Discussion Paper Series

IZA DP No. 18887

August 2026

Coverage-Loss Risk and Labor Supply: Evidence from Medicaid Continuous Coverage

Sankar Mukhopadhyay
University of Nevada Reno
and IZA@LISER

The IZA Discussion Paper Series (ISSN: 2365-9793) ("Series") is the primary platform for disseminating research produced within the framework of the IZA@LISER Network, an unincorporated international network of labour economists coordinated by the Luxembourg Institute of Socio-Economic Research (LISER). The Series is operated by LISER, a Luxembourg public establishment (établissement public) registered with the Luxembourg Business Registers under number J57, with its registered office at 11, Porte des Sciences, 4366 Esch-sur-Alzette, Grand Duchy of Luxembourg.

Any opinions expressed in this Series are solely those of the author(s). LISER accepts no responsibility or liability for the content of the contributions published herein. LISER adheres to the European Code of Conduct for Research Integrity. Contributions published in this Series present preliminary work intended to foster academic debate. They may be revised, are not definitive, and should be cited accordingly. Copyright remains with the author(s) unless otherwise indicated.

IZA@LISER NETWORK
ISSN: 2365-9793

IZA DP No. 18887 August 2026

Coverage-Loss Risk and Labor Supply: Evidence from Medicaid Continuous Coverage*

Abstract

This paper studies how reducing the risk of income-related Medicaid disenrollment affects labor supply. The Families First Coronavirus Response Act of 2020 conditioned enhanced federal Medicaid funding on states maintaining continuous coverage, prohibiting most disenrollments regardless of income changes during the COVID-19 public health emergency. Using five biennial waves of the Panel Study of Income Dynamics (2015–2023) and a fixed pre-policy treatment definition based on Medicaid coverage in the 2019 PSID interview, we estimate event-study models with individual fixed effects. Treatment status is held fixed throughout the panel to avoid conditioning on post-policy Medicaid enrollment, which may respond to the policy itself. The baseline Medicaid group experienced a 4.6-hour (or about 14%) larger change in weekly hours worked than the comparison group ($p = 0.006$), robust across alternative specifications. Employment responses are concentrated among married spouses and partners (17.2 percentage points, $p = 0.015$), consistent with a household-level channel in which secondary-earner employment is discouraged under ordinary Medicaid rules because additional household earnings can trigger coverage loss. Among expansion-state enrollees, the probability of having taxable family income above 138 percent of the federal poverty threshold rises by 14.8 percentage points ($p = 0.015$), supporting the coverage-loss-risk interpretation.

JEL classification

I13, I18, J22

Keywords

Medicaid, continuous coverage, labor supply, income

Corresponding author

Sankar Mukhopadhyay
sankarm@unr.edu

* *Data availability*: The data used in this paper are publicly available. The code used to clean data and obtain results is available from the corresponding author.

Acknowledgment: The author thanks the Nevada Agricultural Experiment Station for their support. Microsoft Copilot was used as an AI-assisted tool for grammar, clarity, and consistency editing. The author takes full responsibility for the content, analysis, and conclusions of the paper.

1. Introduction

Public insurance programs that condition eligibility on income create a discrete penalty for earning more: when household income crosses the eligibility threshold, coverage is lost. In case of Medicaid, this creates a notch in the budget constraint, where modest earnings gains can lead to a loss of insurance (Moffitt 1992; Yelowitz 1995). The resulting implicit tax on additional earnings can discourage both work effort among employed individuals and employment entry among potential (secondary) workers (Moffitt and Wolfe 1992; Yelowitz 1995). Early evidence from Medicaid eligibility expansions confirms that relaxing the income cutoff increases labor-force participation, particularly among married women (Yelowitz 1995). More recent evidence from a large Tennessee disenrollment episode shows that loss of public insurance increases employment, with gains concentrated in jobs that carry private employer-sponsored coverage (Garthwaite, Gross, and Notowidigdo 2014). Yet the literature on Medicaid expansions to new populations finds little or no effect on employment or hours (Kaestner et al. 2017; Baicker et al. 2014).

This paper studies a related policy margin that has received less attention: whether reducing the *risk* that higher earnings will trigger Medicaid loss—without changing the formal eligibility rule or expanding coverage to new groups—affects labor supply among current enrollees. The Families First Coronavirus Response Act (FFCRA) of 2020 created a near-ideal setting to study this question. The act conditioned enhanced federal Medicaid funding on states maintaining "continuous coverage," prohibiting most income-related disenrollments for the duration of the COVID-19 public health emergency. Medicaid enrollment surged from 64.5

million in February 2020 to 73.2 million in December 2020 to 79.6 million in December 2021 and peaked at 87.4 million in April 2023[2].

The policy temporarily severed the link between current earnings and the risk of losing Medicaid coverage without raising the statutory income threshold, without expanding eligibility to previously ineligible groups, and while leaving all other program features unchanged. Therefore, it represents a natural experiment in which a key source of coverage-loss risk was temporarily reduced for existing enrollees. We study labor supply responses to the continuous coverage provision using nationally representative individual-level panel data from the Panel Study of Income Dynamics (PSID), spanning the 2015–2023 survey waves. We focus on working-age individuals with high school or less education. Treatment status is defined by Medicaid coverage status during the 2019 PSID interview, the last survey round conducted before the continuous coverage went into effect. This treatment status is held fixed throughout the panel. Individuals covered by Medicaid at that baseline wave form the treatment group; individuals not covered form the comparison group. This design avoids conditioning treatment on post-policy Medicaid enrollment, which can itself respond to changes in labor supply, income, or the policy. The resulting estimates are interpreted as intent-to-treat effects of pre-policy Medicaid exposure. Event-study specifications with individual fixed effects identify the estimates from within-person changes over time, so that time-invariant differences between the two groups do not drive the results. The key identifying assumption is that, conditional on the fixed effects and time-varying controls, the two groups would have followed similar trends in

[2] https://www.kff.org/medicaid/medicaid-enrollment-and-unwinding-tracker/

labor market outcomes absent the policy. Pre-policy interaction coefficients support this assumption.

The paper yields three main findings, organized around two conceptually distinct coverage-loss-risk channels that correspond to the predictions derived in Section 2. The first finding concerns weekly hours worked. The baseline Medicaid group experienced a 4.6-hour (about 14 percent) larger 2018-to-2022 change in weekly hours than the comparison group (p-value 0.006). This is our primary result. The effect survives when restricting to observations with positive hours (3.1 additional hours, p-value 0.027), indicating that the response is not driven solely by movement from zero to positive work. Robustness checks using non-winsorized hours, state-by-year fixed effects, exclusion of post-2020 Medicaid expansion states, exclusion of essential workers, and controls for differential post-pandemic recovery by baseline household composition all produce similar estimates, ranging from 4.23 to 4.59 hours. Appendix Table A1 confirms robustness to alternative pre-policy treatment definitions.

The second finding concerns employment. The baseline Medicaid group experienced a 7.0 percentage-point larger 2018-to-2022 change in current employment than the comparison group (p-value 0.023). Decomposing by household structure reveals that the employment response is concentrated among married spouses or partners (17.2 percentage points, p-value 0.015) rather than married reference persons (6.5 percentage points, not significant). This pattern is consistent with a household-level secondary-earner channel: because Medicaid eligibility is based on household income and family size under modified adjusted gross income rules, a secondary earner's entry into employment can push household income above the eligibility threshold and trigger coverage loss. Continuous coverage suspended that risk, lowering the effective cost of secondary-earner employment in married households. The hours response and the employment response are concentrated in different subgroups: hours responses are larger among non-married individuals (6.0 hours, p-value 0.005), who face coverage-loss risk individually, while employment responses concentrate among married spouses or partners, who would otherwise risk the household's coverage by entering the labor market.

The third finding provides mechanism evidence. Among baseline Medicaid enrollees in states that had expanded Medicaid under the Affordable Care Act, the probability of having taxable family income above 138 percent of the federal poverty threshold rises by 14.8 percentage points by 2022 (p-value 0.015), with no evidence of differential pre-trends (pre-trend p-value 0.333). The threshold-crossing analysis is restricted to expansion states because 138 percent FPL is the relevant eligibility limit only where the ACA adult expansion applies. This result is consistent with the interpretation that, once the insurance penalty for crossing the income cutoff was suspended, individuals became more willing to accept earnings that would previously have jeopardized their coverage. Because annual survey income may differ from the income used in administrative eligibility determinations, this result is interpreted as suggestive mechanism evidence rather than a direct test of administrative eligibility status.

Together, the three findings support a coherent interpretation rooted in coverage-loss risk. The continuous coverage provision increased labor supply among baseline Medicaid enrollees, with the response operating through an individual-level hours channel for non-married workers and a household-level secondary-earner employment channel for married spouses and partners. Mechanism evidence from threshold crossing supports the view that individuals were more willing to accept higher earnings once the coverage penalty was suspended.

This paper contributes to three strands of the literature. First, it adds to studies of how the *design* of public insurance programs—rather than their generosity or coverage levels—shapes labor supply (Moffitt 1992; Moffitt and Wolfe 1992; Yelowitz 1995; Garthwaite, Gross, and Notowidigdo 2014). The continuous coverage provision is especially useful for this purpose because it changed the *behavioral consequence* of higher earnings without altering the formal eligibility rule or program scope. Second, the paper contributes to the literature on household labor supply and insurance effects (Moffitt and Wolfe 1992; Currie and Madrian 1999). The household-role heterogeneity in the employment results provides new evidence that household-level eligibility rules affect secondary-earner labor supply. Third, the paper contributes to the growing literature on Medicaid continuous coverage and eligibility redeterminations (Daw et al. 2024; Corallo et al. 2021; Sugar et al. 2021). The results suggest that reducing coverage-loss risk through continuous-eligibility policies can increase labor supply on both the hours and employment margins, with implications for ongoing policy debates about the design of eligibility redetermination procedures for nonelderly adults.

The remainder of the paper is organized as follows. Section 2 presents the conceptual framework. Section 3 describes the data and sample. Section 4 presents the empirical specification, baseline results, robustness checks, heterogeneity analysis by marital status and household role, and mechanism evidence from threshold crossing. Section 5 concludes.

2. Conceptual Framework

This section adapts the standard static model of labor supply under means-tested transfers to Medicaid's continuous coverage provision. It follows frameworks developed by Moffitt and Wolfe (1992) and Yelowitz (1995). In these models, public insurance affects labor supply because its value enters the household budget constraint and eligibility is linked to earnings or family income. The discrete loss of Medicaid at an income cutoff creates a nonconvexity that can discourage employment or additional work (Moffitt 1992, 2002; Moffitt and Wolfe 1992; Yelowitz 1995). Moffitt and Wolfe (1992) also emphasize that the value of Medicaid varies across families, implying heterogeneous labor-supply responses to changes in insurance eligibility.

Our adaptation differs from Yelowitz (1995) in the policy margin being studied. Yelowitz considers eligibility expansions that allowed mothers to increase earnings without losing public insurance for their children. The continuous coverage provision did not raise the statutory income threshold or extend eligibility to a new population. Instead, it temporarily prevented income-related disenrollment among individuals who were already enrolled. The relevant treatment is therefore a reduction in the immediate risk that higher household earnings would lead to Medicaid loss. We refer to this as a reduction in *coverage-loss risk*.

2.1 Individual labor supply and the Medicaid notch

Consider an individual who chooses employment $e \in \{0,1\}$ and, conditional on employment, hours $h \geq 0$. Earnings are weh, where w is the wage. Let y denote nonlabor income, F the fixed cost of employment, and $M > 0$ the individual's valuation of Medicaid coverage. Preferences are represented by

$$U(c, h, e) = u(c) - v(h) - Fe \qquad (1)$$

We assume $u' > 0$, $u'' < 0$, $v' > 0$, and $v'' > 0$. Under ordinary Medicaid eligibility rules, consumption is

$$c = y + weh - Fe + M1\{y + weh \leq \bar{Y}\} \qquad (2)$$

Here, $\bar{Y}$ is the relevant income-eligibility threshold. Equation (2) is a version of the Medicaid-notch formulation in Yelowitz (1995). At $\bar{Y}$, the individual loses a benefit valued at M, producing a discrete downward shift in available resources. If the unconstrained choice of hours places earnings just above $\bar{Y}$, the individual may reduce hours sufficiently to remain eligible. The maximum hours consistent with Medicaid eligibility are

$$\tilde{h} = \frac{\bar{Y} - y}{w}.$$

An individual whose unconstrained optimum h^* exceeds $\tilde{h}$ may nevertheless choose $\tilde{h}$ when

$$u(y + w\tilde{h} + M) - v(\tilde{h}) > u(y + wh^*) - v(h^*).$$

The discrete insurance loss can therefore reduce hours even when employment is unchanged. This is the intensive-margin implication of the Medicaid notch developed by Yelowitz (1995).

The same discontinuity can affect employment. An individual who enters work receives earnings but incurs the fixed cost F and may lose Medicaid if earnings place income above $\bar{Y}$. Employment is chosen only if

$$u(y + wh - F) - v(h) > u(y + M).$$

Compared with a setting in which Medicaid is retained, the loss of M reduces the net return to employment. Thus, income-linked Medicaid eligibility can affect both the decision to work, and the number of hours worked. This distinction is important empirically because the paper estimates hours including zeros, hours conditional on positive work, and current employment separately.

During continuous coverage, Medicaid enrollment was maintained despite subsequent increases in income. For a baseline enrollee, the relevant budget constraint becomes

$$c^{CC} = y + weh - Fe + M.$$

Continuous coverage removes the discrete insurance loss from the immediate labor-supply decision. It does not necessarily cause every enrollee to increase labor supply. Individuals well below the threshold, individuals whose unconstrained choices are already below the threshold, and individuals who place little value on Medicaid may not respond. The average effect depends on both the fraction of enrollees for whom the constraint was relevant and the magnitude of their responses, consistent with the heterogeneous valuation emphasized by Moffitt and Wolfe (1992).

Prediction 1. Continuous coverage increases labor supply among baseline Medicaid enrollees whose preferred employment status or hours would otherwise expose them to Medicaid loss. This response may occur through entry into employment, higher hours among employed workers, or both.

The three outcomes in Table 2 correspond to different components of this prediction. Hours including zeros capture changes arising from both employment and hours. Hours conditional on positive work isolate responses among observations with positive hours. Employment captures the extensive margin directly.

2.2 Household income and secondary-earner employment

Medicaid eligibility under modified adjusted gross income rules is determined using household income and family size rather than the earnings of a single person. The individual model therefore does not fully capture the incentives facing married households.

Consider a two-adult household that chooses employment e_j and hours h_j for adults $j \in \{1,2\}$. Household earnings are

$$E = w_1 e_1 h_1 + w_2 e_2 h_2.$$

Consumption under ordinary eligibility rules is

$$C = y + E - F_1 e_1 - F_2 e_2 + M_H \mathbf{1}\{y + E \leq \bar{Y}\},$$

where M_H is the household value of Medicaid coverage. As in Moffitt and Wolfe (1992), the value of public insurance can differ across households.

Suppose adult 1 is the primary earner and adult 2 is a potential secondary earner. The primary earner's income may place the household near $\bar{Y}$ before the secondary earner makes an employment decision. Entry by the secondary earner can then trigger Medicaid loss for one or more household members. The discrete loss of household coverage can discourage secondary-earner employment even when the secondary earner's wage would otherwise compensate for the fixed cost of work.

During continuous coverage, household consumption becomes

$$C^{CC} = y + E - F_1 e_1 - F_2 e_2 + M_H.$$

The secondary earner can enter employment or increase earnings without causing an immediate loss of Medicaid. This mechanism is related to, but distinct from, conventional employment lock. Employment lock ordinarily describes working or choosing a particular job to obtain employer-sponsored insurance. Garthwaite, Gross, and Notowidigdo (2014), for example, find that loss of public insurance increased employment, particularly in jobs offering employer-provided coverage. In the present setting, the household may instead avoid secondary-earner employment to preserve public insurance. Continuous coverage relaxes that household constraint.

> **Prediction 2.** If household coverage-loss risk constrains the labor supply of potential secondary earners, continuous coverage will produce a larger employment response among secondary earners than among primary earners in married households.

The theory does not imply that every labor-supply effect must be larger for married or non-married individuals. Two countervailing forces operate. Non-married individuals cannot rely on a spouse's insurance and may therefore value Medicaid more highly, strengthening their hours response. In married households, potential access to employer-sponsored insurance through a spouse may reduce the value of Medicaid, but aggregation of both spouses' earnings can make secondary-earner employment especially consequential for Medicaid eligibility. The overall married–non-married difference is therefore theoretically indeterminate. The more informative predictions concern the margin and household role of the response: hours may respond among non-married workers, while secondary-earner employment may respond among married spouses or partners.

2.3 Income-threshold crossing

The model also yields a prediction for household income. Under ordinary rules, some individuals or households choose employment or hours so that income remains below $\bar{Y}$. Continuous coverage removes the immediate Medicaid loss associated with crossing the threshold. Consequently, among households for which the threshold was binding, income should become more likely to exceed the cutoff.

In states that expanded Medicaid under the Affordable Care Act, 138 percent of the federal poverty level is the relevant income threshold for the adult expansion group. The threshold-crossing analysis is therefore restricted to expansion states.

> **Prediction 3.** Among individuals covered by Medicaid at baseline in expansion states, continuous coverage increases the probability that taxable family income exceeds 138 percent of the federal poverty threshold.

This prediction is closely related to the logic in Yelowitz (1995): permitting additional earnings without loss of public insurance should make individuals more willing to choose labor-supply positions above the former eligibility cutoff. In the present setting, however, the policy holds formal eligibility rules fixed and temporarily suspends income-related disenrollment.

2.4 Alternative channels and interpretation

The framework emphasizes coverage-loss risk, but Medicaid can affect labor supply through other channels. More stable insurance may improve access to care, facilitate management of health conditions, reduce financial risk, or reduce the administrative costs associated with eligibility renewals (Baicker et al., 2014; Mukhopadhyay, 2022; Daw et al., 2024). The broader health-insurance literature finds that health and access to insurance can affect employment, hours, labor-force participation, job choice, and job mobility (Currie and Madrian 1999). These channels could increase labor supply even for individuals who were not positioned near the Medicaid income threshold.

Our empirical analysis does not separately identify every possible pathway. The evidence is instead interpreted by considering the joint pattern of results. A positive employment or hours effect alone need not establish that the Medicaid notch was responsible. However, three findings are particularly consistent with the coverage-loss-risk framework: a positive hours response among non-married individuals, an employment response concentrated among married spouses or partners, and an increased probability of crossing the 138 percent FPL threshold in expansion states. The results should therefore be interpreted as consistent with the theoretical mechanism outlined above, rather than as proof that the Medicaid notch is the only possible mechanism.

3. Data

We use data from the Panel Study of Income Dynamics (PSID). The PSID is a long-running, longitudinal study of a representative sample of U.S. families and individuals. The PSID interviewed respondents annually from 1968 to 1997 and biannually thereafter. We use the five most recent rounds of data (from 2015 to 2023). Since our focus is on labor market outcomes, we restrict our attention to working-age adults (between the ages of 19 and 64) with high school or less education, and to reference persons and their spouses or partners (if present). We restrict the analysis to reference persons and spouses or partners because hours worked are reported only for these household members. The education restriction focuses the analysis on a population with substantial exposure to Medicaid and limits comparisons across workers with markedly different labor-market opportunities. It also reduces potential confounding from the relatively strong post-pandemic wage growth experienced by less-educated workers compared with more-educated workers. Because income may itself respond to the policy through changes in employment and hours, we do not use post-period income to define the analytical sample.

For the main hours analysis, we further restrict the sample to individuals in the labor force (currently working, temporarily laid off, or unemployed). Treatment status is defined using Medicaid coverage in the 2019 survey wave, the last survey round conducted before the continuous coverage provision took effect, and is held fixed throughout the panel. This approach avoids defining treatment using post-policy Medicaid enrollment, which may itself respond to labor supply, income, or the continuous coverage policy. After imposing these restrictions and requiring non-missing treatment status, the main hours analysis sample contains 9,590 person-year observations from 2,621 individuals. Employment models use the same age, education, household-role, and treatment restrictions but are not conditioned on labor-force participation.

PSID respondents report hours worked and family income for the prior calendar year. Consequently, the 2015 interview provides hours and income for 2014, the 2017 interview provides outcomes for 2016, and so forth. We therefore index labor-market and income outcomes by *work year*, while Medicaid coverage and demographic characteristics are indexed by interview year. In particular, treatment is based on Medicaid coverage reported at the 2019 interview, while the reference-period hours and income reported in that interview pertain to calendar year 2018.

Table 1 presents weighted characteristics of the main hours-analysis sample in the 2019 PSID interview, separately for individuals covered and not covered by Medicaid. Even though we restricted our sample to those with high school or less education, as one would expect with observational data, the respondents on Medicaid are different than those who are not covered by Medicaid. Individuals covered by Medicaid are younger, more likely to be female, less likely to be married or currently working, and have more children. They also work fewer weekly hours and have substantially lower taxable family income relative to the federal poverty threshold[3].

The Medicaid group's mean taxable family income-to-poverty ratio is 1.35, close to the 138 percent FPL Medicaid eligibility threshold applicable to non-elderly adults in expansion

[3] For hours worked and income to poverty ratio, we Winsorize these variables at 5% to reduce the effects of outliers. This means that bottom 5% of observed values are replaced by the 5th percentile value and top 5% of observed values are replaced by the 95th percentile value.

states, compared with 3.54 for the non-Medicaid group. Thirty-six percent of the Medicaid group has taxable family income above 138 percent FPL, compared with 85 percent of the comparison group. Observing some Medicaid-covered respondents above this threshold is not necessarily anomalous. Medicaid eligibility is generally determined using income and household circumstances at the time of enrollment, whereas the PSID income measure covers the full calendar year. Income fluctuations, employment changes, changes in household composition, and measurement error can therefore cause annual survey income to differ from the income used in an administrative eligibility determination. Such transitions and temporary coverage gaps, commonly described as churn, are widely documented (Corallo et al. 2021; Sugar et al. 2021).

The substantial baseline differences shown in Table 1 underscore why cross-sectional comparisons between the two groups would be inappropriate. The empirical analysis instead includes individual fixed effects and identifies the estimates from changes within individuals over time. The event-study specification also allows us to assess whether the baseline Medicaid and non-Medicaid groups followed similar outcome trends before implementation of the continuous coverage provision.

Figure 1 provides an initial descriptive comparison of weighted mean weekly hours worked by fixed baseline Medicaid status. The Medicaid group works fewer hours than the comparison group throughout the pre-policy period, but the difference between the groups is relatively stable before 2020. Hours decline for both groups in 2020, possibly driven by COVID-19 disruptions. By 2022, hours remain below their 2018 level for the comparison group, whereas hours for the baseline Medicaid group recover more fully. These unadjusted trends are consistent with a relative post-policy increase in hours among individuals covered by Medicaid at baseline, but they do not by themselves establish a causal effect. Section 4 evaluates the pattern formally using the regression-adjusted event-study specification and pre-policy trend tests.

4. Results

4.1 Empirical specification and pre-policy trends

The main estimating equation is an event-study specification of the form

$$H_{it} = \alpha_i + \lambda_t + \sum_{k \neq 2018} \gamma_k 1[t = k]\, Mcaid_{2019,i} + X_{it}\delta + \varepsilon_{it} \quad (1)$$

where H_{it} is weekly hours worked by respondent i in year t; $Mcaid_{2019,i}$ is an indicator equal to one if individual i was covered by Medicaid in the 2019 survey wave. Because PSID respondents report hours worked in the prior calendar year, the 2015 interview wave captures hours worked in 2014, and so on throughout the panel. All outcome variables are accordingly indexed to work years rather than interview years. This treatment variable is defined before the continuous coverage period and is held fixed throughout the panel. α_i are individual fixed effects, λ_t are year fixed effects, and X_{it} includes time-varying controls for age, number of children, marital status, years of education, and state of residence. The γ_k coefficients measure the difference in year-specific outcomes between the baseline Medicaid and non-Medicaid groups, relative to the 2018 reference year. All specifications are estimated with PSID individual longitudinal weights and standard errors clustered at the individual level.

The key identifying assumption is that conditional on individual and year fixed effects and the time-varying controls, individuals covered by Medicaid in the 2019 survey wave would have followed similar trends in weekly hours worked as individuals not covered by Medicaid in that baseline wave, absent the continuous coverage provision. This assumption is never directly testable, but it is supported by the pre-policy interaction coefficients.

Figure 2 presents the event-study plot for weekly hours worked. The pre-policy interactions in 2014 and 2016 are close to zero, and the joint F-test for their equality to zero yields a p-value of 0.705, providing no evidence of differential pre-policy trends. The same pattern holds across all robustness specifications; Table 3 reports the pre-trend p-value for each column. The 2020 interaction coefficient is also small and statistically insignificant. This is not surprising: the continuous coverage provision was not in place for the first three months of 2020 and was followed by a lockdown period that disrupted labor markets in ways unrelated to Medicaid incentives. We treat 2020 as a transition year and interpret the 2022 coefficient as the main policy-period estimate. The 2022 interaction is positive and statistically different from zero.

4.2 Baseline results

Table 2 presents the main results for weekly hours worked, weekly hours conditional on positive hours, and current employment status. We discuss each outcome in turn.

Weekly hours worked.

Column 1 shows the main hours estimate. This is the same outcome as Figures 1 and 2. Individuals covered by Medicaid in the 2019 survey experienced a 4.59-hour (standard error 1.66; p-value 0.006) larger change in weekly hours worked than individuals in the non-Medicaid comparison group. To interpret the magnitude, the baseline Medicaid group in the 2019 wave worked approximately 32.3 hours per week on average. A 4.59-hour larger relative change represents roughly 14 percent of that baseline. The estimate is both statistically significant at the one percent level and economically meaningful.

Weekly hours worked, positive hours only.

Because the main hours outcome includes individuals with zero reported annual hours—who are in the labor force but not currently working—Column 2 restricts the sample to observations with positive hours. The 2022 interaction remains positive and statistically significant at 3.14 hours (standard error 1.42; p-value 0.027) or about 8.4% increase, given the pre-intervention mean of the Medicaid group. This result establishes that the main hours effect is not driven solely by movement from zero to positive hours. Some portion of the effect reflects genuine increases in hours worked among individuals who were already employed.

Employment.

Column 3 shows the employment result. The pre-policy Medicaid group experienced a 7.0 percentage-point (standard error 0.031; p-value 0.023) larger change in current employment than the non-Medicaid comparison group. The joint pre-trend p-value for the employment model is 0.689, providing no evidence of differential pre-policy trends. We examine the mechanism behind this employment result in Section 4.3 below.

4.3 Robustness

Table 3 reports robustness checks for the main hours result. The estimated 2022 interaction is stable in sign, magnitude, and statistical significance across all specifications.

Column 1 repeats the baseline estimate from Table 2 (4.59 hours, $p = 0.006$). Column 2 replaces winsorized hours with non-winsorized weekly hours and finds a similar estimate (4.24 hours, $p = 0.014$). Column 3 adds state-by-year fixed effects, which flexibly absorb pandemic-era state-specific labor market shocks and variation in state-level policy responses; the estimate is 4.46 hours ($p = 0.003$). Column 4 excludes states that expanded their Medicaid programs after January 1, 2020—Utah, Idaho, Nebraska, Oklahoma, Missouri, South Dakota, and North Carolina—to address the concern that our estimates may partly reflect Medicaid expansion effects rather than continuous coverage effects; the estimate remains at 4.38 hours ($p = 0.008$). Column 5 excludes essential workers, whose labor supply during and after the pandemic may have been unusual due to legal exemptions from stay-at-home orders and heightened labor demand; the estimate is 4.39 hours ($p = 0.039$).

Columns 6 and 7 address the concern that the 2022 hours recovery among the pre-policy Medicaid group may reflect differential post-pandemic recovery among households with children rather than Medicaid incentives. Medicaid recipients have more children on average than the comparison group (1.42 vs. 0.73; Table 1), and school and childcare closures in 2020–2021 may have disproportionately depressed hours for parents, producing a mechanically larger rebound in 2022 unrelated to Medicaid coverage-loss risk. Column 6 addresses this concern by adding interactions between year indicators and the baseline number of children, measured in the 2019 survey wave. Column 7 uses an indicator for any child in the 2019 survey wave interacted with year. In both columns the 2022 treatment interaction remains positive and statistically significant. The pre-trend p-values in Columns 6 and 7 remain large (0.754 and 0.749, respectively). These results indicate that differential household-composition recovery does not account for the main hours finding.

Across all seven robustness columns, the estimated 2022 interaction ranges from 4.23 to 4.59 hours. All estimates are statistically significant at the five percent level or better, and all joint pre-trend p-values are 0.69 or higher. Appendix Table A1 additionally shows that the main hours result is robust to four alternative pre-policy treatment definitions: Medicaid coverage in both 2017 and 2019 (5.04 hours, $p < 0.05$), Medicaid coverage in either 2017 or 2019 (3.01 hours, $p < 0.05$), the main 2019 treatment with comparison-group individuals who later enter Medicaid excluded (4.62 hours, $p < 0.01$), and contemporaneous Medicaid coverage at the time of each interview as the treatment variable (3.70 hours, $p<0.05$). The consistency of the estimates across alternative identification approaches and sample restrictions supports the interpretation that the observed labor supply response reflects a genuine behavioral response to Medicaid's continuous coverage provision.

4.4 Heterogeneity and the role of household structure

Table 4 examines whether the hours and employment responses differ systematically by baseline marital status and by household role among married individuals. The table reveals a clear and interpretable pattern: hours responses are concentrated among non-married individuals, while employment responses are concentrated among married spouses and partners. This pattern is consistent with the view that continuous coverage affects labor supply through two related but distinct channels, depending on household structure.

Hours by baseline marital status.

Among individuals who were not married in the 2019 survey wave—and therefore faced Medicaid's income constraint without the option of spousal coverage—the 2022 hours interaction is 6.02 hours and statistically significant at the one percent level (standard error 2.13; p-value 0.005). The corresponding pre-trend p-value is 0.521, indicating no evidence of differential pre-policy trends in this subgroup.

Among individuals who were married in the 2019 survey wave, the 2022 hours interaction is 3.11 hours and not statistically significant (standard error 2.79; p-value 0.265). A pooled triple-interaction test comparing the married and non-married treatment effects produces a triple interaction of −2.65 hours (standard error 3.45; p-value 0.442), meaning the married–non-married difference is not statistically distinguishable at conventional levels. The pattern is nonetheless descriptively consistent with the theoretical prediction: unmarried individuals, who have no access to spousal coverage and therefore face the Medicaid eligibility constraint more directly, exhibit a larger hours response when coverage-loss risk is removed.

Employment by baseline marital status and household role.

Among non-married individuals, the 2022 employment interaction is 4.7 percentage points and not statistically significant (standard error 0.044). Among married individuals, however, the increase is 11.5 percentage points and statistically significant at the five percent level (standard error 0.049). The significant employment response among married individuals but not non-married individuals is the reverse of the hours pattern.

Splitting the married sample by household role reveals the source of the employment response. Among married reference persons, the 2022 employment interaction is 6.5 percentage points and not statistically significant (standard error 0.069). Among married spouses and partners (the group most likely to be secondary earners) the increase in employment is 17.2 percentage points and statistically significant at the five percent level (standard error 0.071; p-value 0.015).

The concentration of the employment response among married spouses and partners is consistent with a household-level coverage-loss-risk mechanism, discussed in the theoretical framework. Medicaid eligibility for non-elderly adults is based on household income and family size under modified adjusted gross income (MAGI) rules. In married households, a spouse's decision to enter employment or increase earnings can push household income above the eligibility threshold and trigger Medicaid disenrollment for the enrollee. Continuous coverage suspended this risk by prohibiting income-based disenrollment, potentially making it less costly for secondary earners to enter employment. This interpretation is supported by the mechanism analysis in Section 4.5.

Taken together, the heterogeneity results indicate that reduced Medicaid coverage-loss risk affects labor supply through two distinct channels depending on household structure. For non-married individuals, the primary response is along the hours margin, consistent with the individual-level notch or implicit-tax mechanism: once the earnings penalty associated with Medicaid loss is removed, individuals who were previously constraining their hours near the eligibility threshold work more intensively. For married spouses and partners, the primary response is along the employment margin, consistent with a household-level secondary-earner mechanism: continuous coverage makes it less risky for a secondary earner to accept employment that might push household income above the Medicaid threshold.

4.5 Mechanism: crossing the Medicaid income threshold

We next examine whether continuous coverage allowed baseline Medicaid-covered families to move into income ranges that would have been discouraged under normal Medicaid rules. Under standard Medicaid rules, individuals near the 138 percent FPL eligibility threshold face an implicit penalty: exceeding the threshold triggers loss of coverage. Continuous coverage suspended this penalty by prohibiting income-based disenrollment. If the labor supply responses reflect reduced coverage-loss risk, then the baseline Medicaid group should also become more likely to have income above the eligibility threshold during the continuous coverage period.

To test this prediction, we construct an indicator equal to one if taxable family income exceeds 138 percent of the federal poverty threshold (in expansion states), conditional on family size, and estimate the same fixed-treatment event-study specification used for hours. Before discussing the results, we should note that reported annual income in the PSID is measured with error and reflects income over the full calendar year, whereas Medicaid eligibility is determined using contemporaneous household income and circumstances. As a result, the indicator may not precisely reflect administrative eligibility status. We interpret the results as suggestive mechanism evidence rather than as a direct test of administrative Medicaid eligibility.

Figure 3 presents the event-study plot for income above the 138 percent FPL threshold. Because the 138 percent FPL threshold is the Medicaid eligibility income limit for non-elderly adults in states that expanded Medicaid under the ACA, the threshold-crossing analysis is restricted to states that had expanded Medicaid by the baseline period and had not newly expanded in 2020 or later. In non-expansion states, income thresholds for Medicaid eligibility vary by category and are typically well below 138 percent FPL, making the threshold test less directly interpretable. The 2014 and 2016 pre-policy interactions are small and not jointly significant (pre-trend p-value = 0.333), consistent with similar pre-policy trends for the two groups on this outcome. The 2020 interaction is also small and not statistically significant, consistent with the transition interpretation for that year.

The 2022 interaction is 14.8 percentage points (standard error 0.061; p-value 0.015). This is an economically meaningful and precisely estimated effect. The baseline Medicaid group became substantially more likely to have taxable family income above 138 percent of poverty by 2022, relative to the non-Medicaid comparison group and relative to their own pre-policy position. This pattern is consistent with the interpretation that, once the risk of coverage loss was suspended, baseline Medicaid-covered individuals were more willing to accept earnings or employment that crossed the eligibility threshold.

This threshold-crossing result complements the hours and employment findings in several ways. First, it directly links the labor supply responses to the Medicaid eligibility notch rather than to other pandemic-related mechanisms. If the hours and employment effects were driven solely by, say, differential wage increases or differential exposure to pandemic shocks, there is no clear reason why they should be concentrated precisely at the Medicaid income threshold. Second, it provides evidence on the mechanism for both the hours and employment channels: both are consistent with reduced income-related coverage-loss risk, whether operating through individual earnings or through household income. Third, the magnitude of the threshold-crossing effect—nearly 15 percentage points—suggests that the elasticity of income relative to the eligibility cutoff is substantial for this population once the coverage penalty is removed.

Taken as a whole, the results in Sections 4.2 through 4.5 tell a consistent and internally coherent story. Weekly hours worked increased for the baseline Medicaid group relative to the comparison group, with the effect robust to a wide range of robustness checks. Employment also increased, concentrated among married spouses and partners. These labor supply responses are accompanied by a sizable increase in the probability of crossing the Medicaid income threshold, consistent with reduced reluctance to accept higher earnings once the coverage penalty was suspended. The heterogeneity patterns are consistent with two related mechanisms—an individual-level notch response for hours among non-married individuals and a household-level secondary-earner response for employment among married spouses and partners—both rooted in reduced Medicaid coverage-loss risk.

5. Conclusion

Under the Families First Coronavirus Response Act of 2020, states were required to maintain Medicaid enrollment for most beneficiaries regardless of changes in income, temporarily severing the link between higher earnings and the risk of losing coverage. Unlike Medicaid expansions, which extend eligibility to new populations, the continuous coverage provision did not change who could enroll. It changed a behavioral constraint facing individuals who were already enrolled: the implicit tax on additional earnings that arises from income-linked disenrollment.

Using five biennial survey waves of the Panel Study of Income Dynamics (2015–2023), we estimate the effects of this provision on labor supply. The results support three main findings. First, the baseline Medicaid group experienced a 4.6-hour larger 2018-to-2022 change in weekly hours worked than the comparison group. The estimate is statistically significant at the 1 percent level (p-value 0.006) and is robust across alternative specifications. Second, the baseline Medicaid group also experienced a larger 2018-to-2022 increase in current employment. The employment response is concentrated among spouses or partners in married households (17.2 percentage points, p-value 0.015) rather than married reference persons (6.5 percentage points, not significant). This pattern is consistent with a household-level channel in which secondary-earner employment entry is discouraged under ordinary Medicaid rules because a second earner's income can push household income above the eligibility threshold. Continuous coverage suspended that risk. The hours and employment responses therefore reflect related but distinct mechanisms: hours responses are concentrated among non-married individuals, who face the eligibility constraint individually, while employment responses are concentrated among married spouses or partners, who face it through household income aggregation. Third, among baseline Medicaid enrollees in states that had expanded Medicaid under the Affordable Care Act, the probability of having taxable family income above 138 percent of the federal poverty threshold rises by 14.8 percentage points by 2022 (p-value 0.015). This result is consistent with the interpretation that, once the insurance penalty for crossing the income cutoff was suspended, individuals became more willing to accept earnings that would previously have jeopardized their coverage.

The joint pattern of results is consistent with a coverage-loss-risk framework rooted in the Medicaid notch. The conceptual framework in Section 2 builds on Moffitt and Wolfe (1992) and Yelowitz (1995) to model how income-linked Medicaid eligibility creates a discrete labor-supply distortion. Removing the risk of coverage loss—through continuous coverage—should expand the set of employment and earnings choices that individuals are willing to make.

We do not claim that the coverage-loss-risk mechanism is the only operative channel. Continuous Medicaid coverage may also have improved health, reduced financial risk, and lowered the administrative burden of repeated eligibility redeterminations, any of which could independently support labor supply (Currie and Madrian 1999). The empirical design does not separate these channels. However, two features of the evidence strengthen the coverage-loss-risk interpretation relative to a purely health-based or administrative-burden explanation. First, the employment response is concentrated precisely among married spouses or partners—the subgroup for whom household income aggregation is most likely to make employment entry consequential for Medicaid eligibility. Second, the threshold-crossing result shows that the labor-supply response is associated with income movements around the specific cutoff where Medicaid loss occurs, rather than a general shift in income or employment throughout the distribution.

The findings speak to several active policy questions about the design of Medicaid eligibility rules. The continuous coverage period ended on March 31, 2023. States then resumed eligibility redeterminations, leading to large-scale disenrollments—often described as Medicaid "unwinding"—as many enrollees were found ineligible or were removed due to procedural failures. The results suggest that reintroducing income-linked disenrollment rules may reduce labor supply through the same coverage-loss-risk channels that continuous coverage relaxed.

The household-role heterogeneity in the employment results points to a dimension of Medicaid's labor-supply effects that has received less attention in the empirical literature: the effect on secondary-earner employment decisions in married households. Under modified adjusted gross income rules, Medicaid eligibility depends on household income, so a secondary earner's entry into employment can trigger coverage loss for the household. Policy designs that reduce or eliminate this household income aggregation for secondary earners—or that allow spouses to have separate Medicaid determinations—could further increase employment among this group.

The threshold-crossing results provide direct evidence that the 138 percent FPL Medicaid eligibility threshold creates a benefit cliff that discourages some Medicaid-covered households from accepting higher incomes. When the cliff was temporarily suspended, they moved across it. This suggests that the effective marginal tax rate implied by the Medicaid benefit cliff is large enough to affect behavior for a meaningful fraction of enrollees near the threshold. Smooth benefit phase-outs, rather than discrete disenrollment at a single threshold, could reduce this distortion without requiring continuous coverage as a permanent feature of program design.

One limitation of this paper is that we study a single policy episode during an unusual macroeconomic period. The COVID-19 pandemic disrupted labor markets in ways that may have interacted with continuous coverage through channels unrelated to the Medicaid notch. The robustness checks in Table 3 address several specific concerns—differential household-composition recovery, state-specific pandemic-era shocks, and essential-worker exposure—but cannot rule out all pandemic-related confounds. Replication using administrative Medicaid data linked to earnings records from a different policy episode, such as a state-level continuous eligibility experiment, would strengthen causal inference.

Nonetheless, our findings contribute to a body of evidence showing that social-insurance *design* (not only the generosity or breadth of coverage) shapes labor market behavior. Prior work has documented that eligibility cutoffs (Yelowitz 1995), benefit withdrawal rates (Moffitt 1992, 2002), and work requirements (Sommers et al. 2020) all affect the labor-supply incentives of public insurance recipients. This paper adds evidence that the *disenrollment rules* attached to income thresholds—specifically the risk that additional earnings or employment will trigger immediate coverage loss—are themselves consequential for labor supply, operating on both the hours and employment margins and through both individual-level and household-level channels.

References

Baicker, K., Finkelstein, A., Song, J., Taubman, S., 2014. The impact of Medicaid on labor market activity and program participation: Evidence from the Oregon Health Insurance Experiment. *Am. Econ. Rev*. 104 (5), 322–328.

Corallo, B., Garfield, R., Tolbert, J., Rudowitz, R., 2021. Medicaid Enrollment Churn and Implications for Continuous Coverage Policies. KFF (Kaiser Family Foundation), December 14, 2021. https://www.kff.org/medicaid/medicaid-enrollment-churn-and-implications-for-continuous-coverage-policies/

Currie, J., Madrian, B.C., 1999. Health, health insurance and the labor market. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, Vol. 3C. Elsevier, Amsterdam, pp. 3309–3416.

Daw, J.R., MacCallum-Bridges, C.L., Kozhimannil, K.B., et al., 2024. Continuous Medicaid eligibility during the COVID-19 pandemic and postpartum coverage, health care, and outcomes. *JAMA Health Forum* 5 (3), e240004.

Garthwaite, C., Gross, T., Notowidigdo, M.J., 2014. Public health insurance, labor supply, and employment lock. *Q. J. Econ*. 129 (2), 653–696.

Kaestner, R., Garrett, B., Chen, J., Gangopadhyaya, A., Fleming, C., 2017. Effects of ACA Medicaid expansions on health insurance coverage and labor supply. *J. Policy Anal. Manag.* 36 (3), 608–642.

Moffitt, R.A., 1992. Incentive effects of the U.S. welfare system: A review. *J. Econ. Lit.* 30 (1), 1–61.

Moffitt, R.A., 2002. Welfare programs and labor supply. In: Auerbach, A.J., Feldstein, M. (Eds.), Handbook of Public Economics, Vol. 4. Elsevier, Amsterdam, pp. 2393–2430.

Moffitt, R.A., Wolfe, B.L., 1992. The effect of the Medicaid program on welfare participation and labor supply. *Rev. Econ. Stat*. 74 (4), 615–626.

Mukhopadhyay, S. (2022). The effects of Medicaid expansion on job loss induced mental distress during the COVID-19 pandemic in the US. *SSM-Population Health*, *20*, 101279.

Sommers, B.D., Chen, L., Blendon, R.J., Orav, E.J., Epstein, A.M., 2020. Medicaid work requirements in Arkansas: Two-year impacts on coverage, employment, and affordability of care. *Health Aff.* 39 (9), 1522–1530.

Sugar, S., Peters, C., De Lew, N., Sommers, B.D., 2021. Medicaid Churning and Continuity of Care: Evidence and Policy Considerations Before and After the COVID-19 Pandemic. ASPE Issue Brief HP-2021-10. U.S. Department of Health and Human Services, Washington, DC. https://aspe.hhs.gov/sites/default/files/documents/5f6e4d78d867b6691df12d1512787470/medicaid-churning-ib.pdf

Yelowitz, A.S., 1995. The Medicaid notch, labor supply, and welfare participation: Evidence from eligibility expansions. *Q. J. Econ*. 110 (4), 909–939.

Table 1. Baseline descriptive statistics by fixed pre-policy Medicaid status

Variable	**Medicaid in 2019**		**Non-Medicaid in 2019**	
	Mean	**SD**	**Mean**	**SD**
Age	39.34	11.72	44.88	12.02
Female	0.54	0.50	0.38	0.49
Married	0.40	0.49	0.62	0.48
Number of children	1.42	1.50	0.73	1.10
Years of education	11.03	2.13	11.35	1.72
Currently working	0.76	0.43	0.92	0.28
Weekly hours worked, including zeros	32.33	16.51	41.22	12.54
Taxable family income-to-poverty ratio	1.35	1.28	3.54	2.31
Taxable family income above 138% FPL	0.36	0.48	0.85	0.36
2019-wave observations	450		1,806	

Notes: The table reports weighted characteristics from the 2019 PSID interview for observations meeting the main hours-model sample restrictions in that wave. The sample includes labor-force participants aged 19–64 with a high school education or less and is restricted to reference persons and spouses or partners. Treatment group is defined by Medicaid coverage reported in the 2019 interview and is held fixed in all years. Demographic characteristics and Medicaid coverage are measured at the 2019 interview; weekly hours and taxable family income refer to the prior calendar year, 2018. Weekly hours include zero reported hours and are winsorized at the 5th and 95th percentiles. The income-to-poverty ratio equals taxable family income divided by the federal poverty threshold for the applicable family size and is winsorized at the 5th and 95th percentiles. "Above 138% FPL" is an indicator equal to one when the taxable family income-to-poverty ratio exceeds 1.38. Means and standard deviations are weighted using PSID individual longitudinal weights. Observation counts are unweighted.

Figure 1: Average number of hours worked: by treatment status

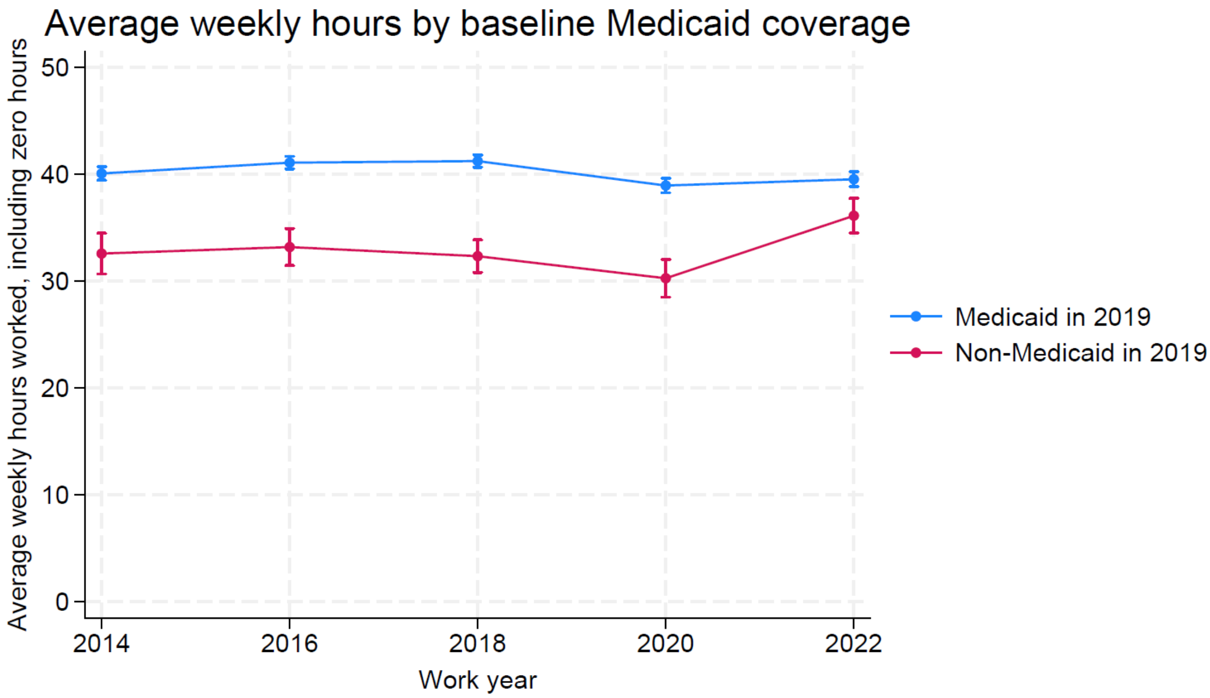

Note: Treatment group is defined by Medicaid coverage in the 2019 PSID interview (the last pre-policy wave) and held fixed in all years. 2018 is the omitted reference year. Work years are one year prior to the corresponding PSID interview year (e.g., work year 2022 is reported in the 2023 interview). 95% confidence intervals are based on standard errors clustered at the individual level.

Figure 2: Regression adjusted event study plot for hours worked

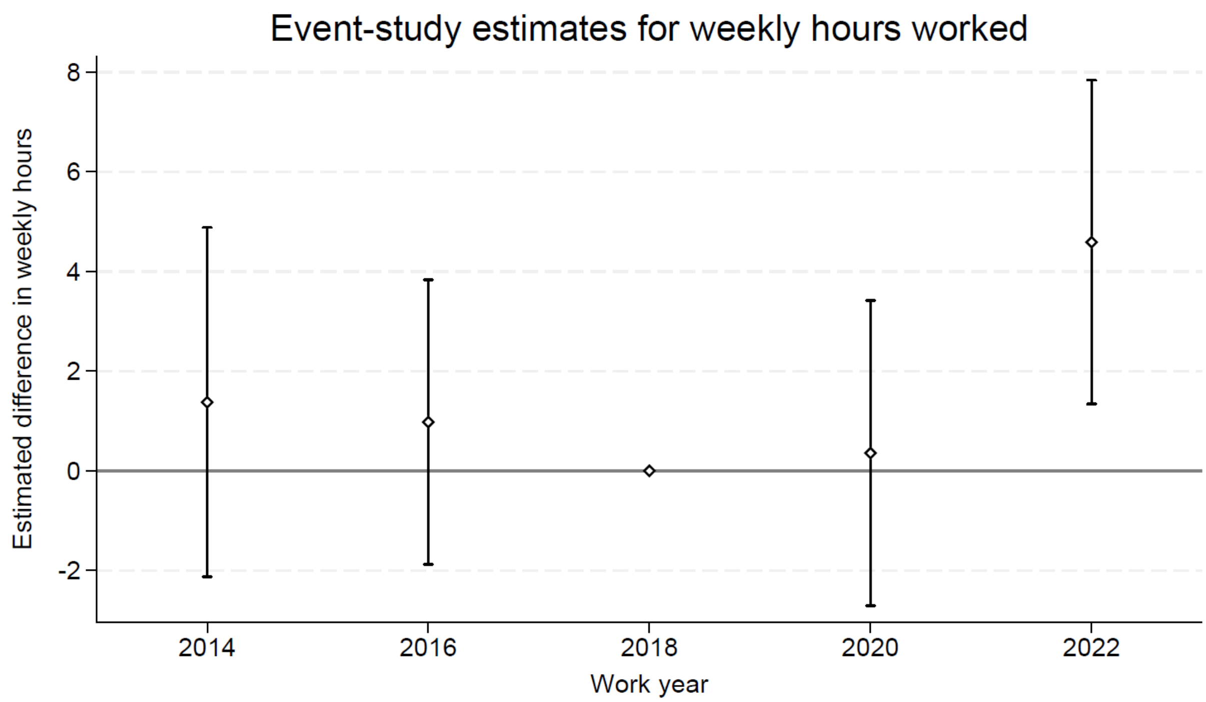

Note: Treatment group is defined by Medicaid coverage in the 2019 PSID interview (the last pre-policy wave) and held fixed in all years. 2018 is the omitted reference year. Work years are one year prior to the corresponding PSID interview year (e.g., work year 2022 is reported in the 2023 interview). All regressions include individual, state, marital-status, and education fixed effects and control for age and number of children. Estimates are weighted using PSID individual weights. 95% confidence intervals are based on standard errors clustered at the individual level.

Table 2. Effects of baseline Medicaid coverage on hours worked and employment

	(1)	(2)	(3)
	Hours incl. zeros	Hours positive only	Currently working
Baseline Medicaid x 2022	4.589***	3.138**	0.070**
	(1.658)	(1.421)	(0.031)
Baseline Medicaid group means	32.33	37.16	0.76
Pre-trend p-value	0.705	0.854	0.689
Observations	9,590	8,634	12,913

Note: Treatment group is defined by Medicaid coverage in the 2019 PSID interview (the last pre-policy wave) and held fixed in all years. For hours worked (columns 1 and 2), the reported coefficient is the interaction between baseline Medicaid coverage and the 2022 work year; 2018 is the omitted reference year. Work years are one year prior to the corresponding PSID interview year (e.g., work year 2022 is reported in the 2023 interview). Column 1 includes zero reported hours; Column 2 restricts the sample to observations with positive hours. The employment model is not conditioned on labor-force participation. All regressions include individual, state, marital-status, and education fixed effects and control for age and number of children. Estimates are weighted using PSID individual weights. Standard errors are clustered at the individual level. The pre-trend p-value reports a joint test that the 2014 and 2016 treatment interactions equal zero.

*** p<0.01, ** p<0.05, * p<0.10.

Table 3. Robustness checks for weekly hours worked

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Base	No Winsor	State by year FE	No post-2020 expansion	No essential worker	Children by year	Any child by year
Baseline Medicaid x 2022	4.589***	4.244**	4.458***	4.376**	4.385**	4.243***	4.232***
	(1.658)	(1.910)	(1.549)	(1.790)	(1.882)	(1.644)	(1.634)
Pre-trend p-value	0.705	0.694	0.788	0.797	0.769	0.738	0.766
Observations	9,590	9,590	9,590	8,256	6,400	9,590	9,590

Note: Treatment group is defined by Medicaid coverage in the 2019 PSID interview (the last pre-policy wave) and held fixed in all years. The reported coefficient is the interaction between baseline Medicaid coverage and the 2022 work year; 2018 is the omitted reference year. Work years are one year prior to the corresponding PSID interview year (e.g., work year 2022 is reported in the 2023 interview). The base specification includes zero reported hours and uses winsorized hours. Column 2 uses non-winsorized hours. Column 3 includes state by year FE. Column 4 excludes states that expanded Medicaid during or after 2020. Column 5 excludes essential workers. Column 6 includes interactions between year indicators and the baseline number of children and Column 7 includes interactions between year indicators and the baseline presence of any children. All regressions include individual, state, marital-status, and education fixed effects and control for age and number of children, except where state-by-year fixed effects are added. Estimates are weighted using PSID individual weights. Standard errors are clustered at the individual level. The pre-trend p-value reports a joint test that the 2014 and 2016 treatment interactions equal zero.

*** p<0.01, ** p<0.05, * p<0.10.

Table 4. Heterogeneity by baseline marital status and household role

	(1) Hours: non-married	(2) Hours: married	(3) Working: non-married	(4) Working: married	(5) Working: married reference person	(6) Working: married spouse or partner
Baseline Medicaid x 2022	6.017***	3.106	0.047	0.115**	0.065	0.172**
	(2.134)	(2.787)	(0.044)	(0.049)	(0.069)	(0.071)
Pre-trend p-value	0.521	0.252	0.641	0.318	0.677	0.580
Observations	4,910	4,678	6,798	6,115	3,643	2,472

Note: Treatment group is defined by Medicaid coverage in the 2019 PSID interview (the last pre-policy wave) and held fixed in all years. The reported coefficient is the interaction between baseline Medicaid coverage and the 2022 work year; 2018 is the omitted reference year. Work years are one year prior to the corresponding PSID interview year (e.g., work year 2022 is reported in the 2023 interview). Column 1 reports hours estimates for individuals not married in the 2019 PSID interview; Column 2 reports hours estimates for individuals married in the 2019 PSID interview. Column 3 reports employment estimates for non-married individuals; Column 4 reports employment estimates for married individuals; Column 5 restricts the married subsample to individuals who were reference persons (household heads) in the 2019 PSID interview; Column 6 restricts the married subsample to individuals who were spouses or partners in the 2019 PSID interview. All employment models are restricted to reference persons and spouses or partners, for consistency with the hours models in which annual hours are reported only for these household members. Sum of Columns 1 and 2 observations do not sum exactly to the full hours baseline in Table 2 because two observations that are non-singletons in the full sample become singletons when the sample is restricted to non-married individuals.

Figure 3: Regression adjusted event study plot for change in percentage of people with Medicaid with income above 138% of FPL in Expansion states

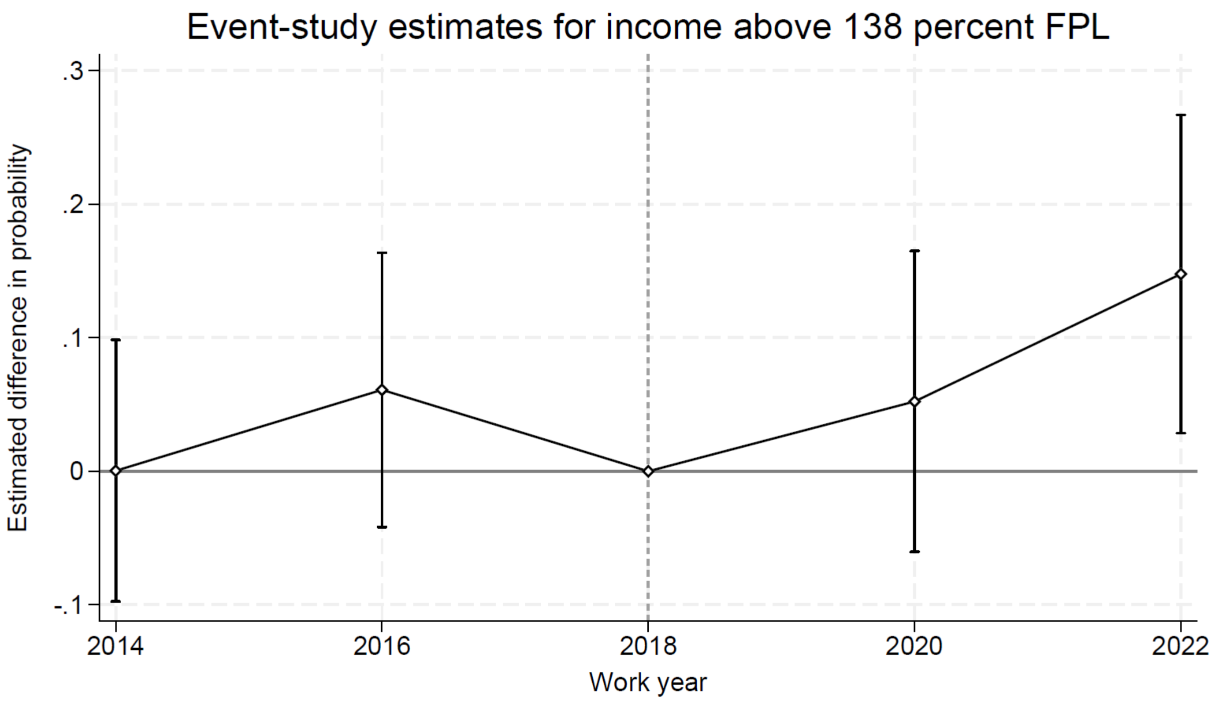

Note: Treatment group is defined by Medicaid coverage in the 2019 PSID interview (the last pre-policy wave) and held fixed in all years. The analysis is restricted to states that had expanded Medicaid by the baseline period and had not newly expanded in 2020 or later, because the 138 percent FPL threshold is the binding eligibility limit only in expansion states. The outcome is an indicator equal to one if taxable family income exceeds 138 percent of the federal poverty threshold. 2018 is the omitted reference year, shown as zero. Work years are one year prior to the corresponding PSID interview year. Regressions include individual, state, marital-status, and education fixed effects and control for age and number of children. Estimates are weighted using PSID individual weights. 95 percent confidence intervals are based on standard errors clustered at the individual level.

Appendix Table A1. Robustness to alternative treatment definitions

	(1) Main: 2019 Medicaid	(2) Medicaid in 2017 and 2019	(3) Medicaid in 2017 or 2019	(4) Exclude control entrants	(5) Medicaid at the time of interview
Baseline Medicaid x 2022	4.589***	5.044**	3.009**	4.623***	3.699**
	(1.658)	(2.136)	(1.461)	(1.660)	(1.507)
Pre-trend p-value	0.705	0.840	0.806	0.742	0.628
Observations	9590	7811	9135	8416	9570

Note: The dependent variable is weekly hours worked, including zero reported hours. The reported coefficient is the interaction between the alternative treatment definition and the 2022 work year; 2018 is the omitted reference year. Column 1 is the preferred specification, defining treatment by Medicaid coverage reported in the 2019 PSID interview, the last survey round before the continuous coverage provision took effect, and holding this classification fixed throughout the panel. Column 2 defines treatment using stable pre-policy Medicaid status in both 2017 and 2019 and excludes pre-policy switchers. Column 3 defines treatment as Medicaid coverage in either 2017 or 2019; the control group is non-Medicaid in both 2017 and 2019. Column 4 uses the main 2019 treatment definition but excludes baseline non-Medicaid individuals who enter Medicaid after 2019. Column 5 uses contemporaneous Medicaid coverage at the time of each interview as the treatment variable; unlike Column 1, treatment status in Column 5 can change over time for the same individual. Because post-policy Medicaid enrollment may respond to labor supply, income, or the continuous coverage policy, Column 5 is less preferred but is included to show the original specification. All regressions include individual, state, marital-status, and education fixed effects and control for age and number of children. Estimates are weighted using PSID individual weights. Standard errors are clustered at the individual level. The pre-trend p-value reports a joint test that the 2014 and 2016 treatment interactions equal zero.
*** $p<0.01$, ** $p<0.05$, * $p<0.10$.