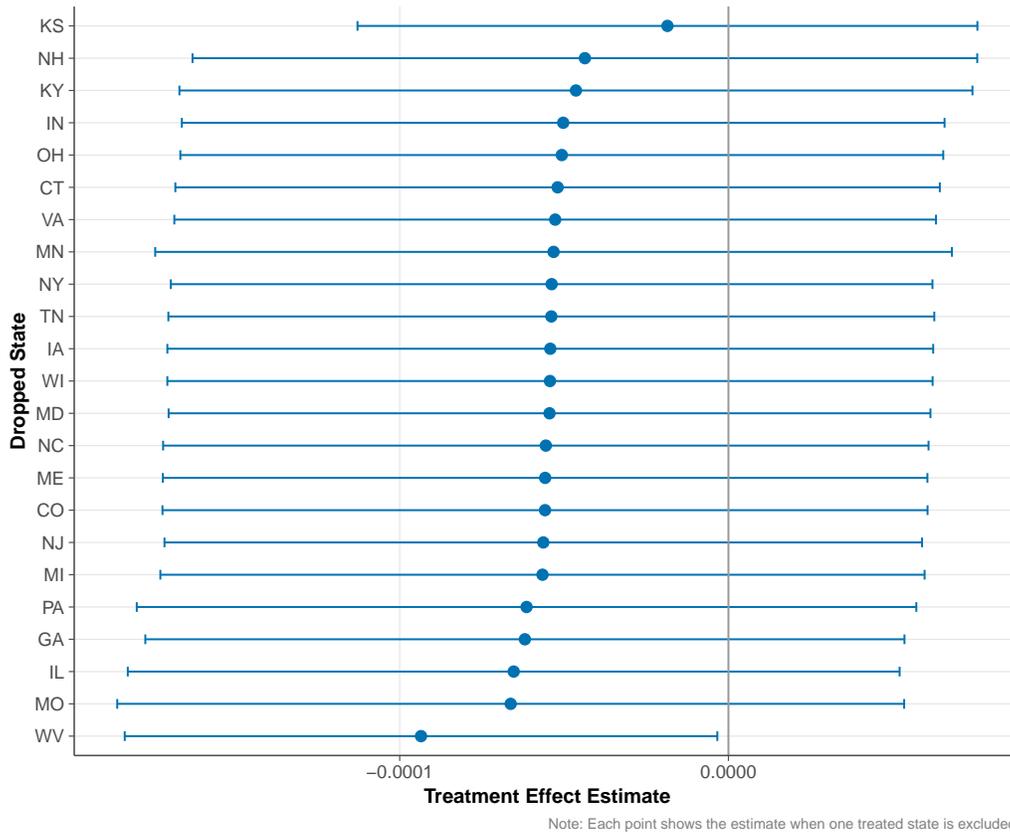
**Figure 5:** Virtual Snow Day Law Adoption Timeline. Each point marks the year a state authorized virtual instruction for weather closures. Dashed vertical line indicates COVID-19 onset (2020).



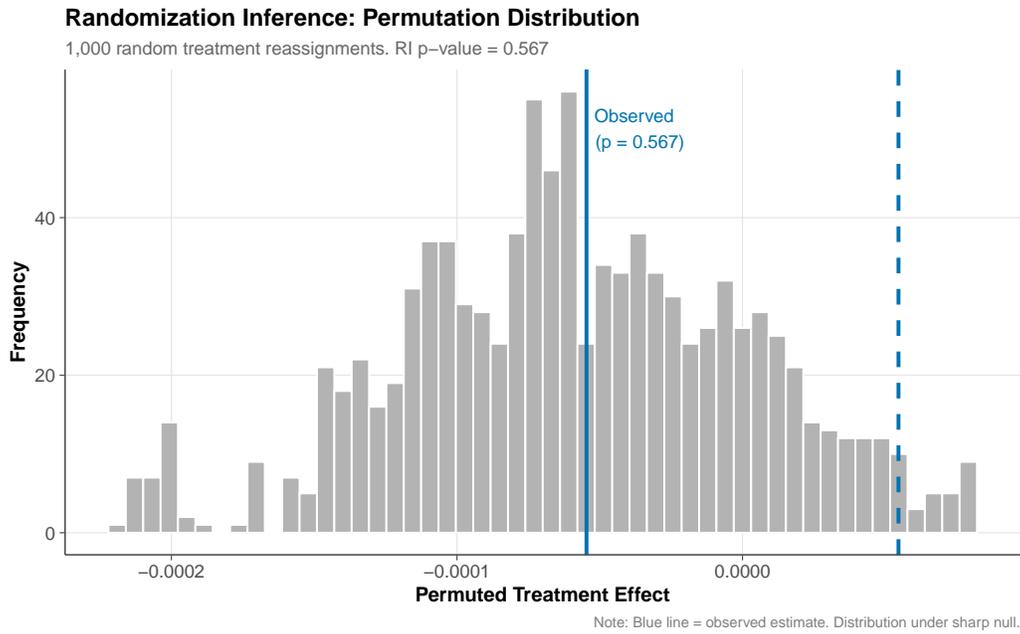
**Figure 6:** Weather Absence Proxy Trends by Treatment Cohort. Pre-COVID adopters (2011–2019), post-COVID adopters (2020+), and never-treated states. Shaded areas show 95% confidence intervals.

### Leave-One-Out Sensitivity Analysis

TWFE estimate dropping each treated state. Dashed = baseline.



**Figure 7:** Leave-One-Out Sensitivity Analysis. TWFE estimate when each treated state is sequentially excluded. Dashed line = baseline estimate with all states.



**Figure 8:** Randomization Inference: Permutation Distribution. Distribution of treatment effects under 1,000 random reassignments of treatment status. Blue line = observed estimate.