

The Impact of Minimum Wage Increases on Incomes, Employment, and Poverty: Monthly Evidence from the \$15 Era

Claude · August 2026

Experimental draft — not for citation or circulation

Claude is an AI system developed by Anthropic; this paper was written by the Claude Fable 5 model generation. This paper was commissioned by Jason Furman. Contact: Claude, c/o Jason Furman, Harvard Kennedy School, 79 John F. Kennedy Street, Cambridge, MA 02138. The views expressed herein are those of the author. JEL codes: I32, J23, J31, J38.

ABSTRACT

Between 2018 and 2024, 28 states holding 57 percent of the U.S. population raised their binding minimum wage from an average of \$9.57 to \$14.17, while 21 states left theirs effectively unchanged near the \$7.25 federal floor. This paper studies the wave in what is, to our knowledge, the first minimum wage evidence from the redesigned Survey of Income and Program Participation — 4.0 million person-months measuring employment, family income, poverty, and program receipt monthly — using staggered event studies, two-way fixed effects, and within-person designs. The floor binds: the 10th-percentile wage elasticity is about 0.25 (randomization $p = 0.02$), and it is unchanged when the survey’s seam months — the December-January boundaries where its reference years meet — are excised. Employment of low-education workers does not detectably fall; the implied own-wage elasticity of -0.04 (s.e. 0.15) bounds broad-group losses only. The raw monthly poverty effect is zero, and drift-adjusted paths lean negative but are fragile. Two mechanical findings follow. Aggregated to the year grain, the identical design on the identical person-months yields a poverty “effect” of about one percentage point that matches its own pre-treatment value — base-period anchoring on a dip-shaped path. And within-person job-loss estimates of 5.2 percentage points are reproduced 60 to 100 percent by placebos run where no increase occurred, though the placebo arms are thin and the netted intervals wide. Monthly measurement makes visible the drift that annual designs average into policy effects.

I. Introduction

Between January 2018 and December 2024, 28 states raised their binding minimum wage — the higher of the state and federal floors — by at least 50 cents in a single statutory step, taking their population-weighted average binding minimum from \$9.57 to \$14.17, a 48 percent increase.¹ Over the same seven years, 21 states — the no-change states, in the vocabulary used throughout this paper — left their binding minimum effectively unchanged, moving less than 25 cents over the entire window; the federal floor itself has been \$7.25 since 2009. In all, states enacted 173 statutory increases over the window, including phase-ins to \$15 or beyond in California, New York, New Jersey, Illinois, Maryland, and Florida. Measured by the number of states moving — the treated states hold 57 percent of the U.S. population — the size of the increases, and the share of the workforce exposed, this is the largest state-led minimum wage experiment in U.S. history.

The literature that should be able to say what this experiment did agrees on one result and disagrees about nearly everything downstream of it. Minimum wage increases raise wages at the bottom of the distribution; that finding survives across designs from Card and Krueger (1994) onward. On employment, bunching estimates covering 138 state increases find low-wage employment essentially unchanged (Cengiz, Dube, Lindner, and Zipperer 2019), while an assembly of the literature’s preferred estimates finds 79.2 percent of employment elasticities negative (Neumark and Shirley 2022). On poverty, event-study estimates put the long-run poverty elasticity between -0.22 and -0.46 (Dube 2019b), while dynamic difference-in-differences estimates on nearly four decades of March CPS data rule out elasticities more negative than -0.129 (Burkhauser, McNichols, and Sabia 2023) — the two ranges do not overlap. On transfers, one side finds that program phase-outs claw back roughly a third of gross income gains (Dube 2019b); the other (Clemens 2016) finds that transfers changed little after the 2008-2009 federal increases.

This paper measures the \$15-era wave in data none of these literatures has used: the redesigned Survey of Income and Program Participation (SIPP), which from 2018 through 2024 provides monthly employment status, monthly family income, monthly income-to-needs ratios, and monthly program amounts for 4,007,970 person-months. The redesign spans the entire wave, identifies states, and measures precisely the objects — income, employment, and poverty, at

¹ Throughout this paper the District of Columbia is counted as a state, so “51 states” describes the full panel and DC’s \$17.50 is the highest “state” minimum in December 2024.

monthly frequency — that the disputes above turn on. The prior SIPP footprint in this literature is a decade old, drawn from the classic panels, and largely concentrated on the recession-era federal episode (Clemens and Wither 2019; Clemens 2016; Sabia and Nielsen 2015).

The floor binds. After a state's first large step, its 10th-percentile hourly wage rises by 3.0 log points (s.e. 2.3) on the raw event-study path and 6.6 log points (s.e. 3.8) net of the pre-trend. Expressed as elasticities with respect to the minimum, the event-study estimates are 0.24 (s.e. 0.18) raw and 0.53 (s.e. 0.31) net of the pre-trend, and a two-way fixed effects regression of the 10th-percentile log wage on the log minimum gives 0.247 (s.e. 0.093) — the same point estimate as the raw event-study elasticity, with half the standard error. Both are reported throughout. A randomization test that reassigns states' entire minimum wage paths puts a p-value of 0.02 on the first stage, and the estimate is unchanged (0.245, s.e. 0.094) when every December and January month is excised from the sample, the direct check on the survey-seam concern Section V takes up. The 25th-percentile elasticity is 0.095 (s.e. 0.036), so the effect fades with distance from the floor, which is the signature of a binding minimum rather than a booming low-wage labor market.

Employment does not detectably fall. The employment rate of workers with a high school degree or less, ages 16 to 64, moves by +0.3 percentage points (s.e. 1.2) after treatment, and the two-way fixed effects employment elasticity is -0.009 (s.e. 0.038). Weekly hours among those employed decline slightly (elasticity -0.044, s.e. 0.024), so the intensive margin absorbs a little of the adjustment. Dividing the employment elasticity by the wage elasticity gives an implied own-wage employment elasticity of -0.04 (s.e. 0.15), with a 95 percent lower bound of -0.33; the same ratio built entirely from the event-study estimates is +0.19 (s.e. 0.70), the same answer with an interval too wide to be informative. The bound is a statement about the broad low-education group: it rules out group-level elasticities near -1 and is consistent with the near-zero bunching estimate (Cengiz et al. 2019) and with the -0.04 median for broad worker groups across the modern literature (Dube 2019a). It is not, by itself, a bound on losses among the workers the floor actually binds — only about 9 percent of the group's employed members are paid within a dollar of the minimum, and a bound diluted by that share cannot exclude affected-worker losses on the Clemens and Wither (2019) scale. The person-level design of Section VIII is the evidence on that question.

Poverty is where the evidence leans favorable without deciding. The raw monthly poverty effect is zero: +0.03 percentage points (s.e. 0.49). Net of a linear pre-trend, the path is negative, -1.38 percentage points (s.e. 1.01), but that adjusted estimate is sensitive to the window the trend is fit on and should not be leaned on. The share of people below 150 percent of needs — where minimum wage families actually sit — is negative in the raw estimate (-1.0 percentage point, s.e. 0.6) and in every adjusted variant, though the raw estimate is imprecise and carries drift of its own. The confidence intervals contain both Dube’s central range and Burkhauser, McNichols, and Sabia’s bound. On the magnitude, these data do not adjudicate.

What the data do show is why the annual literature cannot settle the question. Aggregate the identical person-months to the year grain and run the identical design, and the monthly zero becomes a poverty increase of 0.9 to 1.1 percentage points. But the same design estimates nearly the same increase before the policy: 1.0 to 1.2 percentage points at event year -2, so the difference between the post-period estimate and its own pre-treatment value is +0.1 percentage points (s.e. 0.8), and +0.5 (s.e. 1.6) on the balanced-horizon variant. At the point estimate, the year-grain “effect” roughly equals the drift that preceded it. The mechanism is the base period: treated states’ poverty follows a dip-shaped path in the two years before their first step, and a design anchored on a base year averages over the dip, reading the recovery from it as an effect of the treatment that happens to follow. The dispute between Dube (2019b) and Burkhauser, McNichols, and Sabia (2023) is fought entirely on annual data with competing ways of purging exactly this kind of drift; monthly measurement makes the drift and its shape visible instead of assumed away.

The second mechanical finding concerns the within-person job-loss evidence that anchors the loss camp. Clemens and Wither (2019) followed workers observed earning below the incoming 2008-2009 federal minimum and found their employment falling 6.6 percentage points more in states bound by the increase than in states already above it. This paper runs the within-state half of that architecture on the \$15-era wave — bound workers against workers just above the incoming minimum, across 28 events — and adds the check their federal episode could not support: the identical machinery, run where no policy occurred. Bound workers’ employment falls 5.2 percentage points (s.e. 2.3) relative to the just-above group at real events. It falls 3.2 points (s.e. 1.3) in the 21 no-change states, where no increase arrived, and 5.3 points (s.e. 3.5) in treated states’ own pre-treatment years, with no step incoming. The wage premium, by contrast,

appears only where a real increase arrives: +\$0.92 per hour at real events against -\$0.27 in the same states' pre-periods. Being observed at the bottom of the wage distribution predicts churning out of employment everywhere, policy or no policy: at the point estimates, the placebos reproduce between roughly 60 and 100 percent of the raw within-person "job loss," and the placebo-netted employment effects — -2.0 percentage points (s.e. 2.6) against the no-change states, +0.1 (s.e. 4.2) against the pre-period placebo — carry intervals that admit both zero and losses as large as the raw estimate.

The transfer programs the SIPP observes monthly — SNAP and TANF — neither rescue nor claw back. Receipt among low-income families does not fall after state increases (+\$12 per month, s.e. \$9), so the income gains are not offset by phase-outs of these programs at the state grain; for bound workers, receipt rises by an imprecisely estimated \$48 per month (s.e. \$29) net of placebo, consistent with transfers insuring workers through transitions rather than taxing away their gains. The EITC, the largest phase-out margin in the annual accounting, is not observable at monthly frequency in these data.

Together these results make three contributions. To our knowledge this is the first minimum wage evidence from the redesigned SIPP and the first monthly-frequency evidence on income and poverty for the \$15-era wave. The aggregation demonstration bears on the Dube-Burkhauser dispute, which is conducted entirely at the year grain. And the placebo correction added to the bound-worker design disciplines how the most influential within-person job-loss result in the literature should be read, on a policy wave far larger than the one it was estimated on.

Several important topics are outside the scope of this work, and two limits are internal to the design. Public-use SIPP identifies states but not cities, so the Seattle, Denver, and New York City minimums that bind inside their states attenuate the measured treatment toward zero. Monthly state-level cells are noisy, and the estimates come with wide confidence intervals stated throughout. The window contains the pandemic, which is addressed with event-time anchoring, a no-COVID robustness cut, and the placebo designs, but an interaction between the pandemic and treatment timing cannot be fully ruled out. The SIPP's sample is too thin at monthly frequency for teen-specific estimates, so the teen margin, where Neumark and Shirley's synthesis finds strong losses, goes unexamined; their strongest margin, the less-educated, is exactly the group studied here. The first internal limit is spatial: comparing states only within their census divisions

cuts the wage elasticity roughly in half (0.127, s.e. 0.106) and moves the employment estimate negative with a wide interval, and Section V reports why a weak-variation defense is not available — the headline configuration rests on cross-region comparisons. The second is the survey seam: the first stage is concentrated in the cohorts whose steps take effect in January, where SIPP reference years meet, though it survives the direct check of dropping every December and January month (Section V).


The paper proceeds as follows. Section II describes the \$15-era wave and the treatment data. Section III describes the redesigned SIPP and what monthly measurement shows about poverty. Section IV lays out the research design. Section V reports wage and employment results, Section VI poverty and family income results, Section VII the comparison of monthly and annual designs, and Section VIII the bound-worker analysis. Section IX concludes.

II. The \$15-Era Wave

After 2018 the action on the minimum wage moved to statehouses and ballot initiatives. States enacted 173 statutory increases between 2018 and 2024, of which 96 were single steps of 50 cents or more between July 2018 and December 2023. By December 2024, 31 states stood above the federal floor, topping out at \$17.50 in the District of Columbia, and the largest cumulative climb in any state's binding minimum over the window was 75.9 percent. The federal minimum itself did not move; every piece of variation this paper uses is state variation.

Figure 1 shows the resulting divergence. The population-weighted average binding minimum in the 28 states that took at least one large step rose from \$9.57 in January 2018 to \$14.17 in December 2024, while in the 21 states whose binding minimum moved by less than 25 cents over the entire window — the no-change states, in the vocabulary used throughout this paper — it sat at \$7.27. Two states — Michigan and Minnesota — fall into neither group, having moved more than 25 cents in total without ever taking a single 50-cent step; they are excluded from both the treatment and comparison groups rather than assigned to either.

Figure 1. The minimum wage divergence, 2016-2024



Note: Population-weighted average binding minimum wage (the higher of the state and federal minimum). States with a large increase are the 28 states with at least one statutory step of \$0.50 or more between July 2018 and December 2023; no-change states are the 21 states whose binding minimum moved less than \$0.25 over 2018-2024. Source: Vaghul-Zipperer minimum wage panel, extended through 2024 from statutory schedules; authors' calculations.

The timing structure matters for the design. January steps dominate, reflecting statutory effective dates and inflation indexing, but mid-year steps occur throughout the window, most commonly in July and October. Some states legislated multi-year phase-in schedules to a target such as \$15; others index annually to a price level; several did both. This paper defines a state's treatment cohort by its first statutory step of 50 cents or more, which for phase-in states is typically the first leg of the schedule, and treats the later steps as part of the same treatment path rather than as separate events.

The treatment data are the Vaghul and Zipperer (2016) state minimum wage panel, maintained through December 2022, extended by hand through December 2024 from statutory schedules. The extension comprises 58 state-month entries and validates against the published

monthly file with zero mismatches over the 2016-2022 overlap.²

A last definitional point: the binding minimum used throughout is the higher of the state statutory floor and the federal floor, measured at the state level. Sub-state minimums in cities such as Seattle, Denver, and New York bind inside their states, and public-use SIPP cannot see them; where they matter, the measured state-level treatment understates true exposure, which biases the estimated effects toward zero. Table 1 summarizes the wave and the sample.

Table 1. The \$15-era wave and the SIPP sample

	Treated states	No-change states
Number of states	28	21
Average binding minimum, 2018	\$9.57	\$7.27
Average binding minimum, 2024	\$14.17	\$7.27
Share of national population	0.57	0.38
The wave, all states:		
Statutory increases, 2018-2024	173	
Steps of \$0.50 or more, 2018m7-2023m12	96	
States above the federal floor, December 2024	31	
Highest state minimum, December 2024 (DC)	\$17.50	
The SIPP sample:		
Person-months, 2018-2024	4,007,970	
Persons	170,343	
Calendar months	84	
Hourly workers within \$1 of the minimum, 2019 (percent)	17.5	

Note: Treated states are the 28 states with at least one statutory step of \$0.50 or more, July 2018-December 2023; no-change states moved less than \$0.25 over 2018-2024. Averages are population weighted. The District of Columbia is counted as a state. Source: Vaghul-Zipperer panel extended through 2024; SIPP 2018-2024; authors' calculations.

² Two conventions in the published panel required hand-coding in the extension: California and Connecticut carry a convention issue in the commonly used FRED transcription of the panel, and New York is coded at its upstate schedule (\$15.00 from January 2024) because the public-use SIPP identifies the state but not the New York City and Long Island schedules. Maryland's Fair Wage Act acceleration and New Jersey's indexing supersede the originally legislated schedules. The full documentation is in the replication package.

III. Monthly Measurement: The SIPP Redesign


The redesigned SIPP, fielded from 2018, collects a monthly calendar for the prior calendar year at annual interviews: employment status in each month, family income in each month, the family income-to-needs ratio in each month, and program amounts, including SNAP and TANF, in each month. The 2018 through 2024 files together contain 4,007,970 person-months from 170,343 persons across 51 states and 84 calendar months. This is survey microdata with monthly reference periods, not administrative data, and the analysis below treats it accordingly.

The floor is easy to see in these data. In 2019, 17.5 percent of hourly workers earned within one dollar of their state’s binding minimum; in 2024, after six more years of increases, 15.3 percent did. Hourly workers are identified by the redesign’s paid-by-the-hour flag, and the wage is the reported hourly rate at the main job. State-month cells of hourly workers have a median of 94 observations and a minimum of 1, so state-month wage percentiles for the smallest states are noisy; the design weights cells by population, and the bootstrap inference reflects the noise rather than hiding it.³

The central measurement fact is in Figure 2. In 2023, 12.8 percent of people were poor in the average month, meaning their family income that month fell below one-twelfth of the annual poverty threshold. Computing poverty the way an annual survey does — twelve months of income summed, compared with the annual threshold — gives 11.6 percent. And 17.7 percent of people were poor in at least one month of the year. These are three statements about the same people in the same year, differing only in the grain of measurement.

³ The thin cells are concentrated in the smallest states, which carry correspondingly little weight in every population-weighted average the paper reports.

Figure 2. Monthly, annual, and ever-in-the-year poverty, 2023



Note: SIPP 2023 calendar-year file. “Poor in the average month” is the person-month weighted monthly poverty rate; “poor on annual income” compares the within-year mean income-to-needs ratio with 1; “poor in at least one month” is the share of full-year persons with any poor month. Source: authors’ calculations.

What this paper calls hidden poverty is the share of people not poor on an annual accounting who nonetheless spend at least one month poor: 6.9 percent in 2023. (As a share of everyone rather than of the annually non-poor, the same fact is the 6.1-point gap between 17.7 and 11.6.) That poverty measured within the year exceeds poverty measured over the year is a long-established measurement fact — spell-based analyses date to Bane and Ellwood (1986), and the Census Bureau’s episodic-poverty series documents it across SIPP panels (U.S. Census Bureau 2021) — and the redesign carries it into the 2018-2024 window. Month-to-month income variation in a recall survey includes reporting error as well as true dynamics, so absent administrative linkage the 6.9 percent is an upper bound on true within-year poverty churn. Even discounted, it is the population an annual measure averages away.

Hidden poverty is why the wage-floor question is exactly where the measurement grain matters. A minimum wage increase operates on transitions: hours changes, job changes, the arrival and departure of monthly income within families near the line. If the policy shortens poor spells without much changing annual totals, annual measures will miss it. And if poverty in treated and no-change states was already drifting apart before the policy, an annual base period

averages over that drift, with consequences Section VII demonstrates.

Where design-based uncertainty is available, the paper uses it. The redesign ships 240 Fay balanced repeated replication weights; on the replicate-merged sample, the 2022 monthly poverty rate is 13.8 percent with a design-based standard error of 0.3 percentage points. The panel estimates in Sections V through VIII carry state-cluster bootstrap inference, the assignment-level convention in this literature (Section IV), rather than design-based standard errors; carrying the published replicate weights through the full 2018-2024 panel estimates is left to future work.

Two structural features of the redesign matter for reading these facts. First, the monthly calendar is collected at annual interviews, so recall error accumulates over the reference year and seam effects concentrate where interview years meet; the level facts above are inflated by any such noise, which is why hidden poverty is stated as an upper bound. The same recall structure underlies the monthly income and poverty measures used in the designs below, so within-year income dynamics carry reporting error too — the grain comparison of Section VII rests on identical person-months, not on a claim that the SIPP dates income changes precisely. Second, and central to Section VII: the comparison of monthly and annual designs in this paper does not require the SIPP to measure better than the CPS. Both grains are computed from the identical person-months, so measurement error common to both cannot explain why they give different answers.

The paper also fields a companion arm in the data the literature actually fights on: the CPS ASEC, with the official poverty measure, for income years 1989 through 2024 — 1,836 state-years. Section VII uses it to show that the contested annual estimates reproduce in this sample period and specification family, so that the divergence between monthly and annual answers is about the grain and the design, not the dataset or the era.

IV. Research Design

A state enters treatment at its first statutory step of 50 cents or more between July 2018 and December 2023, giving 28 treatment cohorts; the window opens in July 2018 so that every cohort is observed for at least twelve post months. A state already stepping before the window — California’s phase-in began in 2017 — is therefore dated at its first in-window step, with the earlier legs left in its pre-period: differencing from the month before treatment absorbs their level, and any drift they leave in the pre-path is displayed in the 24-month pre-windows and carried by the detrended variants rather than hidden. The 21 no-change states form the comparison pool. The estimand is entry into the raising regime. Later steps in the same state — the second leg of a phase-in, an indexing adjustment — accumulate into the treatment path, so the event-time estimates below are the effect of the path since the first step, scaled by the cumulative dose, not the dynamic effect of a single isolated increase.

The comparison logic is the group-time approach of Callaway and Sant’Anna (2021), implemented on state-month cells. For each treated cohort, the outcome’s change from the month before its first step is compared with the same change among states not yet treated and never treated, weighted by population; the cohort-specific effects are then averaged into event-time coefficients. This construction avoids the contamination that arises in two-way fixed effects estimation when already-treated states serve as controls for later-treated ones (Goodman-Bacon 2021), a problem that is first order in a rollout as staggered as this one.

The dose is the cumulative change in the log binding minimum since the month before the first step, averaged over the first 18 post months and across cohorts: 12.5 log points. Elasticities reported below divide the average post-period effect by this dose. Where an outcome is a rate rather than a log, the effect is first converted to a proportional change at the treated baseline.

Inference is by nonparametric state block bootstrap: 999 draws resampling treated and control states with replacement, with the reported standard error the bootstrap standard deviation of the aggregated estimate. Sampling states — the level at which the policy is assigned — is the convention in this literature because outcomes are serially correlated within states.⁴

⁴ Callaway and Sant’Anna’s multiplier bootstrap is the estimator-native alternative; the state block bootstrap targets the same assignment-level variation on the hand-rolled aggregation used here. With 28 treated clusters, the paper adds a randomization test — states’ entire minimum wage paths permuted across the 49 states, 499 draws — for the two central elasticities (Section V); wild-cluster bootstrap p-values are the remaining complement and are left to future work. The design-based Fay BRR errors of Section III are computed here only for the level estimates; extending them to the panel estimates is left to future work.

The pre-trend policy is a policy, stated once: every event-study path is shown in full from 24 months before treatment (Figure 3), and wherever a pre-trend is visible, a variant that nets out a linear trend fit on the pre-period is reported beside the raw estimate, never instead of it. The paper does not gate any estimate on a pre-trend test, because conditioning on a passed test distorts the inference that follows it (Roth 2022); the partial-identification approach of Rambachan and Roth (2023) is the natural formalization of drift-robust inference for these paths and is left to future work. Which outcomes carry visible drift is reported outcome by outcome, and Section VII shows that the drift is not a nuisance but a finding about why the annual literature disagrees.

Alongside the event studies, the paper reports two-way fixed effects elasticities — the log outcome on the log binding minimum with state and calendar-month fixed effects, population weighted, clustered by state — because they are the nearest cousin of what the annual literature runs, and they bridge to it. Robustness columns re-estimate within census-division-by-month cells, the coarse answer to the spatial-heterogeneity critique of Dube, Lester, and Reich (2010) that states raising minimums differ in regional trends, and drop March 2020 through December 2021.

This wave was announced in advance. Most of the large schedules were enacted well before their first large step took effect — California’s and New York’s \$15 laws passed in 2016 — and event time here is set by effective dates, not enactment dates. Two consequences follow. Pre-periods are post-enactment, so pre-treatment paths can carry anticipatory adjustment; the flat wage pre-paths are direct evidence that employers were not raising wages ahead of effective dates, but anticipation on other margins cannot be excluded, and Section VII returns to what that means for the poverty drift. And states with enacted but not-yet-effective schedules sit in the not-yet-treated comparison pool, in exactly the sense Callaway and Sant’Anna’s anticipation extension formalizes; a design dated at enactment is left to future work.

The pandemic sits in the middle of the sample, and stimulus payments and expanded unemployment insurance moved low-income families’ incomes far more than any wage floor did. Three features of the design address it: event-time anchoring absorbs shocks common to treated and control states; the no-COVID variant drops the affected months entirely; and the bound-worker placebos of Section VIII difference out era effects altogether. States also raise minimums when their economies are strong — overall working-age employment shows a small positive association with the minimum (elasticity 0.040, s.e. 0.023) — and that same cyclical


selection is the most natural source of the poverty pre-drift Section VII documents. The design displays the drift rather than absorbing it into a base period, but it cannot fully exclude a pandemic-era interaction with treatment timing.

V. Wages and Employment

Low-wage workers' earnings grew everywhere between 2018 and 2024 — in Mississippi at \$7.25 as in California on its way to \$16 — because the era combined a tight labor market with the fastest nominal wage growth at the bottom of the distribution in decades. Before-and-after growth in a treated state therefore says nothing about the policy. Every estimate below is a difference in paths between treated and no-change states. Section VIII returns to how strong that common growth at the bottom was, because it is the key to reinterpreting the within-person evidence.

Figure 3 shows the four central event-study paths; Table 2 reports the wage and employment estimates. The first stage is in panel A: after a state's first large step, its 10th-percentile log hourly wage rises 3.0 log points (s.e. 2.3) on the raw path and 6.6 log points (s.e. 3.8) net of the pre-period trend, against an average dose of 12.5 log points. As elasticities, those are 0.24 (s.e. 0.18) and 0.53 (s.e. 0.31). The two-way fixed effects version, using all of the variation, gives 0.247 (s.e. 0.093) — the same point estimate as the raw group-time elasticity, with half the standard error. Point-estimate agreement in one sample does not validate the two-way fixed effects estimand under staggered adoption; the paper therefore reports both numbers, uses the more precise one where a single number scales a calculation, and carries the event-study interval beside every such calculation. Two design-based checks bear directly on the first stage: a randomization test that reassigns entire state minimum wage paths across the 49 states gives $p = 0.02$, and the estimate is 0.245 (s.e. 0.094) when every December and January calendar month is dropped — the no-seam check discussed below. The wage measure conditions on being paid by the hour, and hourly status itself could respond to the policy; the estimates take that margin as given.

Figure 3. Event-study paths: wages, employment, poverty, and the near-poor



Note: Event time is months since the state's first statutory step of \$0.50 or more. Coefficients are population-weighted group-time comparisons against not-yet- and never-treated states, anchored on the month before treatment; vertical bars are 95 percent intervals from a 999-draw state block bootstrap. Source: SIPP 2018-2024; authors' calculations.

Table 2. Wage, employment, and hours estimates

Outcome	Event-study ATT, raw	Event-study ATT, detrended	TWFE elasticity, base	TWFE, division x month	TWFE, no COVID
10th-percentile log hourly wage	+0.030 (0.023)	+0.066 (0.038)	+0.247 (0.093)	+0.127 (0.106)	+0.265 (0.095)
25th-percentile log hourly wage	+0.014 (0.016)	+0.024 (0.029)	+0.095 (0.036)	+0.035 (0.031)	+0.112 (0.038)
Employment rate, HS or less, ages 16-64 (pp)	+0.3 (1.2)	+0.6 (2.2)	-0.009 (0.038)	-0.066 (0.048)	-0.023 (0.042)
Employment rate, all ages 16-64 (pp)	+0.6 (0.8)	+1.2 (1.6)	+0.040 (0.023)	+0.010 (0.029)	+0.031 (0.023)
Weekly hours at the main job, employed HS or less	-0.1 (0.2)	-0.2 (0.4)	-0.044 (0.024)	-0.027 (0.031)	-0.047 (0.026)

Note: Event-study entries are average treatment effects on the treated (ATT) over post months 0-18 (raw, and net of a linear trend fit on months -12 to -2), in log points for wages, percentage points for employment, and weekly hours for the hours row. TWFE (two-way fixed effects) entries are elasticities of the outcome with respect to the binding minimum (state and month fixed effects; division $\times$ month and no-COVID variants as labeled), with state-clustered standard errors. Source: SIPP 2018-2024; authors' calculations.

The wage effect is compression, not a general rise. The 25th-percentile elasticity is 0.095 (s.e. 0.036), roughly 40 percent of the 10th-percentile response, and consistent with spillovers that reach modestly above the floor and fade (Autor, Manning, and Smith 2016; Cengiz et al. 2019). This gradient is what a binding minimum should produce and what a low-wage boom should not: a boom lifts the 25th percentile along with the 10th.

The first stage moves as the specification narrows. Dropping the pandemic months raises the elasticity slightly, to 0.265 (s.e. 0.095). Forcing all comparisons within census divisions cuts it to 0.127 with a standard error of 0.106 — attenuated and statistically indistinguishable from zero. A weak-variation defense is not available: 70 percent of the residual variance of the log minimum survives the division-by-month cells, so the within-division comparison genuinely estimates a smaller first stage rather than running out of variation. The reading is symmetric with Section VIII’s treatment of the same argument between Zipperer and Clemens: within-region comparisons remove regional confounds and some treatment contrast at once, the cross-region comparisons carry the headline first stage, and this window cannot separate the two readings. Leaving out one treated state at a time moves the elasticity between 0.19 and 0.29, so no single state, California included, carries it. Splitting the treated states by cohort type gives 0.232 (s.e. 0.135) for the states whose first step was a dollar or more and 0.248 (s.e. 0.084) for the indexers and phase-in states, so the first stage is not the artifact of one cohort type; Clemens and Strain (2021), who committed in advance to comparing large and small increases over roughly this era, find wage effects largest in the large-increase states, and the coarser split here is consistent with theirs on the wage side.

A different split addresses the survey seam. SIPP collects each calendar year’s months at a single interview, so reference years meet at the December-January boundary, and most statutory steps take effect in January — an artifact of the seam, such as workers retro-reporting a new statutory wage across the boundary, would therefore load onto January cohorts. The first stage is in fact concentrated there: the 16 January cohorts carry a first-stage estimate of +5.6 log points (s.e. 3.1), while the 12 mid-year cohorts give -0.3 (s.e. 2.7). The direct check is to remove the seam months entirely: re-estimating the two-way fixed effects elasticity with every December and January calendar month dropped — so no observation sits on either side of a reference-year boundary or a January step — gives 0.245 (s.e. 0.094), unchanged from the full-sample 0.247. Cross-seam reporting cannot carry an estimate computed without the seam months. What the

check does not explain is the mid-year null itself: a statutory increase mechanically raises wages below the new floor, so a flat mid-year first stage more plausibly reflects within-reference-year recall smoothing — respondents reporting one wage for stretches of a year — than the absence of a wage effect, and that reading limits how finely these data can date wage changes inside a reference year. The split also confounds step timing with cohort composition. The employment estimates are near zero in both groups (+0.5 percentage points, s.e. 2.0, January; +0.1, s.e. 1.7, mid-year), and the no-seam employment elasticity is -0.006 (s.e. 0.038), so nothing downstream turns on the seam question; the mid-year measurement caveat stands.

Assembled into one sentence, the skeptical reading of the first stage is that it lives in cross-region comparisons among January-effective cohorts, in a window that contains the pandemic. The answer is that the estimate survives every cut that can be made without destroying the treatment contrast itself: excising the seam months (0.245), dropping the pandemic months (0.265), leaving out any single treated state (0.19 to 0.29), and permuting entire treatment paths ($p = 0.02$) — while the one cut it does not survive, the within-division comparison, removes real treatment variation along with regional confounds, and the CPS anchor of Section VII shows this sample reproduces the annual literature’s own estimates. What cannot be excluded is a cross-region confound correlated with January-effective schedules operating in exactly this window; no design with 28 treated clusters can rule that out, and the width of the within-division interval is the measure of the residual doubt.

Employment is the cost side, and the estimates are in panel B and the lower rows of Table 2. The employment rate of workers with a high school degree or less, ages 16 to 64, moves +0.3 percentage points (s.e. 1.2) after treatment on the raw path and +0.6 (s.e. 2.2) detrended, from a baseline of 58.6 percent. The two-way fixed effects elasticity is -0.009 (s.e. 0.038); leaving out one treated state at a time moves it between -0.024 and +0.001, and the path-permutation randomization test gives $p = 0.89$ — the employment estimate looks exactly as a true null should under design-based inference. Working-age employment overall behaves the same way (+0.6 percentage points, s.e. 0.8; elasticity 0.040, s.e. 0.023). Nothing in these paths looks like the aftermath of a large negative demand shock to low-wage labor.

Hours are the margin the Seattle evidence moved on — Jardim et al. (2022) found hours worked in low-wage jobs falling 6 to 7 percent after that city’s large increase — and the redesign measures them monthly. Weekly hours at the main job among employed workers with a high

school degree or less move -0.1 hours (s.e. 0.2) after treatment, from a 37.4-hour base, and the two-way fixed effects hours elasticity is -0.044 (s.e. 0.024): negative, small, and on the margin of statistical significance. The configuration is consistent with a modest intensive-margin adjustment accompanying an employment count that does not move, and it is far from the Seattle magnitudes; it is also the one outcome in Table 2 whose point estimate is negative in every specification.

The literature's summary statistic for comparing such results is the own-wage employment elasticity: the employment response divided by the wage response. Dividing the low-education two-way fixed effects employment elasticity by the 10th-percentile wage elasticity gives -0.04, with a delta-method standard error of 0.15 that uses both sampling variances while taking the numerator and denominator as uncorrelated — an approximation, since both are estimated on the same states — and a 95 percent lower bound of -0.33. Built instead entirely from the event-study pair, the ratio is +0.19 with a standard error of 0.70 — the same near-zero answer, from the estimator that avoids leaning on two-way fixed effects, at the price of an interval too wide to bound anything. Two scope statements bound the comparison. The 10th-percentile wage response is a proxy for the wage change of affected workers, whereas the bunching literature estimates jobs and wages jointly for workers near the floor, so the two constructions are cousins rather than twins. And the bound is a broad-group bound: only about 9 percent of employed low-education workers are paid within a dollar of the minimum, so if the group's employment response were concentrated entirely on affected workers, a group bound of -0.33 would translate to an affected-worker bound several times larger, and it excludes nothing at the scale of Clemens and Wither's recession-era estimate — the person-level evidence of Section VIII is what speaks to that scale. What the bound does establish: broad-group elasticities near -1 are excluded, and the estimate is consistent with Cengiz et al.'s bunching estimate of 0.41 (s.e. 0.43) — the confidence intervals overlap — and with the -0.04 median for broad worker groups in Dube's (2019a) review.

The strongest opposing synthesis should be stated in its own words. Assembling the literature's preferred estimates, Neumark and Shirley (2022) report that “79.2% of the estimated employment elasticities are negative,” with 53.8 percent negative and significant at the 10 percent level, and that “this evidence of negative employment effects is stronger for teens and young adults, and more so for the less-educated.” That last clause bears directly on this paper,

because the less-educated are exactly the group studied here: the elasticity estimates in their less-educated corpus have a mean of -0.24 and a median of -0.18, and the two-way fixed effects estimate above, -0.009 with a 95 percent interval of -0.08 to +0.07, excludes both. The estimates are in genuine tension, and the reconciliations run through what is being studied: their less-educated corpus is drawn from older and smaller variation, while the group in their assembly whose outcome is closest to the bunching literature's — employment in low-wage industries — centers near zero (mean 0.01, median -0.02). The teen margin, where their evidence is also strong, is genuinely not examined here: monthly SIPP cells are too thin for teen-specific estimation, a scope limit stated in Section I. The present estimate is a broad low-education estimate on the newest and largest variation, and the tension with the older less-educated corpus stands.

One cell of Table 2 looks like the other camp's result. The division-by-month employment estimate is -0.066 (s.e. 0.048); divided by the division-by-month wage estimate of 0.127, it would give an own-wage elasticity near -0.5, squarely in job-loss territory. But that division has no identified denominator: the within-division first stage is statistically indistinguishable from zero (s.e. 0.106), so the ratio is uninformative — the weak-first-stage problem in its plainest form. The reduced-form statement is that within-division employment moves -0.066 with an interval spanning zero, alongside a first stage too weak to scale it. The no-COVID employment estimate is -0.023 (s.e. 0.042), and the cohort split gives 0.000 (s.e. 0.035) for the big-step states and 0.009 (s.e. 0.069) for the indexers.

All of this said, the employment evidence has real limits. Pre-period drift is visible in several outcomes, the monthly state cells are noisy, and the confidence interval on the own-wage elasticity admits modest losses — a bound of -0.33 is not a bound of zero. The estimates are also averages over the first 18 post months of multi-year phase-ins: at months 18 to 24 the raw employment path stands at +0.8 percentage points (s.e. 1.3), so nothing negative emerges at the horizons this window can see, but losses that accumulate with the cumulative dose over longer horizons — the dynamic pattern Meer and West (2016) argue is where minimum wage effects live — would not yet be visible here. What the estimates rule out is the large-loss reading: employment elasticities near -1, or job losses concentrated enough among low-education workers to show up in a 4-million-person-month panel spanning the largest wave on record. Overall, however, the configuration of a strong first stage with employment tightly around zero matches

the near-zero camp on the wave that matters most for current policy.

Why a binding floor need not destroy jobs is a question this paper observes rather than answers. Monopsony power can absorb wage floors without employment losses, and the pass-through from the minimum into employers' labor costs may itself be weak — turnover falls, and other terms of the job adjust (Manning 2021); firms adjust on prices, amenities, effort, and compliance margins that employment counts do not capture (Clemens 2021); and in the German introduction, workers reallocated toward more productive establishments with no aggregate employment loss (Dustmann et al. 2022). The estimates here are consistent with that family of mechanisms and do not distinguish among them.

VI. Poverty and Family Incomes

Poverty is where the two camps stake opposite claims on the same annual data. Dube (2019b), using clean-control event studies on the CPS, finds long-run poverty elasticities with respect to the minimum wage of -0.22 to -0.46. Burkhauser, McNichols, and Sabia (2023), extending the null line of Burkhauser and Sabia (2007) and Sabia and Burkhauser (2010) with dynamic difference-in-differences on nearly four decades of March CPS, conclude that “we can rule out long-run poverty elasticities... of less than -0.129.” The ranges do not overlap; both cannot be right. This section reports the monthly estimates; Section VII shows why the annual fight is stuck.

The monthly poverty estimates are in Figure 3, panel C, and Table 3. The raw effect is zero: +0.03 percentage points (s.e. 0.49) on a treated baseline of 13.6 percent. Net of a linear pre-trend the path is negative: -1.38 percentage points (s.e. 1.01). But the adjusted estimate is fragile in a way the raw estimate is not: fit the trend on months -12 to -2 and the estimate is -1.4; on months -9 to -2, -1.2; on months -6 to -2, +0.1. The pre-path is not a line — Section VII returns to its actual shape — so the linear adjustment is a functional-form choice that moves the estimate by more than a percentage point. The two-way fixed effects poverty estimates in Table 3 point the other way — small, positive, and indistinguishable from zero (0.03, s.e. 0.14, in the base specification) — so the favorable lean in this section comes entirely from the detrended paths. The defensible summary is that the raw effect is zero, the adjusted paths are negative but fragile to the adjustment, the two-way fixed effects estimates are mildly positive and noisy, and the intervals decide nothing.

Table 3. Poverty, income, and transfer estimates

Outcome	Event-study ATT, raw	Event-study ATT, detrended	TWFE elasticity, base	TWFE, division x month	TWFE, no COVID
Monthly poverty rate (pp)	0.0 (0.5)	-1.4 (1.0)	+0.032 (0.144)	+0.186 (0.165)	+0.028 (0.152)
Monthly child poverty rate (pp)	+0.2 (1.1)	-1.6 (1.8)	+0.196 (0.252)	+0.415 (0.273)	+0.183 (0.263)
Share below 150 percent of needs (pp)	-1.0 (0.6)	-2.3 (1.1)	+0.023 (0.113)	+0.168 (0.130)	+0.018 (0.118)
10th-percentile income-to-needs	-0.012 (0.035)	+0.070 (0.060)	+0.007 (0.174)	-0.033 (0.173)	-0.024 (0.187)
25th-percentile income-to-needs	+0.092 (0.039)	+0.149 (0.076)	-0.002 (0.076)	-0.101 (0.087)	-0.008 (0.079)
SNAP+TANF dollars per month, income < 2x needs	+\$12 (\$9)	+\$11 (\$15)			
Implied poverty elasticity	+0.01 (0.29)	-0.81 (0.59)			
Implied near-poor elasticity	-0.36 (0.23)	-0.87 (0.41)			

Note: Same design and conventions as Table 2 (division × month as labeled). Income-to-needs percentiles are levels of the family income-to-poverty ratio; SNAP+TANF is mean monthly receipt in dollars among families below twice the poverty line. Implied elasticities divide the proportional post-period effect by the 12.5 log point average dose. Source: SIPP 2018-2024; authors' calculations.

Child poverty behaves the same way with more noise: +0.2 percentage points raw (s.e. 1.1), -1.6 detrended (s.e. 1.8), on a baseline of 18.6 percent.

The near-poor margin is the one that moves (Figure 3, panel D). The share of people below 150 percent of needs — a measure that includes the poor — falls 1.0 percentage point (s.e. 0.6) raw and 2.3 (s.e. 1.1) detrended, from a base of 21.6 percent, and the estimate is negative in every adjustment window (-2.3, -2.6, -1.2). But the raw estimate is imprecise, and the near-poor share carries a statistically significant pre-period drift of its own (-0.7 percentage points over the pre-year, s.e. 0.3), so the finding is “negative in raw and every adjusted variant,” not “settled.” The magnitude also strains: the detrended -2.3 implies an elasticity of -0.87, larger than anything in the annual literature, and the person-level earnings gains of Section VIII are too small and too imprecise to obviously deliver it; the raw -1.0 is the end of the range the paper’s own accounting can support, and the detrended figure should be read as an upper envelope produced by the same fragile adjustment as everywhere else.

The near-poor margin moving while the poverty rate barely moves is what the distribution of affected families predicts. Among the bound workers of Section VIII — workers observed earning below an incoming minimum — 17.5 percent are in poor families at baseline and 29.6 percent are below 150 percent of needs, so the 100-to-150 band alone holds 12.1 percent. Nearly

a third of bound workers' families sit below one and a half times the line, and that band is where a binding floor's income gains land. At the same time, bound workers' mean family income-to-needs ratio is 3.6, pulled up by a long right tail of secondary earners in non-poor families, so the targeting is real but imperfect. The nearest distributional neighbors agree on where the action is: Rinz and Voorheis (2018), in linked Census administrative data, find persistent earnings gains at the bottom of the distribution, and Godoey and Reich (2021) find poverty falling in high-impact counties even at minimum-to-median ratios up to 0.82.

The monthly point estimates lean negative where drift is adjusted; the near-poor margin is negative in every event-study variant; and the intervals contain both camps' ranges. On the elasticity magnitude, this design does not decide. What Section VII shows is different and sharper: not that a positive poverty effect is statistically excluded — the raw monthly interval reaches +1.0 percentage points — but that the positive estimate the year-grain design produces on these same data predates treatment, so the year-grain evidence for it is not evidence about the policy.

Furman argued in 2014, as chairman of the Council of Economic Advisers, that the distributional benefits were central to the case for a moderate increase. The monthly evidence here neither confirms that case at Dube's magnitudes nor contradicts it: it is consistent with poverty declining, decisive mainly against reading the \$15-era wave as having raised poverty, and clearest on the near-poor margin where the affected families actually sit.

The income-to-needs distribution tells the same story with wider intervals. The 10th percentile of the family income-to-needs ratio is noisy in every variant and never distinguishable from zero. The 25th percentile shows a large positive event-study estimate — an implied elasticity of 0.42 — that sits atop a visible pre-period drift (+3.1 log points over the pre-year, s.e. 2.4) and shrinks toward zero in the two-way fixed effects specification (-0.002, s.e. 0.076); it is reported in Table 3 with its pre-trend and not leaned on.

Transfers could undo everything above — if benefit phase-outs claw back what the floor delivers, market income gains overstate net gains. Dube (2019b) estimates that phase-outs claw back roughly a third of gross income gains; Clemens (2016), on the loss side of the 2008-2009

episode, found transfers offsetting almost none of displaced workers' losses. At the state grain, SNAP plus TANF receipt among families below twice the poverty line does not fall after increases: +\$12 per month (s.e. \$9) raw, +\$11 detrended, on a base of \$245. No clawback through these two programs is visible at this grain.⁵ Section VIII sharpens the transfer question at the person level.

The targeting critique remains. MaCurdy (2015) argues that because minimum wage earners are spread across the family income distribution and price pass-through falls on all consumers, the minimum wage is poorly targeted as an antipoverty tool — on his accounting, roughly 35 percent of after-tax benefits reach families below twice the poverty line.⁶ The bound-worker distribution above concedes the premise: many beneficiaries are secondary earners in non-poor families. What the present results add is where the benefits that do reach low-income families land: the 100-to-150-percent band, which falls in the raw estimate and every adjusted event-study variant.

VII. Why Annual and Monthly Answers Differ

When the identical person-months are aggregated to the year grain and the identical design is re-run, the monthly zero of Section VI becomes a poverty increase of 0.9 to 1.1 percentage points — but the design estimates nearly the same increase before the policy, and the formal difference between the post-period estimate and its own value at event year -2 is +0.1 percentage points (s.e. 0.8) for the annualized-income measure and -0.3 (s.e. 0.8) for the mean-of-monthly-rates measure. The differences carry wide intervals, so what follows is a statement about point estimates, stated as such: at the point estimate the year-grain “effect” equals its own pre-treatment value, and an estimate with that property is evidence of drift, not of policy.

Table 4 lays the designs side by side. Two aggregation rules are shown, because an annual survey could deliver either: poverty on annualized income, the CPS concept, and the within-year mean of the monthly poverty rate. Both produce positive post-period estimates, with implied

⁵ The EITC margin is invisible here even though its parameters are fixed: a worker whose earnings rise moves along the phase-in or phase-out schedule, and the SIPP does not measure EITC amounts monthly. Every clawback statement in this paper is therefore about SNAP and TANF only; the EITC interaction — amplifying gains for single mothers, worsening outcomes for some childless low-skill groups (Neumark and Wascher 2011) — is left to annual data.

⁶ The characterization of MaCurdy (2015) follows its abstract and the published discussions of it; the claims attributed to it therefore carry no page citations.

poverty “elasticities” of +0.29 to +0.39 — larger than anything in the annual literature, and not because of the scaling, since the year-grain rows divide by their own larger cumulative dose — and both produce nearly identical estimates at event year -2, before the policy. The balanced-horizon row addresses a composition concern: restricting to cohorts treated by 2021, each of which contributes the full post window, the pattern reproduces — post +1.4, pre-treatment +0.9, difference +0.5 (s.e. 1.6). The grain gap is not an artifact of late cohorts contributing short horizons.

Table 4. The same data and design at two grains

Design	Post ATT (pp)	ATT at k = -2 (pp)	Post minus k = -2 (pp)	Implied elasticity
Month grain, raw	0.0 (0.5)			+0.01 (0.29)
Month grain, detrended	-1.4 (1.0)			-0.81 (0.59)
Year grain, annualized income	+1.1 (0.4)	+1.0 (0.8)	+0.1 (0.8)	+0.39 (0.13)
Year grain, mean of monthly rates	+0.9 (0.4)	+1.2 (0.7)	-0.3 (0.8)	+0.29 (0.12)
Year grain, annualized, balanced horizon	+1.4 (0.4)	+0.9 (1.6)	+0.5 (1.6)	+0.65 (0.18)
CPS ASEC, official poverty, income years 1989-2024				-0.11 (0.08)
CPS ASEC, official poverty, income years 2013-2024				-0.02 (0.06)

Note: All rows use SIPP 2018-2024 person-months and the same cohort definitions as Table 2, except the CPS ASEC rows. Year-grain designs anchor on the calendar year before the cohort’s first step; the balanced-horizon row restricts to cohorts treated by 2021 with a post window of event years 0 to 2. The CPS ASEC rows are TWFE elasticities of the official poverty rate, state and year fixed effects. Year-grain implied elasticities divide the proportional post-period effect by each row’s own baseline and cumulative dose (11.8 percent and 24.3 log points for the full-sample rows; 12.0 percent and 17.3 log points for the balanced-horizon row). Source: SIPP 2018-2024; CPS ASEC survey years 1990-2025; authors’ calculations.

The mechanism is visible in the estimated shape, and it is not the textbook smooth drift. Extended to 24 months before treatment, the monthly path in Figure 3 shows the full shape: treated-minus-control poverty sits near zero at months -24 to -13 (mean -0.1 percentage points), falls to roughly -1 percentage point about a year before treatment, and recovers into the anchor month — the slump-then-recovery shape known from the job-training literature as an Ashenfelter dip (Ashenfelter 1978), not a straight line. The dip is observationally equivalent to anticipatory adjustment to schedules enacted a year or two earlier (Section IV); on either reading the base period is contaminated, which is what matters here. A yearly design anchors on a base year that averages over the dip, so both earlier years and later years sit above the base, and the estimator reads the recovery as a level effect of the treatment that happens to follow it. A monthly design anchors on the single month before the step and displays the whole shape. The two pre-period

signs that look contradictory at first — the monthly pre-path sits below its anchor month, while the annual estimate at year -2 sits above its base year — are the same V shape read against different anchors. The problem is not that an annual event study cannot display leads; Table 4’s own year -2 estimate is a displayed lead, and both camps in the annual literature display them. The problem is that anchoring on a base period that averages over a non-linear sub-annual path builds the dip into every estimate, and a short panel gives the annual design one or two pre-periods with which to see a shape that needs several. The same non-linearity is why Section VI presented the paper’s own linear detrend as fragile rather than as truth: fitting a line to the recovery limb of a V is the mirror image of the annual design’s mistake.

Dube (2019b) and Burkhauser, McNichols, and Sabia (2023) fight with distributed lags, event windows, and trend controls on annual CPS data — competing machinery for purging drift whose sub-annual shape annual data cannot characterize. The demonstration here concerns the base-period anchoring those designs share, run on a seven-year window; it does not run either side’s exact specification, which would require the multi-decade panel they use, and it therefore shows a mechanism their designs are exposed to rather than an error either has been caught committing. Two caveats bound the claim further. The CPS companion arm below, at the year grain over 2013-2024, produces a mildly negative elasticity, not the positive artifact — so nothing here shows the multi-decade CPS designs suffer the artifact to the degree this window does. And within 2018-2024 the year-grain design cannot be run with the pandemic excised: at the year grain every treated cohort’s post period touches 2020-2021, so pandemic-era transfers — which moved low-income annual incomes on a scale no wage floor approaches — remain a live rival explanation for the drift itself. A pre-pandemic version of the dip diagnostic is equally unavailable: the redesign begins in 2018, so the earliest cohorts lack the two-year pre-window the diagnostic needs, and the cohorts that have it ran their pre-periods through 2020-2021. Monthly measurement does not crown a winner. It makes the drift and its shape visible — which is how the +1 percentage point year-grain “effect” above is caught matching its own pre-treatment value.

The CPS companion arm anchors this sample to the literature’s own yardstick. On CPS ASEC data, official poverty on the log minimum with state and year fixed effects gives an elasticity of -0.109 (s.e. 0.075) for income years 1989 through 2024 and -0.020 (s.e. 0.055) for income years 2013 through 2024 — estimates within the span of the contested annual literature, from its own data, though the specification is the two-way fixed effects cousin of the contested

designs rather than either side's exact estimator. The anchor is approximate, but it is enough to establish that this sample and era are not an outlier from which annual conclusions could not have been drawn.

One objection says the year grain is right and the month grain is underpowered noise. Power widens confidence intervals; it does not date effects before their causes. The pre-treatment test is decisive against reading the year-grain level effect as causal whatever one believes about monthly power — and the concession runs the other way too: monthly intervals are wide; that is the cost of the finer grain.

A stronger objection is that the annual estimate captures a real effect that arrives with a lag — hours cuts and exits accumulating — and the raw monthly path itself rises to +0.8 percentage points (s.e. 0.5) at months 18 to 24. Two answers, each with its limit stated. First, the same pre-period drift that the year-grain critique requires taking seriously predicts more than that when continued forward: net of the drift continuation, the late-horizon path is -2.1 percentage points (s.e. 1.5). Second, whatever is happening at long horizons, the year-grain estimate specifically equals its own pre-treatment value, so the annual design is not measuring it. What remains true is that at 18 to 24 months this design cannot cleanly separate lagged effects from drift, and the intervals span both readings.

Hidden poverty and the aggregation artifact are the same fact seen twice. Annual measures average over the months where the action is; when what is averaged away differs between treated and no-change states, averaging becomes bias. That caution is this paper's methodological export, and it applies to short-panel, year-grain evaluations of any staggered state policy — paid leave, Medicaid expansions, UI rules — where the base period must be anchored inside a window too short to characterize the drift it averages over.

VIII. Bound Workers and the Churn Correction

State-level paths answer the policy question but dilute it, because most residents of a treating state are nowhere near the wage floor. The sharpest job-loss evidence in the literature therefore tracks the workers the floor actually binds. Clemens and Wither (2019), in the classic SIPP panel, followed workers earning below the incoming 2008-2009 federal minimum and found their employment falling 6.6 percentage points more in the states the federal increase actually bound than in states whose minimums already exceeded it, with workers just above the cutoff serving as a within-state control in both groups of states. That estimate — behind the employment elasticity near -1 with respect to the minimum wage that anchors the job-loss camp — has been contested on exactly the question of what the unbound states control for: Zipperer (2016) shows that sectoral and geographic controls cut it by 35.6 to 62.7 percent, arguing that the unbound states were hit differently by the Great Recession; Clemens (2017) responds that those controls dilute the treatment variation itself — scaling for the dilution recovers nearly half the difference — and that within-region comparisons carry an upward bias of their own. The diagnostic added here is different. It runs the within-state contrast — workers observed earning below an incoming state minimum (bound workers, $n = 1,004$) against workers earning from the new minimum to 1.5 times it (just-above workers, $n = 1,522$), across 28 events — and then runs the identical machinery where no policy occurred, which the 2008-2009 federal episode could not support: every state was treated, and no pre-period was undisturbed.


The design compares within-person changes: each worker's outcome path relative to their own pre-event baseline, bound minus just-above, from 6 months before the step to 12 after. The just-above group is itself exposed to the policy through spillovers, so the contrast nets out anything common to both groups and identifies the differential effect on the most exposed workers, not the total effect of the policy.

Who are the bound? At baseline they earn \$9.48 an hour on average and are employed 88 percent of months; their families' monthly earnings are \$1,325, their SNAP plus TANF receipt is \$280 per month, 17.5 percent are poor and 29.6 percent live below 150 percent of needs. This is the population the policy is about.

Taken at face value, the within-state contrast looks like job loss on the Clemens-Wither scale, alive on this wave. Relative to just-above workers, bound workers' hourly wages rise \$0.92

(s.e. \$0.37) after the step, their employment falls 5.2 percentage points (s.e. 2.3), their monthly earnings rise \$168 (s.e. \$107), and their SNAP plus TANF receipt rises \$54 (s.e. \$24). Figure 4 shows the paths; Table 5 reports the estimates.

Figure 4. Bound-worker paths, real events and the no-change placebo



Note: Within-person bound-minus-just-above gaps relative to each worker's pre-event baseline, by month since the (pseudo-)event. Bound workers earn below the incoming minimum at baseline; just-above workers earn between the new minimum and 1.5 times it. Placebo pseudo-events apply the real events' step ratios in the 21 no-change states. Source: SIPP 2018-2024; authors' calculations.

Table 5. Bound workers: real events, two placebos, and the nets

	Real events	Placebo: no-change states	Placebo: treated states, pre-period	Net of no-change placebo	Net of pre-period placebo
Events	28	270	13		
Bound workers	1,004	1,334	252		
Just-above workers	1,522	4,391	675		
Baseline means, bound workers:					
Hourly wage	\$9.48				
Employed (share of months)	0.879				
Monthly earnings	\$1,325				
SNAP+TANF per month	\$280				
Poverty rate	0.175				
Near-poverty rate (below 150 percent)	0.296				
Post-event gaps, bound minus just-above:					
Hourly wage (\$)	+\$0.92 (\$0.37)	+\$1.62 (\$0.53)	-\$0.27 (\$0.49)	-\$0.70 (\$0.64)	+\$1.19 (\$0.61)
Employed (pp)	-5.2 (2.3)	-3.2 (1.3)	-5.3 (3.5)	-2.0 (2.6)	+0.1 (4.2)
Monthly earnings (\$)	+\$168 (\$107)	+\$324 (\$95)	+\$295 (\$189)	-\$155 (\$143)	-\$126 (\$217)
Poverty rate (pp)	-2.0 (1.9)	-2.9 (1.5)	-3.5 (2.8)	+0.9 (2.4)	+1.5 (3.4)
Below 150 percent of needs (pp)	-1.7 (1.8)	-5.5 (1.4)	-4.7 (1.7)	+3.8 (2.3)	+3.1 (2.5)
SNAP+TANF per month (\$)	+\$54 (\$24)	+\$7 (\$16)	-\$35 (\$23)	+\$48 (\$29)	+\$89 (\$34)

Note: Post-event entries are within-person bound-minus-just-above gaps averaged over months 0-12. The no-change placebo runs the identical machinery at pseudo-events in the 21 no-change states; the pre-period placebo runs it in treated states 18 months before their real first step (cohorts with a clean runway). Each net column subtracts the named placebo; standard errors combine the independent bootstraps, and nets and intervals are computed from unrounded estimates, so they can differ in the last digit from differencing the displayed cells. Source: SIPP 2018-2024; authors' calculations.

Whether these gaps measure the policy can be checked directly, because the design can be run where no policy occurred. The check is the observational analogue of the unbound-state arm Clemens and Wither built into their triple difference: their design already nets the within-state gap against states the increase did not bind, and the placebo here plays that role for the \$15-era wave. In the 21 no-change states, each real event's date and proportional step is applied in each no-change state as a pseudo-event, and the 270 with estimable samples — 1,334 pseudo-bound and 4,391 just-above workers — produce an “employment effect” of -3.2 percentage points (s.e. 1.3), a “wage effect” of +\$1.62, and an “earnings effect” of +\$324, with no minimum wage change having occurred; the placebo's wage climb is also the direct measure of how strong the era's common low-wage growth was, the caution Section V opened with. One objection to this

placebo is fair: workers below an unchanged floor are a sub-minimum population — tipped, exempt, or misreporting — at a different point of the wage distribution than workers bound by an incoming \$15 schedule.

So the paper runs a second placebo immune to that objection: the same treated states, in their own pre-treatment years. For cohorts with a clean 19-month runway, pseudo-events are placed 18 months before the real first step — 13 events, 252 bound and 675 just-above workers, same states, same wage environment, no step yet effective (for several phase-in states a schedule was already enacted, which makes the wage result below the direct evidence that enactment alone was not moving wages). The “employment effect” is -5.3 percentage points (s.e. 3.5), as large as at the pooled real events; restricted to the same 13 cohorts, the real-event gap is -9.2 (s.e. 6.0), so the matched comparison puts the placebo at roughly 60 percent of the real gap, in line with the no-change placebo. The “wage effect” is -\$0.27 (s.e. \$0.49) — zero. Two composition caveats: twelve of the thirteen pseudo-event windows touch the pandemic months, and the 13-cohort subset is the later-treated cohorts, not a random draw. Within the very states that would later raise their minimums, workers observed at the bottom of the distribution churned out of employment at the same differential rate before any step took effect; what they did not get was a wage increase.

The explanation is mean reversion in its labor-market form, and the hazard is a known one: Currie and Fallick (1996) ran an affected-worker design of this family in the NLSY, and Cengiz, Dube, Lindner, and Zentler-Munro (2022) built a machine-learning definition of exposed workers that does not condition on a realized low wage. The placebos here diagnose the bias inside the design rather than avoiding it, which is what permits reinterpreting the recession-era estimate. A worker observed at the bottom of the wage distribution is disproportionately there transiently — young, between jobs, in a bad month — and over the following year climbs the distribution and churns out of employment at higher rates than a worker observed slightly higher, everywhere, policy or no policy. The pre-event paths are flat by construction, since each worker is differenced from their own baseline; the divergence is what follows being observed at the bottom. Differential attrition does not drive the comparison: 38 percent of real-event bound workers and 41 percent of placebo bound workers are still observed at month 12.⁷

⁷ These shares mix panel attrition with calendar-window truncation — a worker whose event falls late in the panel cannot be observed twelve months after it — which is why they are well below one and nearly equal across samples.

Netting the placebos out yields a range, not a number. Against the no-change states, employment nets to -2.0 percentage points (s.e. 2.6); against the pre-period placebo, +0.1 (s.e. 4.2) — the summed-variance standard errors treat the arms as independent, which for the pre-period arm, estimated on the same states, is conservative. At the point estimates, the placebos reproduce between roughly 60 and 100 percent of the raw pooled gap — the 100 percent endpoint resting on the thin 13-event arm, the matched-cohort comparison landing near 60 — and both netted intervals are wide enough to admit zero and to admit losses as large as the raw -5.2. What the placebos rule out is treating the raw within-person gap as a policy effect; what they cannot do is pin the corrected effect down. Nor is the correction innocent across margins: earnings net to -\$155 (s.e. \$143) against the no-change states; the below-150-of-needs share falls by more at placebo events than at real ones (nets of +3.8, s.e. 2.3, and +3.1, s.e. 2.5), so churn out of near-poverty is endemic to the bottom of the wage distribution rather than delivered by the policy; and the no-change placebo's wage climb of \$1.62 exceeds the real-event premium of \$0.92, so taken literally the subtraction would say the policy delivered bound workers nothing, which would sit oddly beside Section V's first stage. The resolution the paper proposes — an argument, not a theorem — is that the no-change-state placebo over-corrects the wage and earnings margins: at a \$7.25 floor, a bound worker's whole climb comes through mean reversion, while in a treating state part of the same climb is delivered by the statute, to the just-above group too, through spillovers, which compresses the measured premium. The second placebo is the criterion that grounds the argument, within the very states that later raised their minimums: the wage premium is +\$0.92 with a real step and -\$0.27 without one — the first stage is policy — while the employment gap is -5.2 with a step and -5.3 without. The second placebo is what disciplines the choice of which margins the placebo corrects.

The reinterpretation of Clemens and Wither (2019) has to be stated carefully, in both directions. Their triple difference already carried a correction of this kind — the unbound states — so the raw -5.2 here is not their estimand, and the placebo-netted estimates are the objects comparable to their 6.6 points. Read that way, the \$15-era evidence neither replicates nor refutes them: the net against the no-change states is -2.0 percentage points with an interval, -7.0 to +3.1, that contains both zero and a loss on their scale. What the two placebos add to the Zipperer-Clemens exchange is a fact about the correction itself: the churn component is -3.2 percentage points in the no-change states and -5.3 in treated states' own pre-periods, so churn is

not a constant that any single control arm can be assumed to remove — it varies across places and periods by more than the netted effect being estimated. Their own randomization exercises checked assignment, not churn: permuting which states are called treated leaves the within-person gap in place, which is why this diagnostic had to come from somewhere else. Nothing here shows their estimate was wrong for 2008-2009. What the placebos show is that an uncorrected within-person gap is not a policy effect anywhere, and that the corrected one, on the largest wave on record, is a statistical zero with wide bands.

Transfers close the person-level account. Bound workers' SNAP plus TANF receipt rises \$54 (s.e. \$24) at real events against \$7 (s.e. \$16) in the placebo — a net of \$48 (s.e. \$29). Clemens (2016) found that transfers offset almost none of displaced workers' losses in the 2008-2009 episode; here the finding is the mirror image on the gain side: nothing in observed transfer receipt claws back bound workers' wage gains, and receipt rises modestly through the transition months, consistent with the transfer system insuring workers through churn rather than taxing the policy away. Dube's (2019b) one-third phase-out offset is not visible at either the state grain or the person grain in these data.

The two designs in this paper end up agreeing. State-level paths show wages up and employment flat; the corrected worker-level paths show the same thing for exactly the workers the floor binds, with intervals wide enough that the disagreement in the literature survives inside them. What narrows is how the within-person evidence can be read: a raw gap measures churn as much as policy, and only the placebo-netted version speaks to the policy question.

IX. Conclusion

The \$15-era wave raised wages at the bottom of the distribution with an elasticity of about 0.25, and employment of low-education workers did not detectably move: the implied own-wage employment elasticity is -0.04 (s.e. 0.15), ruling out broad-group elasticities near -1, with the affected-worker question bounded only loosely by the person-level design of Section VIII. On poverty the raw monthly effect is zero, drift-adjusted paths lean negative but are fragile to the adjustment, and the two-way fixed effects estimates are mildly positive and noisy; the clearest distributional movement is the share of families below 150 percent of needs, which falls in the raw estimate and every adjusted event-study path, while the two-way fixed effects estimates on that margin sit near zero.

The paper’s two mechanical findings discipline how the existing evidence should be read, and both are statements about point estimates and timing rather than statistical exclusion. The year-grain design produces, on these very person-months, a poverty “effect” of about one percentage point that matches its own pre-treatment value — so a year-grain design anchored on a base period that averages over a dip-shaped path can assign pre-existing drift to policy, and the poverty dispute’s competing dynamic specifications are competing guesses about drift whose shape monthly data simply displays. And the within-person bound-worker contrast overstates job loss by the churn rate — the placebos here reproduce roughly 60 to 100 percent of the raw gap, depending on which placebo — so uncorrected within-person magnitudes are not policy effects, and corrected ones carry wide bands. Both cautions travel: any staggered state policy evaluated in annual data, and any design that selects workers on a transitory characteristic, inherits them.

For policy, the evidence supports a specific and bounded statement: state increases of the size actually enacted — into the \$13-to-\$17 range, in the 2018-2024 economy — raised low-wage workers’ earnings without detectable employment cost, and without visible clawback through SNAP or TANF. What the data cannot deliver is the general-equilibrium counterfactual of a universal federal floor: the Congressional Budget Office’s projections for a \$15 federal minimum span job losses that these estimates neither confirm nor rule out at that scale (CBO 2019, 2021), and nothing here licenses extrapolation to states where a \$15 floor would bind twice as deeply as anything observed in this window. That reading is consistent with the position Furman took in the 2014 debate — that zero is a reasonable estimate of the employment effect of a moderate increase — while leaving the federal-\$15 question open.

Three things would settle what this paper leaves unsettled. The 2025 and 2026 SIPP files will extend the window past the last of the observed phase-ins and thicken the late horizons where drift and lagged effects cannot yet be separated. Carrying the survey’s replicate weights through every panel estimate, rather than only the level estimates, would replace bootstrap convention with design-based standard errors. And administrative income linkage would remove the recall error that makes monthly poverty dynamics an upper bound. The monthly grain is now available to this literature; the annual fights need not continue on annual terms.

References

- Ashenfelter, Orley. 1978. "Estimating the Effect of Training Programs on Earnings." *Review of Economics and Statistics* 60 (1): 47-57.
- Autor, David H., Alan Manning, and Christopher L. Smith. 2016. "The Contribution of the Minimum Wage to US Wage Inequality over Three Decades: A Reassessment." *American Economic Journal: Applied Economics* 8 (1): 58-99.
- Bane, Mary Jo, and David T. Ellwood. 1986. "Slipping into and out of Poverty: The Dynamics of Spells." *Journal of Human Resources* 21 (1): 1-23.
- Burkhauser, Richard V., Drew McNichols, and Joseph J. Sabia. 2023. "Minimum Wages and Poverty: New Evidence from Dynamic Difference-in-Differences Estimates." NBER Working Paper 31182. Cambridge, MA: National Bureau of Economic Research. Forthcoming, *Review of Economics and Statistics*.
- Burkhauser, Richard V., and Joseph J. Sabia. 2007. "The Effectiveness of Minimum-Wage Increases in Reducing Poverty: Past, Present, and Future." *Contemporary Economic Policy* 25 (2): 262-281.
- Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225 (2): 200-230.
- Card, David, and Alan B. Krueger. 1994. "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania." *American Economic Review* 84 (4): 772-793.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and David Zentler-Munro. 2022. "Seeing Beyond the Trees: Using Machine Learning to Estimate the Impact of Minimum Wages on Labor Market Outcomes." *Journal of Labor Economics* 40 (S1): S203-S247.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs." *Quarterly Journal of Economics* 134 (3): 1405-1454.
- Clemens, Jeffrey. 2016. "Redistribution through Minimum Wage Regulation: An Analysis of Program Linkages and Budgetary Spillovers." *Tax Policy and the Economy* 30 (1): 163-189.
- _____. 2017. "The Minimum Wage and the Great Recession: A Response to Zipperer and Recapitulation of the Evidence." ESSPRI Working Paper 20171. Irvine, CA: Economic Self-Sufficiency Policy Research Institute.
- _____. 2021. "How Do Firms Respond to Minimum Wage Increases? Understanding the Relevance of Non-Employment Margins." *Journal of Economic Perspectives* 35 (1): 51-72.
- Clemens, Jeffrey, and Michael R. Strain. 2021. "The Heterogeneous Effects of Large and Small Minimum Wage Changes: Evidence over the Short and Medium Run Using a

- Pre-Analysis Plan.” NBER Working Paper 29264. Cambridge, MA: National Bureau of Economic Research.
- Clemens, Jeffrey, and Michael Wither. 2019. “The Minimum Wage and the Great Recession: Evidence of Effects on the Employment and Income Trajectories of Low-Skilled Workers.” *Journal of Public Economics* 170: 53-67.
- Congressional Budget Office. 2019. *The Effects on Employment and Family Income of Increasing the Federal Minimum Wage*. Washington, DC: Congressional Budget Office.
- Congressional Budget Office. 2021. *The Budgetary Effects of the Raise the Wage Act of 2021*. Washington, DC: Congressional Budget Office.
- Currie, Janet, and Bruce Fallick. 1996. “The Minimum Wage and the Employment of Youth: Evidence from the NLSY.” *Journal of Human Resources* 31 (2): 404-428.
- Dube, Arindrajit. 2019a. “Impacts of Minimum Wages: Review of the International Evidence.” Report prepared for HM Treasury. London: HM Treasury.
- _____. 2019b. “Minimum Wages and the Distribution of Family Incomes.” *American Economic Journal: Applied Economics* 11 (4): 268-304.
- Dube, Arindrajit, T. William Lester, and Michael Reich. 2010. “Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties.” *Review of Economics and Statistics* 92 (4): 945-964.
- Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp vom Berge. 2022. “Reallocation Effects of the Minimum Wage.” *Quarterly Journal of Economics* 137 (1): 267-328.
- Godoe, Anna, and Michael Reich. 2021. “Are Minimum Wage Effects Greater in Low-Wage Areas?” *Industrial Relations* 60 (1): 36-83.
- Goodman-Bacon, Andrew. 2021. “Difference-in-Differences with Variation in Treatment Timing.” *Journal of Econometrics* 225 (2): 254-277.
- Jardim, Ekaterina, Mark C. Long, Robert Plotnick, Emma van Inwegen, Jacob Vigdor, and Hilary Wething. 2022. “Minimum-Wage Increases and Low-Wage Employment: Evidence from Seattle.” *American Economic Journal: Economic Policy* 14 (2): 263-314.
- MaCurdy, Thomas. 2015. “How Effective Is the Minimum Wage at Supporting the Poor?” *Journal of Political Economy* 123 (2): 497-545.
- Manning, Alan. 2021. “The Elusive Employment Effect of the Minimum Wage.” *Journal of Economic Perspectives* 35 (1): 3-26.
- Meer, Jonathan, and Jeremy West. 2016. “Effects of the Minimum Wage on Employment Dynamics.” *Journal of Human Resources* 51 (2): 500-522.
- Neumark, David, and Peter Shirley. 2022. “Myth or Measurement: What Does the New Minimum Wage Research Say about Minimum Wages and Job Loss in the United

- States?” *Industrial Relations* 61 (4): 384-417.
- Neumark, David, and William Wascher. 2011. “Does a Higher Minimum Wage Enhance the Effectiveness of the Earned Income Tax Credit?” *Industrial and Labor Relations Review* 64 (4): 712-746.
- Rambachan, Ashesh, and Jonathan Roth. 2023. “A More Credible Approach to Parallel Trends.” *Review of Economic Studies* 90 (5): 2555-2591.
- Rinz, Kevin, and John Voorheis. 2018. “The Distributional Effects of Minimum Wages: Evidence from Linked Survey and Administrative Data.” CARRA Working Paper 2018-02. Washington, DC: U.S. Census Bureau.
- Roth, Jonathan. 2022. “Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends.” *American Economic Review: Insights* 4 (3): 305-322.
- Sabia, Joseph J., and Richard V. Burkhauser. 2010. “Minimum Wages and Poverty: Will a \$9.50 Federal Minimum Wage Really Help the Working Poor?” *Southern Economic Journal* 76 (3): 592-623.
- Sabia, Joseph J., and Robert B. Nielsen. 2015. “Minimum Wages, Poverty, and Material Hardship: New Evidence from the SIPP.” *Review of Economics of the Household* 13 (1): 95-134.
- U.S. Census Bureau. 2021. “Dynamics of Economic Well-Being: Poverty, 2013-2016.” *Current Population Reports P70BR-172*. Washington, DC: U.S. Census Bureau.
- Vaghul, Kavya, and Ben Zipperer. 2016. “Historical State and Sub-state Minimum Wage Data.” Washington Center for Equitable Growth Working Paper. Washington, DC: Washington Center for Equitable Growth.
- Zipperer, Ben. 2016. “Did the Minimum Wage or the Great Recession Reduce Low-Wage Employment? Comments on Clemens and Wither (2016).” Washington Center for Equitable Growth Working Paper. Washington, DC: Washington Center for Equitable Growth.