Estimating the Consumption-Smoothing Value of Medicaid Expansion: Evidence and Limits*

Anita Mukherjee Daniel W. Sacks Hoyoung Yoo

May 26, 2026

Abstract

We ask what direct consumption data can reveal about the consumption-smoothing component of the insurance value of the Affordable Care Act Medicaid expansions. Medicaid expansion may reduce consumption risk by limiting exposure to out-of-pocket medical spending, but this channel may be muted when much uninsured care is financed through charity care, bad debt, or other forms of uncompensated care. We combine household expenditure data from the Consumer Expenditure Survey with the staggered timing of state Medicaid expansions to estimate effects on average consumption and on the within-education-group consumption distribution. We then map these estimated distributional effects into a static expected-utility framework. The exercise highlights both what this approach can reveal and where it is limited. Point estimates are largest for individuals without a high school degree and close to zero for higher-education groups, but the estimates are imprecise, including for average consumption. The resulting insurance-value calculations require additional assumptions and are also statistically imprecise. We therefore interpret the results as showing that existing consumption data can be used to bound the consumption-smoothing channel, but are not adequate to precisely estimate this component of the insurance value of Medicaid expansion. The findings clarify what can and cannot be learned from consumption data about this component of Medicaid's value.

Keywords: Medicaid expansion, insurance, quantile difference-in-differences, changes-in-changes, risk premium

JEL codes: I13, H51, G52

*Mukherjee: Wisconsin School of Business, University of Wisconsin – Madison, anita.mukherjee@wisc.edu. Sacks: Wisconsin School of Business, University of Wisconsin – Madison, dan.sacks@wisc.edu. Yoo: Department of Economics, University of Illinois – Urbana-Champaign, hoyoung@illinois.edu. We gratefully acknowledge funding from the Center for Demography of Health and Aging (CDHA) Rapid Pilot Program at the University of Wisconsin – Madison. Andrea Liu provided helpful research assistance. We are also grateful to Hessam Bavafa, Jonathan Kolstad, Connacher Murphy, Matt Notowidigdo, Justin Sydnor, and seminar participants at the Health Economics Initiative 2024 Annual Conference, the 2024 Risk Theory Society Annual Meeting, the Household Finance Seminar, Center for Financial Studies, University of Wisconsin-Madison, and the Finance Department Seminar, Tippie School of Business, University of Iowa. All errors are our own.

1 Introduction

The U.S. health care landscape is distinctive in its financing structure: patients face high prices for care, substantial cost sharing, and high rates of uninsurance. Medical debt is therefore pervasive and frequently delinquent. The Consumer Financial Protection Bureau estimates that medical debt in collections totals over $88 billion, affecting over one-fifth of the adult population. Policymakers have increasingly targeted medical debt, including proposals for large-scale forgiveness and restrictions on credit reporting (Consumer Financial Protection Bureau, 2023). In part to alleviate this financial burden, many states expanded Medicaid following the implementation of the Affordable Care Act (ACA) in 2014.

A large body of literature documents the consequences of the ACA Medicaid expansions. Coverage rose sharply for low-income and nonelderly adults (e.g., Miller and Wherry, 2017; Kaestner et al., 2017; Frean, Gruber and Sommers, 2017); utilization increased across primary, preventive, behavioral, and emergency care (e.g., Simon, Soni and Cawley, 2017; Soni et al., 2018; Meinhofer and Witman, 2018; Miller and Wherry, 2017; Nikpay et al., 2017; Duggan, Gupta and Jackson, 2022); and important health improvements, including reductions in mortality, followed (Miller et al., 2021; Wyse and Meyer, 2023; Goldin, Lurie and McCubbin, 2021). The expansions also improved household financial outcomes, reducing bankruptcy, medical debt, and delinquency (Gross and Notowidigdo, 2011; Mazumder and Miller, 2016; Hu et al., 2018; Brevoort, Grodzicki and Hackmann, 2020; Dodini, 2023).

Despite this evidence, we know relatively little about how the ACA Medicaid expansions affected household consumption or the distribution of consumption. Understanding distributional impacts is important not only because mean impacts can mask important distributional responses (Bitler, Gelbach and Hoynes, 2006), but also because part of the value of health insurance is reducing the risk of low consumption, a tail phenomenon. This makes the distribution of consumption, not just its mean, a natural object of interest. Indeed, although Levy, Buchmueller and Nikpay (2019) estimate the impact of Medicaid expansion on mean consumption, they do not consider its distributional impact. The most closely related evidence comes from Finkelstein, Hendren and Luttmer (2019) (hereafter, FHL), who estimate the value of Medicaid coverage in the Oregon Health Insurance Experiment using medical spending, out-of-pocket spending, health, and related outcomes. Their approach has the advantage of incorporating margins that are central to Medicaid's value, including medical care utilization and health. They do not however directly examine consumption distributions, and we look at the effect of expansion rather than coverage.

In this paper, we ask what we can learn about the impact of the ACA Medicaid expansions on the consumption distribution and about the implied consumption-smoothing component of the insurance value of Medicaid expansion. Using the Consumer Expenditure Survey (CEX) and the staggered adoption of Medicaid expansion across states, we estimate how Medicaid expansion affects average consumption and the distribution of household consumption. In principle, distributional responses could reflect changes in risk, or they could reflect heterogeneous effects across fixed groups. Only the former contributes directly to the consumption-smoothing component of the insurance value of Medicaid expansion. We therefore estimate distributional effects conditional on education group. These distributional impacts paint a richer picture of the effect of Medicaid expansion than mean impacts alone would.

These distributional impacts can be used to recover the consumption-smoothing component of the insurance value of Medicaid expansion under additional, strong assumptions. Intuitively, with a given utility function, if we know how the consumption distribution varies with and without Medicaid expansion, and we observe all the utility-relevant components of consumption, then we can recover the value of Medicaid expansion and decompose it into a mean component and a risk-reducing component; we call the latter the consumption smoothing component of the insurance value of Medicaid expansion. However, we require strong assumptions to ensure that all the utility-relevant components are observed. In particular we require that (i) individuals' ex ante risk is the same within education group,[1] (ii) measured consumption expenditures fully capture the dimensions of consumption relevant for the consumption-smoothing component of the insurance value of Medicaid expansion (ruling out insurance value from healthcare consumption as well as non-expenditure values like leisure or hassle costs), and (iii) there are no intertemporal consumption responses, i.e. no borrowing or savings. Combining these assumptions with a constant relative risk aversion utility function, the impact of Medicaid expansion on the consumption distribution can be used to recover a consumption-smoothing value. This modeling exercise is useful as much for revealing the strength of the assumptions required as for showing how to recover the consumption-smoothing component of the insurance value of Medicaid expansion.

Our empirical work begins by estimating the within-education-group mean and distributional effects of Medicaid expansions. Among the four education groups, individuals with less than a high school education show the largest mean effects, though the estimates are statistically imprecise. We then document how Medicaid expansion affects the consumption distribution within each education group. For non-high school graduates, the point estimates are positive throughout the distribution, with smaller effects in the lower tail and larger effects in the middle to upper part of the distribution. In contrast, the effects for high school graduates, those with some college, and college graduates or higher are close to zero. The confidence intervals are wide, however, producing substantial overall uncertainty.

We turn to interpreting these distributional results through an expected utility framework, maintaining the strong assumptions required to infer the consumption-smoothing component of the insurance value of Medicaid expansion. Our main result is that our point estimates are small but imprecise. For people with less than a high-school education, we find a consumption-smoothing value that is negative but with wide confidence intervals; for people with a high school education, the point estimate is positive, but again the confidence

[1] Others have made similar assumptions, for example Deshpande and Lockwood (2022) in the context of disability insurance.

intervals are wide. We gauge the precision of our estimates by asking what benchmarks they can rule out; they are consistent with both zero or negative values as well as values nearly equal to health benefits or program outlays (depending on the level of risk aversion). We emphasize that these estimates speak to one component of Medicaid expansion's value: the component operating through measured non-health consumption smoothing. Other channels, including improved access to care, better health, reduced mortality, reduced financial distress, Medicaid-paid medical care, leisure, and future consumption, are conceptually distinct and not part of this baseline calculation.

Overall, our analysis clarifies what can be learned from consumption data about the value of Medicaid expansion. Although standard models of insurance indicate that the impact of insurance on the consumption distribution is a central aspect of its value, we find that quantifying this component of the value of Medicaid expansion is challenging. We require strong theoretical assumptions to relate observable aspects of consumption data to the value of Medicaid expansion, and even under these assumptions, available data do not permit precise identification. Other approaches, which rely on assumptions of optimizing behavior, or more detailed measurement of costs and benefits (as in Finkelstein, Hendren and Luttmer (2019)), may be more fruitful.

These findings primarily contribute to the literature documenting the consequences of the ACA Medicaid expansion. We add to that large literature by showing distributional effects and investigating the limits of what can be learned about the consumption-smoothing component of the insurance value of Medicaid expansion. Our findings also relate to evidence on uncompensated care in U.S. health care financing. Garthwaite, Gross and Notowidigdo (2018) shows that hospitals provide more uncompensated care when they serve more uninsured patients, arguing that hospitals serve as insurers of last resort. Finkelstein, Hendren and Shepard (2019) finds that low-income Americans have willingness to pay for health insurance below their own marginal cost, consistent with uncompensated care reducing demand for formal insurance. Similarly, FHL concludes that the value of Medicaid to beneficiaries is likely below its cost, with part of the incidence accruing to providers who would otherwise provide uncompensated care. Our results are consistent with the possibility that some of the incidence or value of Medicaid expansion appears outside measured household consumption. The extent to which Medicaid improves provider finances by reimbursing otherwise uncompensated care, or the possible general equilibrium effects of Medicaid expansion on health care prices, are important but not included in our analysis.

The rest of the paper proceeds as follows. Section 2 introduces the data. Section 3 describes background and empirical strategy. Section 4 presents the mean and distributional consumption estimates. Section 5 maps those estimates into the consumption-smoothing component of the insurance value of Medicaid expansion under additional assumptions, and Section 6 concludes.

2 Data

Our analysis requires consumption data with state identifiers and national coverage in years before and after the 2014 Medicaid expansions.[2] We therefore rely on the Consumer Expenditure Survey (CEX), a long-running, nationally representative survey of household expenditures.

Our primary outcome is total consumption, which provides a comprehensive measure of household spending and allows for flexible responses to Medicaid expansion. As robustness checks, we also examine two narrower consumption measures. The first is well-measured consumption, developed in Meyer and Sullivan (2023), which aggregates expenditure categories for which under-reporting is believed to be limited and includes recurring nondurable outlays

[2] Several data sources meet some but not all of these criteria. Credit- and debit-card based measures, such as those used in Baker (2018), do not cover all households and may underrepresent the low-income populations most likely to benefit from Medicaid expansion. Household scanner data, such as the Nielsen Household Panel, provides high-quality information but focuses mainly on grocery and food expenditures. The Current Population Survey includes questions on food expenditures in its food security module, but lacks comprehensive consumption measures. The Panel Study of Income Dynamics contains rich consumption data but is too small for state-level identification.

and the flow values of major durables such as housing and vehicles.[3] The second is a flexible consumption measure that excludes categories viewed as consumption commitments, such as housing, utilities, and other nondiscretionary expenditures, to capture spending that may respond more readily to economic shocks. Across all consumption measures, we construct inflation-adjusted values using the CPI-U and adjust for household composition using the equivalence scale $(A + 0.7K)^{0.7}$ from Citro and Michael (1995), where A and K denote the number of adults and children in the household.

We construct our sample to include working-age households for whom Medicaid expansion could be relevant. We restrict the sample to people aged 22–64 and classify them into four mutually exclusive education groups: less than high school, high school, some college, and college or more.[4]

Using income and family information reported in the CEX, we estimate that 42%, 26%, 19%, and 7% of people in each respective education category would be income-eligible for the Medicaid expansions we study. The low eligibility rate at higher levels of education does not necessarily mean that Medicaid expansion has no insurance value for these groups; it is possible that it is valuable precisely because it provides some insurance against low probability events.

Nonetheless, low eligibility means that average effects are likely to be small, and partial enrollment among the eligible attenuates the effect of Medicaid expansion further. Decker, Abdus and Lipton (2022) estimates 44% to 46%, which we approximate as 50%. For this reason, in the main text, we show results for the two lower-education groups, as they remain most affected by Medicaid expansions. We provide the analogous analysis for the two higher-education groups in the Appendix. This approach assumes that unobserved heterogeneity is

[3] Our implementation of well-measured consumption follows Meyer and Sullivan (2023) except that we retain communication expenses, which they omit due to long-run changes in communication technology over the period they study.

[4] We define education groups based on the education level of the reference person (EDUC_REF) available in the CEX data. For example, "less than high school" category includes households in which the reference person's highest level of education is either nursery, kindergarten, elementary (grades 1-8), or high school (grades 9-12) without a degree.

limited within an education group.

Table 1 reports summary statistics for income and consumption across the education groups in our sample.[5] Income and consumption increase with education, with households in the lowest education group having the lowest levels of equivalence-scaled income and consumption, and those with a college degree having the highest. Across all education groups, well-measured consumption constitutes a large share of total consumption, and its main components, housing, food at home, and gasoline and motor oil, are broadly similar in relative importance. The table shows large cross-group differences in mean consumption, which we do not interpret as risk facing an individual; it also shows substantial within group variation, which we will at times interpret as individual consumption risk which Medicaid expansion might reduce.

3 Background and Empirical Strategy

3.1 Background: Medical Spending, Uncompensated Care, and Measured Consumption

Two features of the U.S. health care financing system are therefore central for our setting. First, uninsured households can generate large medical bills. Panel A of Figure 1 shows the distribution of medical charges in 2008–2013 for uninsured Americans with income below 138 percent of the poverty line, the group most directly targeted by the ACA Medicaid expansions.[6] Charges are the amount billed, not necessarily the amount paid. While most individuals have near-zero charges, the distribution has a long right tail, with 10 percent facing charges of $7,400 or more. Few uninsured, low-income people can afford to pay such

[5] Consumption in the CEX data is reported on a quarterly basis. Because our analysis is conducted at the annual level, we annualize consumption by multiplying quarterly expenditures by four.

[6] For details on the construction of the figure, see Appendix B. We use the Medical Expenditure Panel Survey (MEPS) for information on medical charges by insurance status.

charges. For example, nearly half of U.S. households report that they could not come up with $2,000 within a month if needed (Lusardi, Schneider and Tufano, 2011).

Second, few low-income uninsured households pay these charges in full. Panel A of Figure 1 also shows that the distribution of amounts paid is shifted far to the left relative to charges. The gap between charges and payments reflects uncompensated care, including charity care and unpaid bills that may become bad debt. This bad debt shows up on households' balance sheets and is reflected in the high levels of medical debt in the U.S. (Kluender et al., 2021). If this uncompensated care has little effect on household resources, then Medicaid expansion may shift the financing of care without producing large changes in measured non-health consumption. If instead unpaid medical bills lead to repayment burdens, collections, lower credit access, or other constraints, then Medicaid expansion may improve measured consumption, especially in adverse states.

How uncompensated care affects consumption is difficult to measure directly. If most uncompensated care reflects charity care or bad debt and households are not otherwise borrowing or saving, then uncompensated care should have little impact on consumption. If uncompensated care leads to unpaid medical debt in collections that lowers credit scores and constrains borrowing, then it may impact consumption. We illustrate two extreme scenarios in Panel B of Figure 1. If uncompensated care has no impact on consumption, then Medicaid expansion should have a relatively modest impact on the consumption distribution. On the other hand, if households end up financing uncompensated care with their own consumption (by slowly paying down their debt, for example), then Medicaid expansion should have a much larger impact, especially at the lower tails of the distribution. Because it is challenging to determine to what extent uncompensated care affects consumption, our approach is to directly estimate the impact of the post-ACA Medicaid expansion on the consumption distribution.

Recent evidence suggests that medical debt may pull down consumption less than expected. Duarte et al. (2025) study the removal of medical debts below $500 from consumer credit reports and find little effect on several credit-market outcomes. Kluender et al. (2025) study medical debt relief using two randomized experiments and similarly find limited average effects on a range of downstream financial and health-related outcomes. These findings do not imply that medical debt is unimportant for all households, but they do suggest that the channel from unpaid medical bills to measured non-health consumption may be weaker than a simple out-of-pocket spending view would imply. Together, the guarantee of emergency care under the Emergency Medical Treatment and Labor Act (EMTALA) and the limited downstream consequences of some unpaid medical bills provide a direct mechanism through which Medicaid expansion may shift the financing of care without generating large detectable changes in measured non-health consumption.

3.2 Empirical strategy: Recovering the counterfactual consumption distribution with a Medicaid expansion design

Our goal is to estimate reduced-form effects of Medicaid expansion on the distribution of measured household consumption. We recover these effects using data on consumption, described above, and a research design leveraging staggered adoption of Medicaid expansion across states. We implement this strategy separately within each education group.

While the ACA originally directed all states to expand their Medicaid program to cover all adults with income up to 138% of the poverty line, the Supreme Court ruled that the federal government could not force states to expand Medicaid. When the main provisions of the ACA went into effect in 2014, 21 states expanded their Medicaid programs. Five states (and Washington, DC) expanded coverage earlier than 2014, and an additional eight states did so between 2015 and 2019. As of 2019, the end of our sample period, 17 states had not yet expanded Medicaid. Our empirical strategy takes advantage of this uneven expansion, following a large literature investigating the effects of the Medicaid expansion.[7]

[7]Within this literature, there is some inconsistency on how states' expansion decisions should be classified and dated. We follow the classification in Miller, Johnson and Wherry (2021).

Figure A.1 shows the timing of state expansion adoptions. While the expansion states are spread throughout the country, the never-expanders are concentrated in the southeast and Great Plains.

Much of the literature focuses on estimating the average impact of Medicaid expansion, largely by means of DID models comparing the difference in mean outcomes between expansion states and nonexpansion states before and after the expansions. Our interest, by contrast, lies in estimating distributional impacts. To this end, we turn to two natural alternatives to DID. The first is quantile regression (in particular, quantile DID or QDID), which requires that the parallel trends assumption hold within quantiles rather than across means. However, as Athey and Imbens (2006) explain, this application has some unappealing features. They propose an alternative method, changes-in-changes (CIC). We use both approaches and find similar results.[8]

We introduce some notation to explain our approach. Let $t = 1, 2$ denote time periods (where $t = 1$ is before an expansion and $t = 2$ is after), and let $C_{i,t}$ denote the consumption of household i in period t. We assume for the moment that states (s) either expand Medicaid at the end of period 1 (denoted by $D_{s(i)} = 1$) or never do so ($D_{s(i)} = 0$). We relax this assumption below. We let $C_{i,t}(0)$ and $C_{i,t}(1)$ denote potential consumption without and with expansion.

The Medicaid expansion literature usually focuses on the average effect of expansion among expanders in the post period:

$$\tau_2 = E\left[C_{i,2}(1) - C_{i,2}(0) \middle| D_{s(i)} = 1\right].$$

While $E\left[C_{i,2}(1) \middle| D_{s(i)} = 1\right]$ is directly observed, researchers recover the missing mean poten-

[8]A third approach would be to estimate unconditional quantile regressions (Firpo, Fortin and Lemieux, 2009), which would measure how Medicaid affects the unconditional consumption distribution. This distribution includes cross-state differences in consumption levels, and so is not the object of interest for us.

tial outcome, $E\left[C_{i,2}(0)\middle|D_{s(i)}=1\right]$, by invoking the parallel trends assumption,

$$E\left[C_{i,2}(0)-C_{i,1}(0)\middle|D_{s(i)}=1\right]=E\left[C_{i,2}(0)-C_{i,1}(0)\middle|D_{s(i)}=0\right],$$

and a no-anticipation assumption,

$$E\left[C_{i,1}(0)\middle|D_{s(i)}=1\right]=E\left[C_{i,1}(1)\middle|D_{s(i)}=1\right],$$

which yields the average untreated potential outcome in the post period as a function of observed quantities:

$$E\left[C_{i,2}(0)\middle|D_{s(i)}=1\right]=E\left[C_{i,1}(1)\middle|D_{s(i)}=1\right]+E\left[C_{i,2}(0)-C_{i,1}(0)\middle|D_{s(i)}=0\right].$$

Our approach extends this logic from mean effects to the distribution of untreated potential outcomes. This distributional object is the main empirical input for the consumption-smoothing calculation in Section 5.

To estimate the distribution of untreated potential outcomes, we invoke parallel trends assumptions for quantiles rather than expectations. Specifically, let $P_{qgt}(0)$ and $P_{qgt}(1)$ be the qth quantile of consumption for expansion group g in period t and policy regime 0 (nonexpansion) or 1 (expansion). Then, for all quantiles q, we assume

$$P_{q12}(0)-P_{q11}(0)=P_{q02}(0)-P_{q01}(0). \tag{1}$$

We call this assumption *parallel trends in quantiles*. This assumption lets us recover the counterfactual untreated consumption distribution, $F_0(q)=P_{q11}(0)+P_{q02}(0)-P_{q01}(0)$.

To implement this identification approach, we estimate quantile regressions of consumption on indicators for the post-expansion period and for the state's having ever expanded Medicaid and their interaction, as in the standard DID model, hence the name *quantile* DID.

In our main specification, we include controls (indicators for education, for household size, top-coded at six, and household type, as well as age). We estimate the QDID model at a set of quantiles denoted by $p_1, \ldots, p_N$. In practice, we estimate the impacts for 20 evenly spaced quantile bins, starting at the 2.5th quantile. We perform inference via the bootstrap, described below.

The assumption of parallel trends in quantiles makes QDID a straightforward application of DID ideas to quantiles rather than means. However, this simplicity comes at the cost of some unappealing features, as Athey and Imbens (2006) explains. QDID recovers a counterfactual quantile (for $D_i = 1$) by adding the observed change in quantile q for $D_i = 0$. To justify this additive approach, we would have to assume that time effects are the same between $D_i = 0$ and $D_i = 1$, which limits our ability to leverage cross-group heterogeneity.[9] Athey and Imbens (2006) therefore develops change-in-changes, an alternative approach to extending DID to estimate distributional impacts that avoids these unappealing features of QDID.

We estimate both CIC and QDID models, which yield similar estimates. To implement CIC, we recover estimates for 20 uniformly spaced counterfactual quantiles, as in our QDID implementation. Estimating CIC requires repeatedly evaluating empirical distribution functions at each point in the support of the outcome. This is computationally demanding for continuous outcomes taking on many values, as our outcome variable does. We therefore coarsen our outcome measure slightly, rounding consumption amounts to the nearest $1. We use the CIC implementation developed by Kranker (2016), again using the bootstrap for inference.

Accounting for staggered adoption: Until now, we have assumed that all states either expanded Medicaid at the same time or not at all. In fact, states expanded at different times.

[9] Another unappealing feature of additive separability is that it is generally not invariant to alternative scalings of outcomes. If the level of consumption satisfies the (additive) assumption of parallel trends in quantiles, the log of consumption does not. Of course, DID suffers from this problem, as well (e.g., Roth and Sant'Anna (2023)).

Pooling groups with different treatment timings generally leads to bias in DID-type models (De Chaisemartin and d'Haultfoeuille, 2020; Goodman-Bacon, 2021; Sun and Abraham, 2021; Callaway and Sant'Anna, 2021), but dropping late expanders would reduce power. To account for states' staggered timing of expansion, we follow the logic of Callaway and Sant'Anna (2021). That is, we first estimate all parameters separately by expansion timing and then aggregate the estimates to obtain an overall average effect.

In the first step, we divide the data into timing groups, defined by the year when each state expanded Medicaid. Let g index these groups, and let $\hat{F}^g_m$ refer to the estimated factual and counterfactual timing group–specific distributions, where $m = 1$ denotes the expansion regime and $m = 0$ denotes the counterfactual nonexpansion regime. We estimate $\hat{F}^g_1$ directly from the data, and we recover $\hat{F}^g_0$ using QDID or CIC, with never-expanders (as of 2020) as the control group for timing group g. For all the timing groups, our data start in 2008 and end in 2019. We use this range to preserve the widest symmetric window for 2014 expanders while avoiding the complex impacts of the COVID-19 pandemic on health and consumption.

In our second step, we aggregate the timing group–specific distributions to overall averages. To do so, we weight each timing group by the share of the post-expansion population that it contains:

$$w^g = \frac{\sum_i w_{it} \cdot 1\{\text{timing} = g\} \cdot 1\{g \leq t\}}{\sum_{g'} \sum_i w_{it} \cdot 1\{\text{timing} = g'\} \cdot 1\{g' \leq t\}}, \tag{2}$$

$$\hat{F}_m = \sum_g w^g \hat{F}^g_m, \tag{3}$$

where w_{it} is the survey weight of observation i in year t. These weighted quantiles are analogous to average effect of treatment-on-the-treated estimates, but for distributions, with weights accounting for the uneven post-period lengths across timing groups.

Inference: We conduct inference via the block bootstrap, resampling states. We use the bootstrap because our distributional objects are nonlinear functions of estimated quantiles and because we are not aware of a cluster-robust variance estimator for CIC.

We implement our bootstrap as follows. First, we draw a bootstrap sample of states, resampling with replacement. Second, we estimate timing group–specific factual and counterfactual distributions for the resampled timing groups. Third, we recalculate the timing group weights within the bootstrap iteration and obtain bootstrap-specific estimates $\hat{F}_{m,b}$.[10] We repeat this procedure 1,000 times. We then report bootstrap confidence intervals as the 2.5th and 97.5th quantiles of the bootstrap estimates.

4 Effects on Consumption: Mean and Distributional Estimates

We begin our discussion of the results by showing trends in the distribution of total consumption (excluding health insurance) for the two low-education groups, less than high school and high school graduates, separately, in Figure 2. The figure plots percentiles of this indicator from 2008 to 2019 for nonexpansion states and for states that expanded in 2014.[11] Several patterns are apparent in the figure. First, even in our relatively homogeneous sample (within each low-education group), consumption is dispersed, with a 90/10 ratio of approximately 3. Second, consumption is generally higher at each percentile in expansion states than in nonexpansion states. Third, consumption fell throughout the distribution in the Great Recession, although it fell more sharply at higher percentiles. Consumption did not return to its 2008 levels until 2019.

Prior to the Medicaid expansions, at each percentile, consumption moved roughly in

[10] Our bootstrap procedure does not stratify on timing group, so some bootstrap samples exclude some timing groups. Thus, the variability across bootstrap iterations reflects, in part, variability in the timing group, and our bootstrap confidence intervals account for uncertainty in which states expanded when, as they should in a design-based approach to uncertainty.

[11] We focus on the 2014 expansion states here to avoid the problems from pooling them with later adopters. We show analogous plots for states that expanded later (2015, 2016, 2017, and 2019) in Appendix Figures A.2–A.5. Note that no states expanded in 2018.

parallel for expansion and nonexpansion states. Individuals at high percentiles of consumption experienced a sharp consumption decline from 2008 to 2009 in both expansion and nonexpansion states. In Panel (a), the figure shows some divergence in consumption across percentiles among individuals with less than a high school education after the Medicaid expansions. The median, 75th, and 90th percentiles appear to diverge somewhat, falling in nonexpansion states but holding steadier in expansion states. These patterns suggest that the Medicaid expansions increased consumption more in the middle of the distribution than in the lower tail. In Panel (b), the consumption distribution for high school graduates shows little divergence across percentiles between Medicaid expansion and nonexpansion states.

We report our key distributional estimates in Figure 3. The figure shows the factual and counterfactual consumption distributions (left-hand side) and the estimated quantile treatment effects and average effects (right-hand side), separately for people with less than a high school degree (top panel) and with high school graduates (bottom panel). We find that Medicaid expansion increases consumption by $538 on average among individuals without a high school degree (panel a), but a smaller $170 among those with a high school degree. The point estimate for the less educated group corresponds to about 3 percent of mean consumption, and for the more educated group about half a percent. These small to modest average effects are imprecisely estimated; the 95% confidence interval is (-500, 1525) for the least education group and (-1024, 1033) for high school graduates. This insignificant finding is consistent with the previous literature. For example, Levy, Buchmueller and Nikpay (2019) finds no significant effect of Medicaid expansion on non-health consumption using the CEX for 2010 through 2016. We note that because Medicaid coverage is not expanding from zero, some individuals in very low-consumption states are likely to have been eligible or enrolled even before the ACA expansions. This can attenuate the reduced-form effects of expansion, especially in parts of the consumption distribution where many individuals were already eligible for Medicaid before the ACA expansions. The larger point estimates in the middle of the distribution are consistent with marginal enrollees being more prevalent in that range.

Turning to the distributional effects, we see positive effects of Medicaid expansion across the consumption distribution for individuals without a high school degree. These effects appear larger at about the median and above than at smaller quantiles. As with the mean, however, these quantile effects are imprecisely estimated, and none of them is significantly different from zero. For individuals with a high school degree, we see quantile effects that are close to zero throughout the distribution; as with the less-than-high-school estimates, they are uniformly not statistically significant. Thus, for less educated people, we find small but noisy effects throughout the distribution, and for people with a high school degree, we find even smaller effects. Results are similar when using the changes-in-changes (CIC) approach, as shown in Appendix Figure A.6. Appendix Figure A.7 shows effects on consumption that are close to zero, both on average and across the distribution, for the higher-education groups (individuals with some college education and those with a college degree or higher).

5 Measuring the insurance value of Medicaid expansion

So far we have estimated the impact of Medicaid expansion on the distribution of consumption, finding noisy estimates throughout the education-group specific distribution. Next we ask what assumptions are necessary to map these estimates to the consumption-smoothing component of the insurance value of Medicaid expansion. We caution that this exercise shows assumptions needed to recover one component of the value of Medicaid expansion; we do not attempt to provide an estimate of the total value of Medicaid, nor of the value of Medicaid enrollment itself.

5.1 Assumptions and model

Mapping the reduced-form effects of Medicaid expansion on measured consumption to a model-based measure of insurance value requires three key assumptions, each of which is strong and debatable.

Assumption 1: We assume that the within-education-group consumption distribution represents the distribution of ex ante consumption states faced by a representative individual in that education group, i.e. it reflects consumption risk not heterogeneity. Because we do not observe the same household across all possible states of the world, this interpretation is necessarily a maintained assumption. It is most plausible when remaining within-group variation reflects risk or transitory states, such as health, employment, or medical-spending shocks, rather than fixed unobserved differences across individuals. To the extent that the within-group distribution still reflects persistent unobserved heterogeneity, the mapping from distributional consumption effects to ex ante insurance value is less direct.

Assumption 2: We assume that the consumption expenditures measured in the CEX are the only aspects of consumption relevant for the insurance value of Medicaid expansion. One way this assumption can fail is that the CEX does not capture third-party financed healthcare consumption. Medicaid expansion increases healthcare consumption; this medical benefit is surely utility-relevant and may generate insurance value. This assumption also fails if Medicaid expansion affects leisure or hassle costs in a way that is relevant for its insurance value.

Assumption 3: We assume that there are no intertemporal consumption responses - no borrowing or savings. This assumption matters because in a recent paper, Jaffe, Malani and Reif (2026) show that access to credit can reduce the incremental value of insurance by allowing households to smooth shocks over time. This reinforces that our static consumption-based mapping should be interpreted as one limited way to summarize the consumption-smoothing content of the estimates, rather than as a complete model of all smoothing technologies available to households. Savings might also arise if households anticipate future taxes to finance Medicaid expansions.

Model: Consider two policy regimes indexed by $m \in \{0, 1\}$, where $m = 0$ corresponds to a regime without Medicaid expansion and $m = 1$ to one with Medicaid expansion. The joint distribution of income, health shocks, and health care financing induces a distribution over measured non-health consumption in each regime. Let $F_m(c)$ denote the consumption distribution in policy regime m and let u denote the utility function. We summarize expected utility under regime m as

$$EU_m = \int_c u(c) dF_m(c).$$

In the model, F_m is the distribution of consumption facing a given individual. We infer F_m from the education-group-specific consumption distribution, relying on Assumption 1 to go from observed distributions to risk, and Assumption 2 to go from observed consumption expenditures (a subset of all risky consumption) to the model's comprehensive consumption concept. Assumption 3 shows up implicitly as the formulation of consumption here is static.

Given this setup, we define the willingness to pay for Medicaid expansion, γ, as the amount of consumption an individual would forgo in the expansion regime to be just indifferent between the expansion and nonexpansion regimes. That is, γ solves

$$\int_c u(c - \gamma) dF_1(c) = \int_c u(c) dF_0(c). \tag{4}$$

We then define the consumption-smoothing component of the insurance value of Medicaid expansion as the difference between γ and the expected consumption gain:

$$\pi = \gamma - \left(\int_c c dF_1(c) - \int_c c dF_0(c) \right). \tag{5}$$

The second term in Equation (5) is the difference in expected measured consumption between the two policy regimes. We call this term the mean value. In this model it captures the average change net consumption gain from Medicaid, net of taxes and behavioral responses. Changes in savings due to anticipated future tax changes are ruled out via our Assumption 3. Although our model abstracts from non-consumption gains in utility, other benefits or costs of Medicaid could be thought of as appearing in this average gain (were they measured). Our focus is not on the average gain but on the gain or loss above the man change.

As the remainder of γ is the additional willingness to pay for Medicaid expansion above its impact on mean consumption, we refer to it as the consumption-smoothing component of the insurance value of Medicaid expansion.[12] It captures the value of reallocating consumption toward states in which marginal utility is high. Under CRRA utility, low-consumption states receive greater weight in expected utility, so changes in the lower part of the consumption distribution are especially important. However, the consumption-smoothing value depends on the full change in the consumption distribution, not only on whether the lower tail increases.

We note that our procedure recovers the consumption-smoothing component of the insurance value of Medicaid *expansion* rather than eligibility or take-up. This is because F_0 and F_1 are distributions under Medicaid expansion rather than under Medicaid *take-up*, π reflects the value of expansion, not of take-up. This distinction is important because Medicaid expansion does not increase eligibility from 0 to 100 percent (many people are eligible without expansion, and many are ineligible with expansion), and because Medicaid eligibility does not imply take-up. Thus, we should expect that our estimate of π is lower than estimates of the value of Medicaid eligibility or take-up, both because we are focused on a single component and because we are measuring the value of expansion rather than eligibility or take-up.

[12]The terms "consumption smoothing component" and "insurance value" might feel redundant to some readers. We include both to emphasize that in general, outside the context of our model, insurance value might include more than consumption smoothing.

5.2 Implementation

Our QDID and CIC estimates give us factual and counterfactual consumption distributions for the post period in Medicaid expansion states. Denote these distributions $\hat{F}_1$ and $\hat{F}_0$. The estimation recovers percentiles $p_1, \ldots, p_N$. Denote the value of consumption under distribution F_m at these percentiles as c_{im}, so that, for example, c_{10} is the p_1 percentile of consumption in the nonexpansion policy regime.

To translate these percentiles to the consumption-smoothing component of insurance value, we assume that the consumption distribution has N points of support and that utility is of the constant relative risk aversion class with risk aversion ρ. We can therefore estimate expected utility under each policy regime:

$$\hat{EU}_m = \sum_{i=1}^{20} (p_i - p_{i-1}) \frac{c_{im}^{1-\rho}}{1-\rho}. \quad (6)$$

To recover $\hat{\gamma}$, we find the consumption offset that equalizes expected utility between the two policy regimes. That is, $\hat{\gamma}$ solves

$$\sum_{i=1}^{N} (p_i - p_{i-1}) \frac{(c_{i1} - \hat{\gamma})^{1-\rho}}{1-\rho} = \sum_{i=1}^{N} (p_i - p_{i-1}) \frac{c_{i0}^{1-\rho}}{1-\rho}. \quad (7)$$

Finally, to estimate this component, π, we subtract off the difference in expected consumption:

$$\hat{\pi} = \hat{\gamma} - \sum_{i=1}^{N} (p_i - p_{i-1}) (c_{i1} - c_{i0}). \quad (8)$$

In implementation, we have $N = 20$ uniformly spaced percentiles.

We report estimates of $\hat{\pi}$ for different levels of ρ. We report risk-aversion parameters of 1, 3, and 5 to show how the calculation varies with this maintained parameter and to facilitate comparison with prior work. Lower values may be more consistent with some evidence from labor-supply behavior (Chetty, 2006), while higher short-run values may be relevant when households face consumption commitments. We therefore do not take a stand on a preferred value and instead emphasize the sensitivity of the calculation to this parameter.

We compute this object separately by expansion timing group and then aggregate across timing groups using the weights defined in Section 3.2. Let $\hat{\pi}^g$ denote the timing group–specific value computed from $\hat{F}_0^g$ and $\hat{F}_1^g$. We aggregate as

$$\hat{\pi} = \sum_g w^g \hat{\pi}^g. \tag{9}$$

Inference uses the same block bootstrap described above, resampling states and recomputing the timing group–specific distributions, weights, and insurance values in each bootstrap iteration.

5.3 Results

We report these estimates and their 95% confidence intervals separately for individuals without a high school degree (Table 2) and for individuals with a high school degree (Table 3). At a coefficient of relative risk aversion of 3, the estimated consumption-smoothing component of insurance value is -\$177 per person per year for individuals without a high school degree and \$7.5 per person per year for high school graduates. Neither estimate is statistically significant.

The estimate is also sensitive to the assumed coefficient of relative risk aversion: the magnitude is less than half as large when risk aversion is 1, and moderately larger in magnitude when risk aversion is 5. For individuals with a high school degree, the estimated consumption-smoothing component of insurance value is close to zero when risk aversion is 3. Overall the point estimate is moderate for individuals without a high-school degree and small for individuals with a high school degree, but for both groups it is imprecise, with wide confidence intervals.

The pattern of small but imprecise insurance-value estimates is not sensitive to the exclusion of covariates, the consumption measure, the estimator, or the sample, as the results in Tables 2 and 3 show. Working with well-measured consumption or flexible consumption, instead of overall consumption (Panels C and D) produces similar point estimates but narrower confidence intervals, consistent with smaller noise in the two other consumption measures than the overall consumption measure. The CIC estimates in Panel E are similar to the results from QDID. Lastly, our estimates are not sensitive to the exact coding of expansion status or handling of consumption. Excluding California or Wisconsin (Panels F and G) does not materially change our estimates; nor does working with a 50-point grid of the consumption distribution (Panel H).

We emphasize two caveats in interpreting these value estimates. First, these calculations do not estimate the total value of Medicaid, nor the value of Medicaid enrollment itself. They isolate one model-based component of the value of Medicaid expansion: the value generated by changes in measured non-health consumption across different policy regimes. Other channels, including improved health, reduced mortality, lower financial distress, access to Medicaid-paid medical care, and possible effects on leisure or future consumption, are outside this calculation. Second, we estimate the effect of Medicaid expansion, averaging over people who are and are not eligible for Medicaid and who do and do not take it up. This averaging picks up the option value of expansion (i.e., that a person could take it up when needed) to the extent that the option (when exercised) shows up in consumption data. But it understates the value of Medicaid to those who actually take it up.

With these caveats in mind, we ask whether the confidence intervals evident in Tables 2 and 3, though wide, are informative enough to rule out meaningful benchmarks. We plot our estimates and 95 percent confidence intervals in Figure 4, as a function of the assumed coefficient of relative risk aversion. Panel (a) presents results for individuals without a high school degree, and Panel (b) presents results for high school graduates. We overlay several benchmark values for comparison.

Many of our benchmarks are derived from estimates of the per-Medicaid-enrollee value of

Medicaid. To make them comparable to our (per capita) estimates, we scale the benchmarks down in two steps to be comparable to our estimated insurance values. For the lowest education group, the benchmarks are scaled down first by 42% (the percent income-eligible for Medicaid, from the CEX) then by 50% (the percent estimated to enroll upon eligibility, based on Decker, Abdus and Lipton 2022). For the next education group, high school graduates, we scale first by 26% and then by the same 50%. These and other calculations for the benchmarks are further detailed in Appendix C.

We start by comparing our estimates to government outlays on the Medicaid expansion population. These estimates are from Guth et al. (2021), who reports $6,110 per beneficiary, which works out to about $1,300 per less-than-high-school educated person in our data or $794 per high school graduate. Our confidence intervals indicate that the consumption-smoothing component of the value of Medicaid expansion is below outlays for high school graduates, and also for less-than-high-school graduates, except at high levels of risk aversion. However, this benchmark is a very high target. Government outlays do not reflect the social cost of Medicaid provision (as they crowd out some care that would be financed in other ways), and they include the mean component of the value of Medicaid.

A more appropriate benchmark would be the social cost of Medicaid, which in some models is the moral hazard cost of insurance. Finkelstein, Hendren and Luttmer (2019) estimate the moral hazard cost of Medicaid coverage at $800 per enrollee. Scaling down by eligibility and take-up rates reduces this to $168 per person for those with less than high school and $104 per person for high school graduates. These levels are comparable to our point estimates for the consumption-smoothing component of the insurance value of Medicaid expansion. That is, our point estimates of the insurance value are similar to the implied moral hazard cost from Finkelstein, Hendren and Luttmer (2019). However, of course, our confidence intervals are also consistent with a wide range of insurance values.

Another estimate of the value of Medicaid comes from its mortality-reducing effects. Miller et al. (2021); Wyse and Meyer (2023) find substantial reductions in mortality from

Medicaid expansion. We scale the per-coverage-year mortality benefits in Wyse and Meyer (2023) by eligibility and take-up, and translate to a dollar benefit with an assumed value of a statistical life-year of $100,000. Figure 4 shows that the mortality benefits of Medicaid lie outside our confidence intervals except at low levels of risk aversion.

As an additional benchmark, we calculate the insurance value that would arise if uninsured individuals were responsible for the full cost of their medical care, i.e., if there were no uncompensated care. We calculate this benchmark using data from the Medical Expenditure Panel Survey (analysis detailed in Appendix B and plotted as gray lines in Figure 4). The point estimates are below this full-cost benchmark (e.g., if the risk aversion is 3, the MEPS-implied insurance value is about $436), but our confidence intervals do not rule it out.

Finally, we compare our estimates to related estimates in the literature, after rescaling to account for limited eligibility and take-up. For example, Finkelstein, Hendren and Luttmer (2019) report an estimate of approximately $30 per year for the pure-insurance component in the Oregon Health Insurance Experiment, while Shupe (2023) estimates a consumption insurance value of approximately $20 per year. Our estimates are consistent with both these values.

These comparisons reveal sharp limits to what we can learn about the consumption-smoothing component of the insurance value of Medicaid expansion. Strong assumptions are required to map consumption distribution effects into insurance values. Even under these assumptions, available consumption data and limited eligibility take-up generate wide confidence intervals on the consumption-smoothing values. Although our point estimates are small, our confidence intervals are consistent with wide ranges, including small or negative values as well as values as large as key costs and benefits of Medicaid expansion, such as the moral hazard costs or mortality reductions.

6 Conclusion

This paper investigates what we can learn about the consumption consequences of the Affordable Care Act Medicaid expansions from repeated cross-sections of consumption data and theoretical structure. The motivation for this exercise is that the mean impact of the Medicaid expansion may miss important changes in the tails of the distribution, and indeed theory predicts reducing tail risk is an important component of the value of insurance coverage and therefore Medicaid expansions.

Our analysis reveals sharp limits to the inference we can make from available data. We find imprecise estimates of the effect of Medicaid expansion on the distribution of income for people with less than a high school degree and for people with a high school degree. For people with more education, we find effects that are close to zero but still imprecise. This imprecision is a reflection of limited available data as well as incomplete eligibility and take-up.

Even setting aside the limited power to learn distributional effects, it is challenging to make inferences about the insurance value of Medicaid expansion from the effects of expansion on the distribution of consumption. To do so, we require strong assumptions. We develop the necessary framework and assumptions to use distributional impacts to recover the consumption smoothing component of the insurance value of Medicaid expansion. The assumptions include a restriction on heterogeneity (that for individuals in the same education group, the observed consumption distribution is the distribution of consumption risk), on the relevant arguments of the utility function (that the consumption smoothing impact of Medicaid expansion is captured in non-health consumption expenditures), and on intertemporal substitution (that there is no borrowing or savings). Each of these assumptions is likely dubious. This modeling exercise reveals the strength of the assumptions required for mapping distributional effects to the insurance value of Medicaid expansion. It provides a framework for interpreting distributional estimates, and a starting point for further work, but not a definitive account of the insurance value of Medicaid expansion, let alone an overall accounting of the value of Medicaid.

Indeed, our framework intentionally isolates one component of the value of Medicaid expansion exposure: the value generated by changes in measured non-health consumption across different policy regimes. It does not estimate the value of Medicaid enrollment itself, nor does it estimate the total value of Medicaid. Medicaid may generate value through improved health, reduced mortality, reduced financial distress, lower hassle costs, improved hospital finances, or changes in future consumption and saving. It may also generate direct utility benefits through increased health care utilization and the option to obtain subsidized medical care. These in-kind healthcare benefits are outside our consumption measure and outside the baseline utility calculation.

Our findings also speak to how health care is financed in the US. A growing literature emphasizes that providers function, to some extent, as insurers of last resort through uncompensated care, and that formal insurance expansions can therefore shift costs away from implicit provider-based insurance and toward public financing (Garthwaite, Gross and Notowidigdo, 2018; Finkelstein, Mahoney and Notowidigdo, 2018). In this light, the modest and imprecise consumption-smoothing estimates are consistent with the possibility that much of the value or incidence of Medicaid expansion appears outside measured household consumption.

Recent work also emphasizes that standard consumption-based approaches may understate insurance value when households face binding liquidity constraints, even in the absence of large observed consumption smoothing (Ericson, Jaspersen and Sydnor, 2025). At the same time, access to credit can substitute for insurance in smoothing shocks over time and thereby reduce the incremental value of insurance (Jaffe, Malani and Reif, 2026). Our framework does not incorporate these intertemporal smoothing mechanisms. The results should therefore be read narrowly: they show that, using existing CEX consumption data and a transparent static consumption-based model, it is difficult to find precise evidence that the ACA Medicaid expansions generated large gains through measured non-health consumption smoothing. This does not imply that Medicaid expansion had little overall value. Rather, it clarifies what this particular data-and-model combination can and cannot establish. In that sense, the paper's contribution is not to deliver a sharp estimate, but to show both why direct consumption data are a reasonable way to study this component of insurance value and why existing data and a static mapping do not take the approach very far.

References

Athey, Susan, and Guido W Imbens. 2006. "Identification and inference in nonlinear difference-in-differences models." *Econometrica*, 74(2): 431–497.

Baker, Scott R. 2018. "Debt and the response to household income shocks: Validation and application of linked financial account data." *Journal of Political Economy*, 126(4): 1504–1557.

Bitler, Marianne P, Jonah B Gelbach, and Hilary W Hoynes. 2006. "What mean impacts miss: Distributional effects of welfare reform experiments." *American Economic Review*, 96(4): 988–1012.

Brevoort, Kenneth, Daniel Grodzicki, and Martin B Hackmann. 2020. "The credit consequences of unpaid medical bills." *Journal of Public Economics*, 187: 104203.

Callaway, Brantly, and Pedro HC Sant'Anna. 2021. "Difference-in-differences with multiple time periods." *Journal of Econometrics*, 225(2): 200–230.

Chernozhukov, Victor, Iván Fernández-Val, and Blaise Melly. 2022. "Fast algorithms for the quantile regression process." *Empirical Economics*, 62(1): 7–33.

Chetty, Raj. 2006. "A new method of estimating risk aversion." *American Economic Review*, 96(5): 1821–1834.

Citro, Constance F., and Robert T. Michael, ed. 1995. *Measuring poverty: A new approach.* National Academies Press.

Consumer Financial Protection Bureau. 2022. "Medical Debt Burden in the United States." https://files.consumerfinance.gov/f/documents/cfpb_medical-debt-burden-in-the-united-states_report_2022-03.pdf.

Consumer Financial Protection Bureau. 2023. "CFPB Proposes to Ban Medical Bills from Credit Reports." Accessed: 2024-08-29.

De Chaisemartin, Clément, and Xavier d'Haultfoeuille. 2020. "Two-way fixed effects estimators with heterogeneous treatment effects." *American Economic Review*, 110(9): 2964–2996.

Decker, Sandra L, Salam Abdus, and Brandy J Lipton. 2022. "Eligibility for and enrollment in Medicaid among nonelderly adults after implementation of the Affordable Care Act." *Medical Care Research and Review*, 79(1): 125–132.

De Nardi, Mariacristina, Eric French, and John B Jones. 2010. "Why do the elderly save? The role of medical expenses." *Journal of Political Economy*, 118(1): 39–75.

Deshpande, Manasi, and Lee M Lockwood. 2022. "Beyond health: Nonhealth risk and the value of disability insurance." *Econometrica*, 90(4): 1781–1810.

Dodini, Samuel. 2023. "Insurance Subsidies, the Affordable Care Act, and Financial Stability." *Journal of Policy Analysis and Management*.

Duarte, Victor, Julia Fonseca, Divij Kohli, and Julian Reif. 2025. "The effects of deleting medical debt from consumer credit reports." National Bureau of Economic Research.

Duggan, Mark, Atul Gupta, and Emilie Jackson. 2022. "The impact of the Affordable Care Act: evidence from California's hospital sector." *American Economic Journal: Economic Policy*, 14(1): 111–151.

Ericson, Keith Marzilli, Johannes G Jaspersen, and Justin R Sydnor. 2025. "Moral Hazard under Liquidity Constraints." National Bureau of Economic Research.

Finkelstein, Amy, Nathaniel Hendren, and Erzo FP Luttmer. 2019. "The value of Medicaid: Interpreting results from the oregon health insurance experiment." *Journal of Political Economy*, 127(6): 2836–2874.

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2019. "Subsidizing health insurance for low-income adults: Evidence from Massachusetts." *American Economic Review*, 109(4): 1530–67.

Finkelstein, Amy, Neale Mahoney, and Matthew J Notowidigdo. 2018. "What does (formal) health insurance do, and for whom?" *Annual Review of Economics*, 10: 261–286.

Firpo, Sergio, Nicole M Fortin, and Thomas Lemieux. 2009. "Unconditional quantile regressions." *Econometrica*, 77(3): 953–973.

Frean, Molly, Jonathan Gruber, and Benjamin D Sommers. 2017. "Premium subsidies, the mandate, and Medicaid expansion: Coverage effects of the Affordable Care Act." *Journal of Health Economics*, 53: 72–86.

Garthwaite, Craig, Tal Gross, and Matthew J Notowidigdo. 2018. "Hospitals as insurers of last resort." *American Economic Journal: Applied Economics*, 10(1): 1–39.

Goldin, Jacob, Ithai Z Lurie, and Janet McCubbin. 2021. "Health insurance and mortality: Experimental evidence from taxpayer outreach." *The Quarterly Journal of Economics*, 136(1): 1–49.

Goodman-Bacon, Andrew. 2021. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics*, 225(2): 254–277.

Gross, Tal, and Matthew J Notowidigdo. 2011. "Health insurance and the consumer bankruptcy decision: Evidence from expansions of Medicaid." *Journal of Public Economics*, 95(7-8): 767–778.

Guth, Madeline, Bradley Corallo, Robin Rudowitz, and Rachel Garfield. 2021. "Medicaid Expansion Enrollment and Spending Leading up to the COVID-19 Pandemic." Accessed: 2024-08-26.

Hu, Luojia, Robert Kaestner, Bhashkar Mazumder, Sarah Miller, and Ashley Wong. 2018. "The effect of the Affordable Care Act Medicaid expansions on financial wellbeing." *Journal of Public Economics*, 163: 99–112.

Jaffe, Sonia, Anup Malani, and Julian Reif. 2026. "Access to credit reduces the value of insurance." *Journal of Public Economics*, 258: 105633.

Kaestner, Robert, Bowen Garrett, Jiajia Chen, Anuj Gangopadhyaya, and Caitlyn Fleming. 2017. "Effects of ACA Medicaid expansions on health insurance coverage and labor supply." *Journal of Policy Analysis and Management*, 36(3): 608–642.

Kearsley, A. 2024. "HHS Standard Values for Regulatory Analysis, 2024." Accessed: 2024-08-26.

Kluender, Raymond, Neale Mahoney, Francis Wong, and Wesley Yin. 2021. "Medical debt in the US, 2009-2020." *JAMA*, 326(3): 250–256.

Kluender, Raymond, Neale Mahoney, Francis Wong, and Wesley Yin. 2025. "The effects of medical debt relief: evidence from two randomized experiments." *The Quarterly Journal of Economics*, 140(2): 1187–1241.

Kranker, Keith. 2016. "Effects of Medicaid disease management programs on medical expenditures: evidence from a natural experiment in Georgia." *Journal of Health Economics*, 46: 52–69.

Levy, Helen, Thomas Buchmueller, and Sayeh Nikpay. 2019. "The impact of Medicaid expansion on household consumption." *Eastern Economic Journal*, 45: 34–57.

Lusardi, Annamaria, Daniel J Schneider, and Peter Tufano. 2011. "Financially fragile households: Evidence and implications." National Bureau of Economic Research.

Mazumder, Bhashkar, and Sarah Miller. 2016. "The effects of the Massachusetts health reform on household financial distress." *American Economic Journal: Economic Policy*, 8(3): 284–313.

Meinhofer, Angélica, and Allison E Witman. 2018. "The role of health insurance on treatment for opioid use disorders: Evidence from the Affordable Care Act Medicaid expansion." *Journal of Health Economics*, 60: 177–197.

Meyer, Bruce D, and James X Sullivan. 2023. "Consumption and Income Inequality in the United States since the 1960s." *Journal of Political Economy*, 131(2): 247–284.

Miller, Sarah, and Laura R Wherry. 2017. "Health and access to care during the first 2 years of the ACA Medicaid expansions." *New England Journal of Medicine*, 376(10): 947–956.

Miller, Sarah, Luojia Hu, Robert Kaestner, Bhashkar Mazumder, and Ashley Wong. 2021. "The ACA Medicaid expansion in Michigan and financial health." *Journal of Policy Analysis and Management*, 40(2): 348–375.

Miller, Sarah, Norman Johnson, and Laura R Wherry. 2021. "Medicaid and mortality: new evidence from linked survey and administrative data." *The Quarterly Journal of Economics*, 136(3): 1783–1829.

Nikpay, Sayeh, Seth Freedman, Helen Levy, and Tom Buchmueller. 2017. "Effect of the Affordable Care Act Medicaid expansion on emergency department visits: Evidence from state-level emergency department databases." *Annals of emergency medicine*, 70(2): 215–225.

Roth, Jonathan, and Pedro HC Sant'Anna. 2023. "When is parallel trends sensitive to functional form?" *Econometrica*, 91(2): 737–747.

Shupe, Cortnie. 2023. "Public health insurance and medical spending: The incidence of the ACA Medicaid expansion." *Journal of Policy Analysis and Management*, 42(1): 137–165.

Simon, Kosali, Aparna Soni, and John Cawley. 2017. "The impact of health insurance on preventive care and health behaviors: evidence from the first two years of the ACA Medicaid expansions." *Journal of Policy Analysis and Management*, 36(2): 390–417.

Soni, Aparna, Kosali Simon, John Cawley, and Lindsay Sabik. 2018. "Effect of Medicaid expansions of 2014 on overall and early-stage cancer diagnoses." *American Journal of Public Health*, 108(2): 216–218.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics*, 225(2): 175–199.

Wyse, Angela, and Bruce Meyer. 2023. "Saved By Medicaid: New Evidence on Health Insurance and Mortality from the Universe of Low-Income Adults." Working paper.

Figure 1: Medical spending and implied consumption among would-be Medicaid eligibles

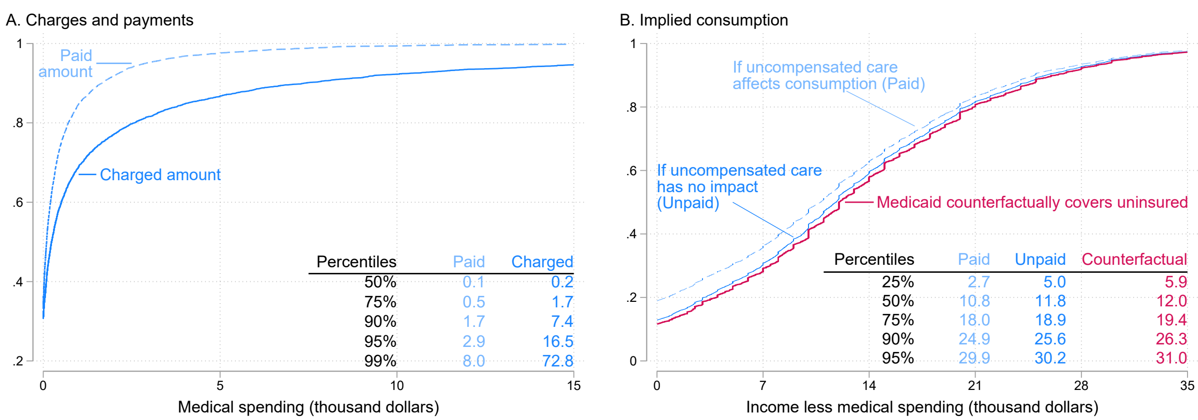

Notes: Panel A presents the cumulative distribution functions of both the medical spending paid and the amount charged among the uninsured. Panel B shows the empirical cumulative distribution function (CDF) of implied consumption for (i) those who paid their medical expenses, (ii) those who did not pay, and (iii) those who were uninsured but counterfactually covered by Medicaid. We use the Medical Expenditure Panel Survey (MEPS) for the years 2008–2013. See Appendix B for details.

Figure 2: Trends in percentiles of annual consumption

(a) Less than High School

(b) High School Graduates

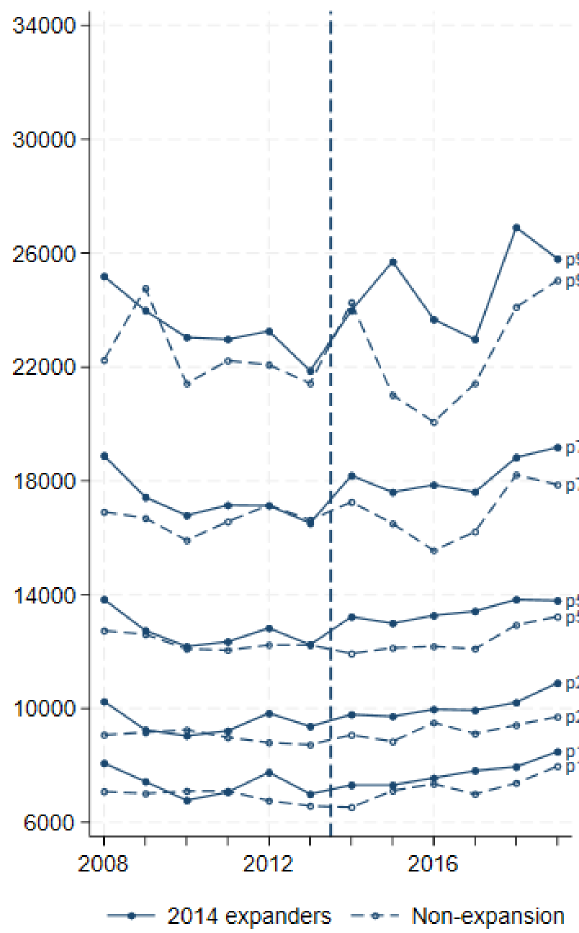

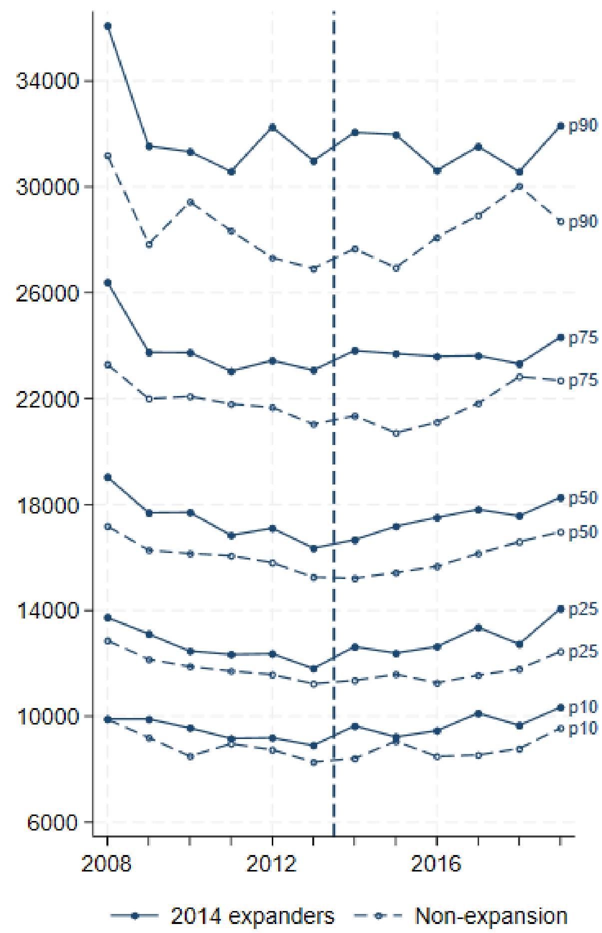

Notes: The figure plots the indicated percentiles of total consumption (excluding health insurance) for the states that expanded Medicaid in 2014 (solid line) and the states that never expanded (dashed line). The sample consists of people aged 22–64, with Panel (a) restricted to those with less than a high school education and Panel (b) to high school graduates (Kaestner et al., 2017). We multiply quarterly consumption from the Consumer Expenditure Survey by four to convert it to an annual basis for consistency.

Figure 3: Estimated impact of Medicaid expansion on the consumption distribution

(a) Less than High School

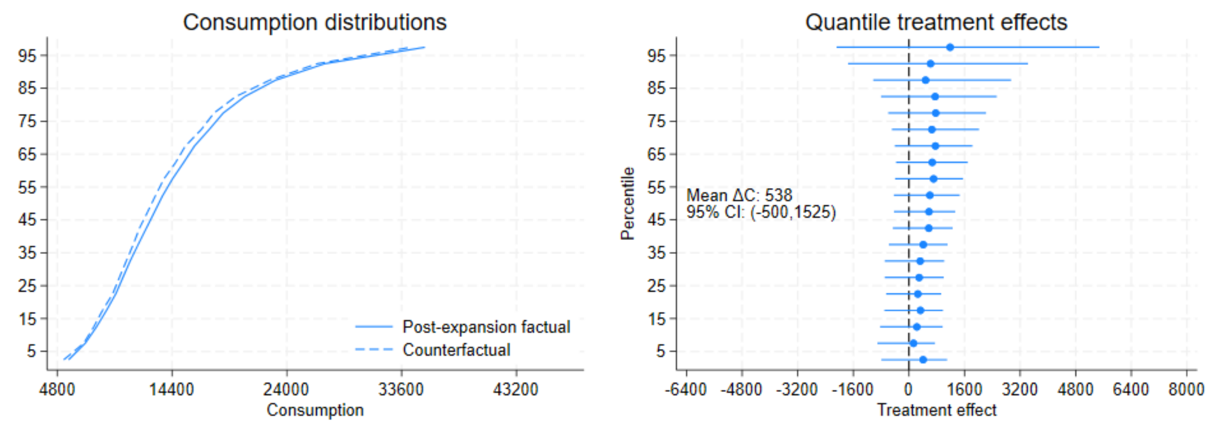

(b) High School Graduates

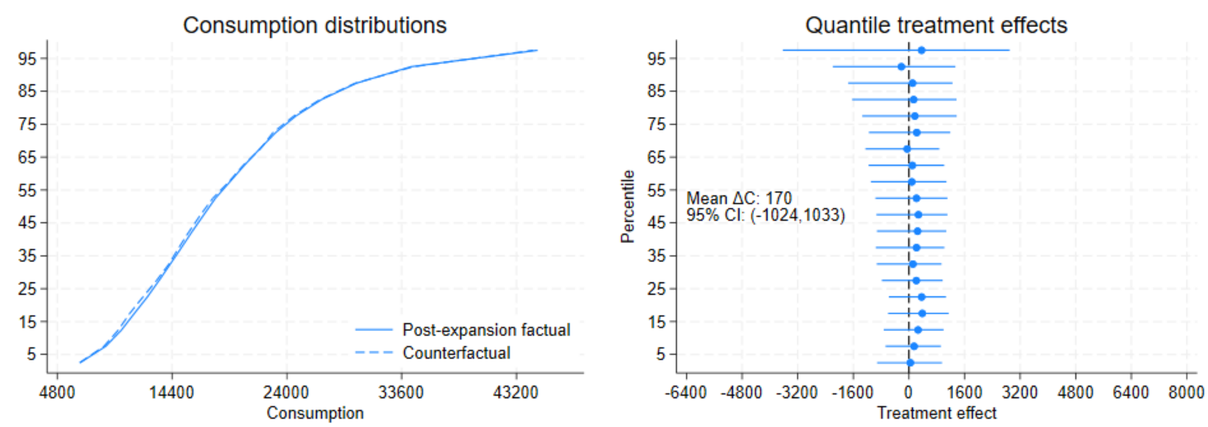

Notes: The figure plots the estimated factual and counterfactual consumption distributions (left panel) and the quantile treatment effects (right panel) estimated with quantile difference-in-differences with covariates (dummies for education, dummies for household size, top-coded at 6, dummies for family type, and age) with 95% confidence intervals calculated via the bootstrap (resampling states). We use the `qrprocess` command in Stata for quantile regression (Chernozhukov, Fernández-Val and Melly, 2022). We estimate effects separately for each timing group, and the figure reports the weighted average effect. The sample consists of individuals aged 22–64, with Panel (a) restricted to those with less than a high school education and Panel (b) to high school graduates (Kaestner et al., 2017). Results for individuals aged 22–64 with some college education and college graduates are presented in Appendix Figure A.7. We multiply quarterly consumption and estimates by four to convert them to an annual basis for consistency.

Figure 4: Consumption insurance value of Medicaid expansion and some benchmarks

(a) Less than High School

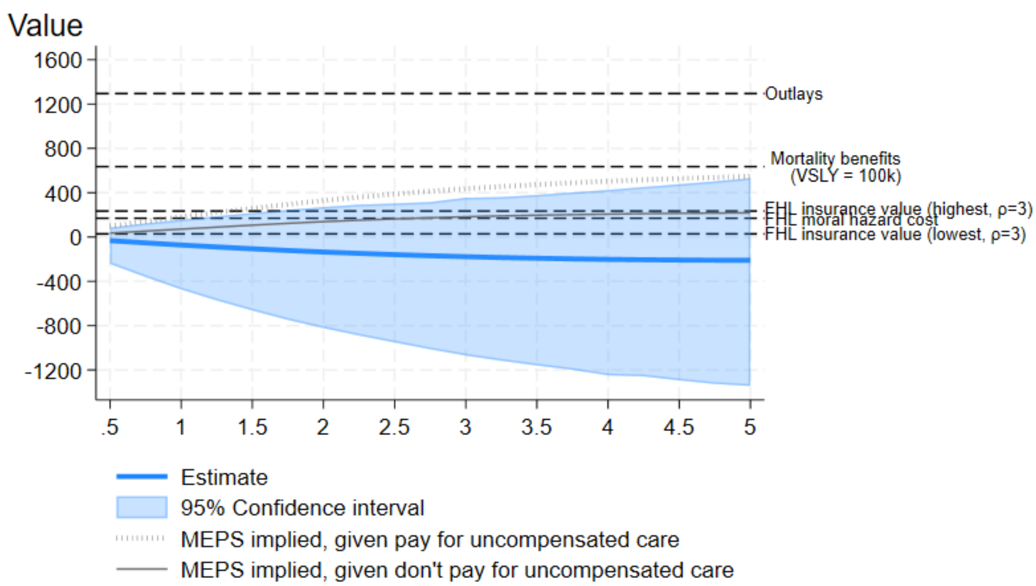

(b) High School Graduates

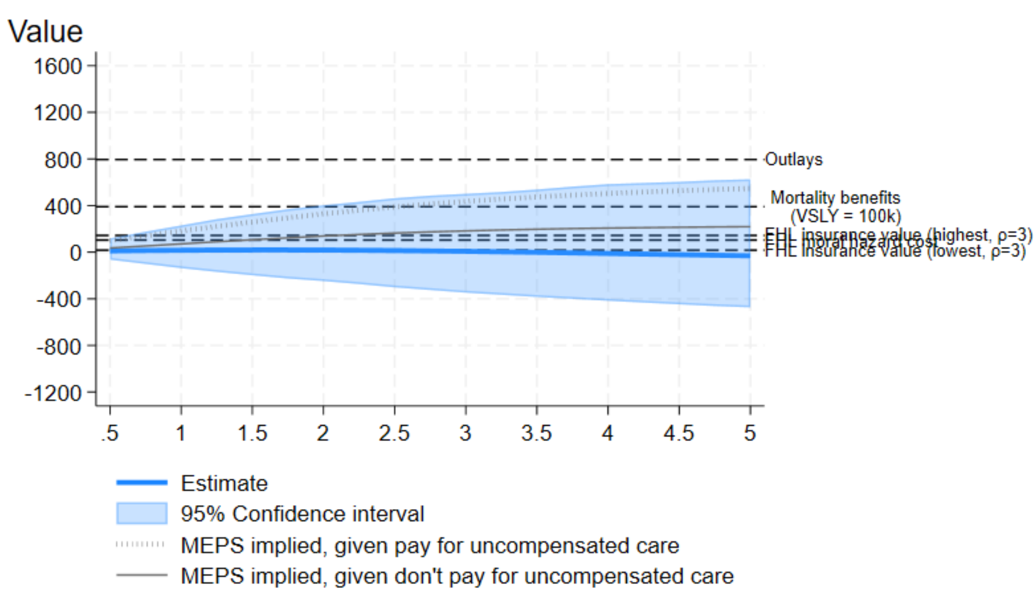

Notes: The figure plots our estimated average consumption insurance value with its 95% confidence interval, with the risk aversion parameter (ρ) on the x-axis, for two low-education groups: individuals aged 22–64, with Panel (a) restricted to those with less than a high school education and Panel (b) to high school graduates. The estimates are weighted using survey weights. Confidence intervals are based on the bootstrap, resampling states. We also report benchmarks for the value of Medicaid expansion, as described in Appendix C and MEPS implied values described in Appendix B. VSLY denotes the Value of Statistical Life Year. FHL denotes Finkelstein, Hendren and Luttmer (2019) and MEPS is the Medical Expenditure Panel Survey.

Table 1: Summary statistics on consumption by education level

	Expansion states		Nonexpansion states	
	Mean	(SD)	Mean	(SD)
A. Less than high school				
Income (annual, before tax)	18,403	(16,533)	17,761	(16,151)
Consumption (excl. health ins)	14,804	(8,246)	13,972	(7,680)
Well-measured consumption	12,668	(5,841)	11,782	(5,232)
# Observations		11,394		7,261
# Households		4,336		2,799
B. High school graduates				
Income (annual, before tax)	29,001	(24,064)	26,271	(21,610)
Consumption (excl. health ins)	19,444	(10,393)	17,891	(9,620)
Well-measured consumption	15,454	(6,565)	14,422	(6,100)
# Observations		24,992		14,666
# Households		9,299		5,583
C. Some college				
Income (annual, before tax)	36,502	(31,007)	32,411	(27,431)
Consumption (excl. health ins)	23,041	(13,065)	20,509	(11,337)
Well-measured consumption	17,443	(7,726)	15,669	(6,655)
# Observations		35,188		22,878
# Households		13,056		8,763
D. College or higher				
Income (annual, before tax)	61,439	(45,628)	53,642	(42,107)
Consumption (excl. health ins)	31,544	(18,930)	27,537	(16,143)
Well-measured consumption	21,853	(9,736)	19,418	(8,534)
# Observations		43,801		23,023
# Households		15,574		8,497

Notes: The sample consists of households from the 2008–2019 Consumer Expenditure Survey, classified by education level: (a) less than high school, (b) high school graduates, (c) some college, and (d) college graduates or higher. Expansion states include 2014–2019 expanders, and nonexpansion states include never-expanders (as of 2020). Well-measured consumption includes the flow value of housing, food at home, gas and motor oil, utilities, the flow value of new vehicles, and communication fees.

Table 2: Annual insurance value of Medicaid expansion (Less than High School)

Risk aversion	$\rho = 3$	$\rho = 1$	$\rho = 5$
A. All consumption (excluding health insurance), QDID, With covariates			
Insurance value	-177.2	-71.5	-210.2
95% CI	(-1070.0, 354.8)	(-473.0, 156.0)	(-1344.6, 531.8)
B. Without covariates			
Insurance value	-98.9	-50.9	-60.9
95% CI	(-973.0, 385.0)	(-419.8, 156.0)	(-1189.7, 615.6)
C. Alt consumption measure: Well-measured consumption			
Insurance value	-188.0	-54.4	-296.5
95% CI	(-761.8, 251.8)	(-296.8, 113.4)	(-1052.1, 391.2)
D. Alt consumption measure: Flexible consumption			
Insurance value	-119.8	-41.0	150.1
95% CI	(-359.1, 513.3)	(-143.8, 49.9)	(-467.2, 1272.0)
E. Alt estimator: CIC, Without Covariates			
Insurance value	-94.6	-42.7	-42.9
95% CI	(-931.3, 453.3)	(-405.1, 194.2)	(-1165.3, 767.5)
F. Excluding CA (given its early expansion)			
Insurance value	-139.2	-76.1	-95.6
95% CI	(-1102.6, 433.9)	(-500.9, 181.7)	(-1389.3, 682.8)
G. Excluding WI (given its high coverage)			
Insurance value	-158.2	-53.7	-202.6
95% CI	(-1038.7, 485.0)	(-450.6, 172.7)	(-1326.9, 691.4)
H. Finer grid (50-point, instead of 20-point grid)			
Insurance value	-180.9	-72.5	-208.0
95% CI	(-1025.5, 388.2)	(-459.2, 149.4)	(-1274.8, 623.0)

Notes: The table reports the implied annual insurance value of Medicaid expansion, for the indicated risk aversion value, consumption measure, estimation approach, sample, and grid granularity. The table reports insurance value for individuals aged 22–64 with less than a high school education. QDID is quantile difference-in-differences, and CIC is change-in-changes. Well-measured consumption includes housing, food, gas, utilities, vehicle, and communication expenses (Meyer and Sullivan, 2023). Flexible consumption excludes housing, utilities, and vehicle expenses from well-measured consumption. We report 95% confidence intervals, calculated via the bootstrap (resampling states), reported in parentheses. For Panels A, C, D, F, G, and H, covariates include dummies for education, dummies for household size (top-coded at 6), dummies for family type, and age.

Table 3: Annual insurance value of Medicaid expansion (High School Graduates)

Risk aversion	$\rho = 3$	$\rho = 1$	$\rho = 5$
A. All consumption (excluding health insurance), QDID, With covariates			
Insurance value	7.5	16.7	-30.8
95% CI	(-347.8, 501.5)	(-137.4, 229.8)	(-473.6, 626.1)
B. Without covariates			
Insurance value	-96.1	-22.8	-166.0
95% CI	(-568.6, 401.2)	(-243.5, 199.2)	(-743.3, 474.4)
C. Alt consumption measure: Well-measured consumption			
Insurance value	37.4	49.4	-24.8
95% CI	(-218.5, 336.6)	(-49.9, 157.0)	(-372.5, 410.0)
D. Alt consumption measure: Flexible consumption			
Insurance value	13.6	26.2	-23.9
95% CI	(-152.0, 208.3)	(-44.9, 109.3)	(-236.9, 262.7)
E. Alt estimator: CIC, Without Covariates			
Insurance value	-72.7	-9.4	-149.7
95% CI	(-543.6, 524.1)	(-226.7, 262.8)	(-724.8, 634.5)
F. Excluding CA (given its early expansion)			
Insurance value	67.3	53.1	17.1
95% CI	(-364.7, 568.7)	(-139.2, 262.3)	(-505.4, 663.5)
G. Excluding WI (given its high coverage)			
Insurance value	56.1	23.1	55.9
95% CI	(-307.7, 544.3)	(-145.8, 239.1)	(-412.7, 670.0)
H. Finer grid (50-point, instead of 20-point grid)			
Insurance value	73.0	43.0	69.9
95% CI	(-339.3, 492.1)	(-124.7, 215.7)	(-468.8, 597.2)

Notes: The table reports the implied annual insurance value of Medicaid expansion for the indicated risk aversion value, consumption measure, estimation approach, sample, and grid granularity. The table reports insurance value for individuals aged 22–64 with a high school degree. QDID is quantile difference-in-differences, and CIC is change-in-changes. Well-measured consumption includes housing, food, gas, utilities, vehicle, and communication expenses (Meyer and Sullivan, 2023). Flexible consumption excludes housing, utilities, and vehicle expenses from well-measured consumption. We report 95% confidence intervals, calculated via the bootstrap (resampling states), reported in parentheses. For Panels A, C, D, F, G, and H, covariates include dummies for education, dummies for household size (top-coded at 6), dummies for family type, and age.

A Appendix Figures and Tables

Figure A.1: Timing of state Medicaid expansions

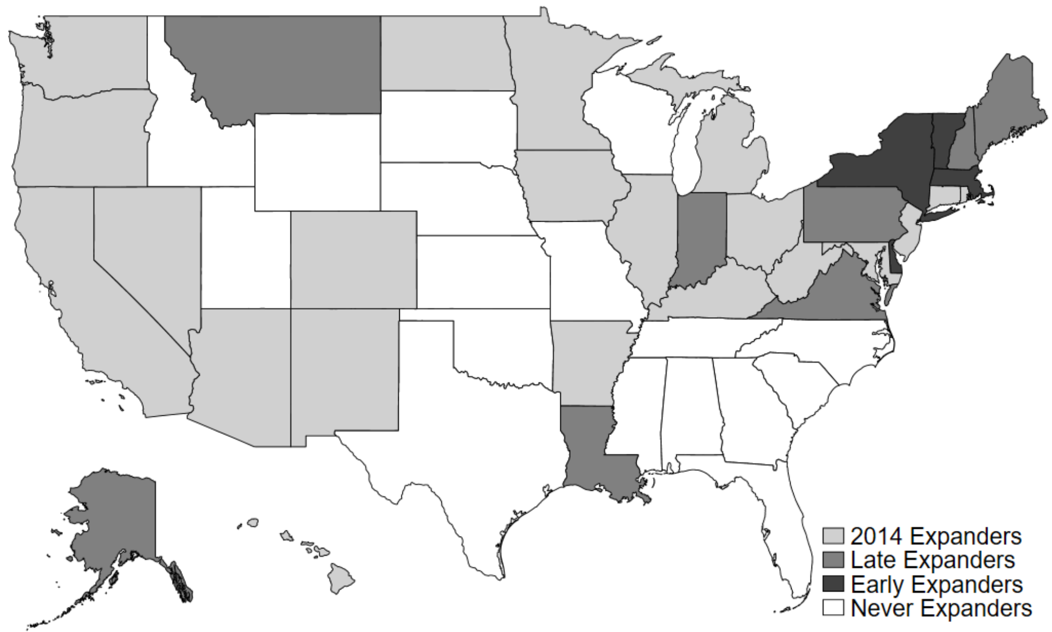

Notes: The figure indicates the timing of states' Medicaid expansion decisions as of 2019. We use the classification in Miller, Johnson and Wherry (2021). Among the late expanders, Indiana, New Hampshire, and Pennsylvania adopted the expansion in 2015, Alaska and Montana in 2016, Louisiana in 2017, and Maine and Virginia in 2019. Other states (e.g., North Carolina) expanded after our sample period ends.

Figure A.2: Percentiles of annual consumption, 2015 expanders

(a) Less than High School (b) High School Graduates

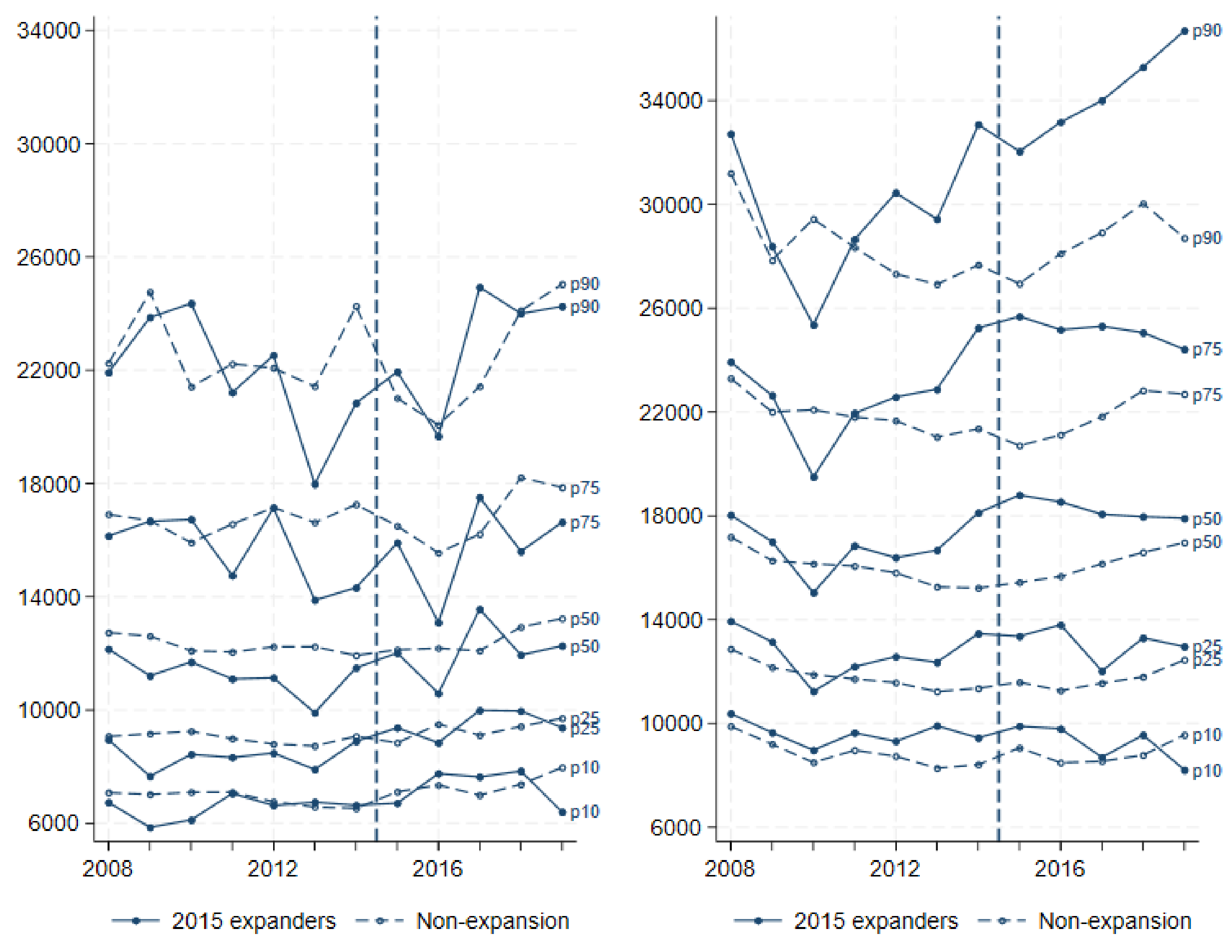

Notes: The figure plots the indicated percentiles of total consumption (excluding health insurance) for the states that expanded Medicaid in 2015 (solid line) and the states that never expanded (dashed line). The sample consists of people aged 22–64, with Panel (a) restricted to those with less than a high school education and Panel (b) to high school graduates (Kaestner et al., 2017). We multiply quarterly consumption from the Consumer Expenditure Survey by four to convert it to an annual basis for consistency.

Figure A.3: Percentiles of annual consumption, 2016 expanders

(a) Less than High School (b) High School Graduates

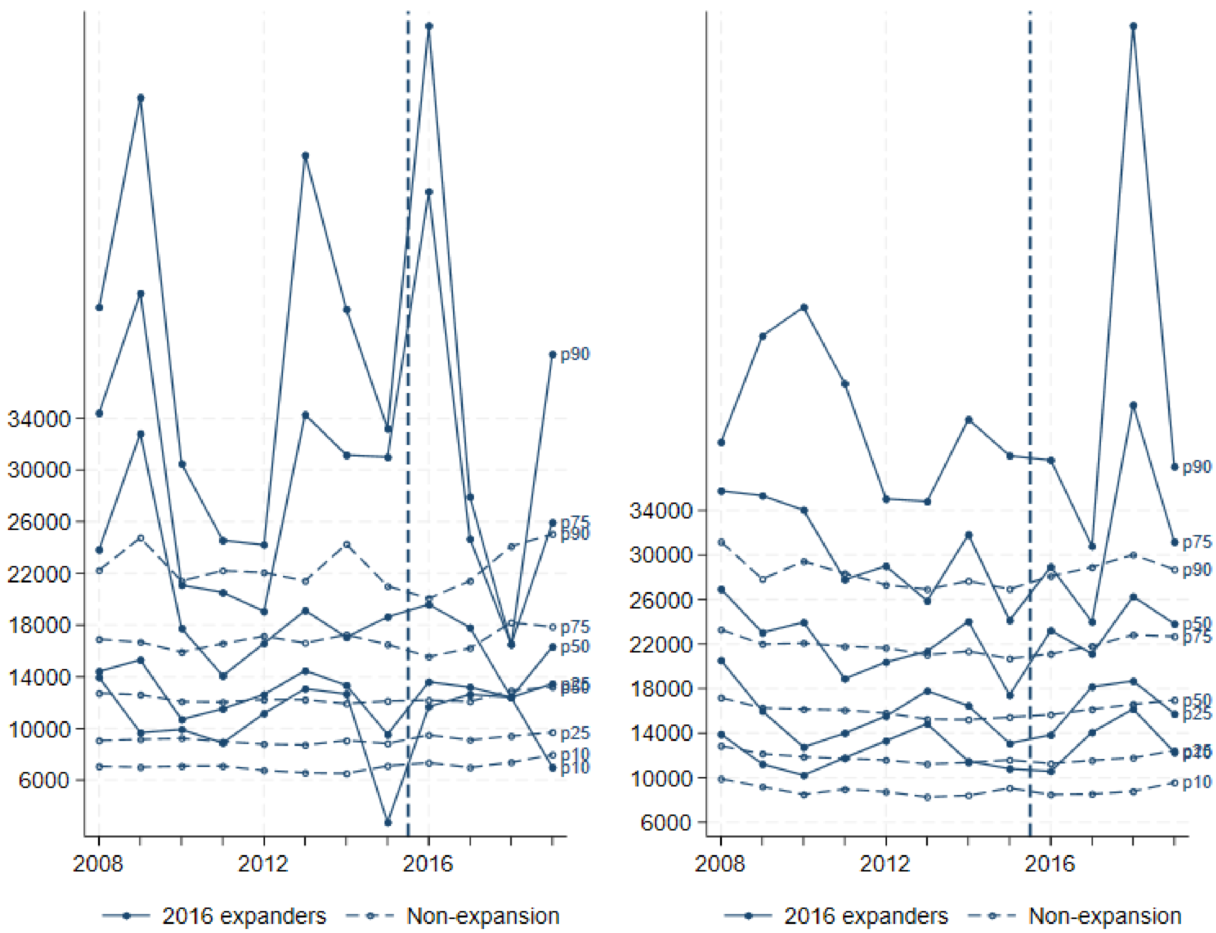

Notes: The figure plots the indicated percentiles of total consumption (excluding health insurance) for the states that expanded Medicaid in 2016 (solid line) and the states that never expanded (dashed line). The sample consists of people aged 22–64, with Panel (a) restricted to those with less than a high school education and Panel (b) to high school graduates (Kaestner et al., 2017). We multiply quarterly consumption from the Consumer Expenditure Survey by four to convert it to an annual basis for consistency.

Figure A.4: Percentiles of annual consumption, 2017 expanders

(a) Less than High School (b) High School Graduates

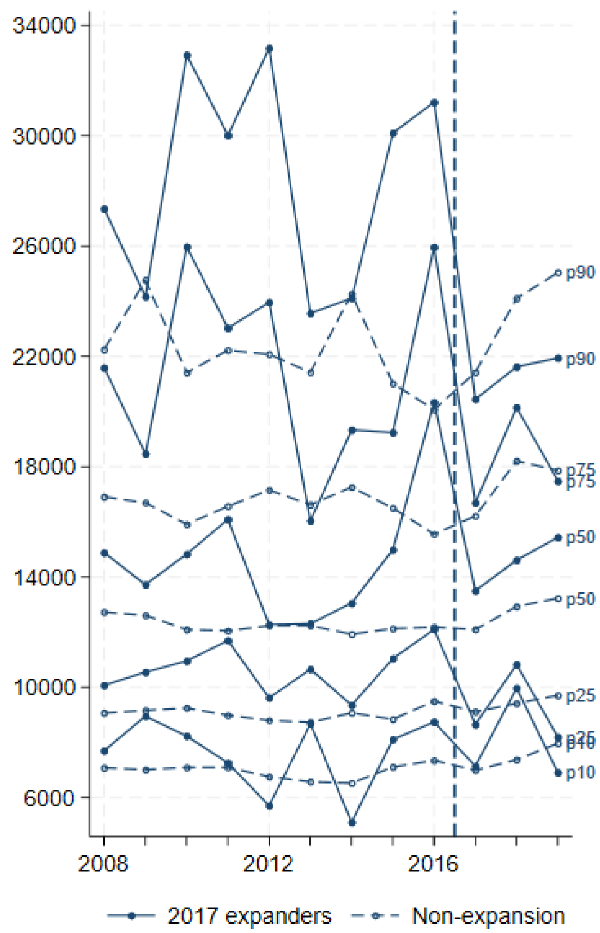

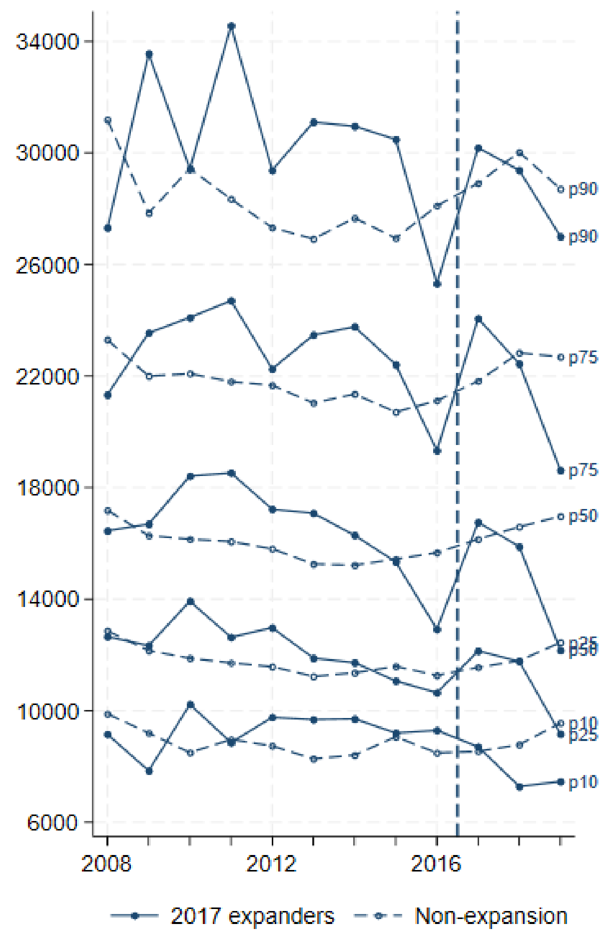

Notes: The figure plots the indicated percentiles of total consumption (excluding health insurance) for the states that expanded Medicaid in 2017 (solid line) and the states that never expanded (dashed line). The sample consists of people aged 22–64, with Panel (a) restricted to those with less than a high school education and Panel (b) to high school graduates (Kaestner et al., 2017). We multiply quarterly consumption from the Consumer Expenditure Survey by four to convert it to an annual basis for consistency.

Figure A.5: Percentiles of annual consumption, 2019 expanders

(a) Less than High School (b) High School Graduates

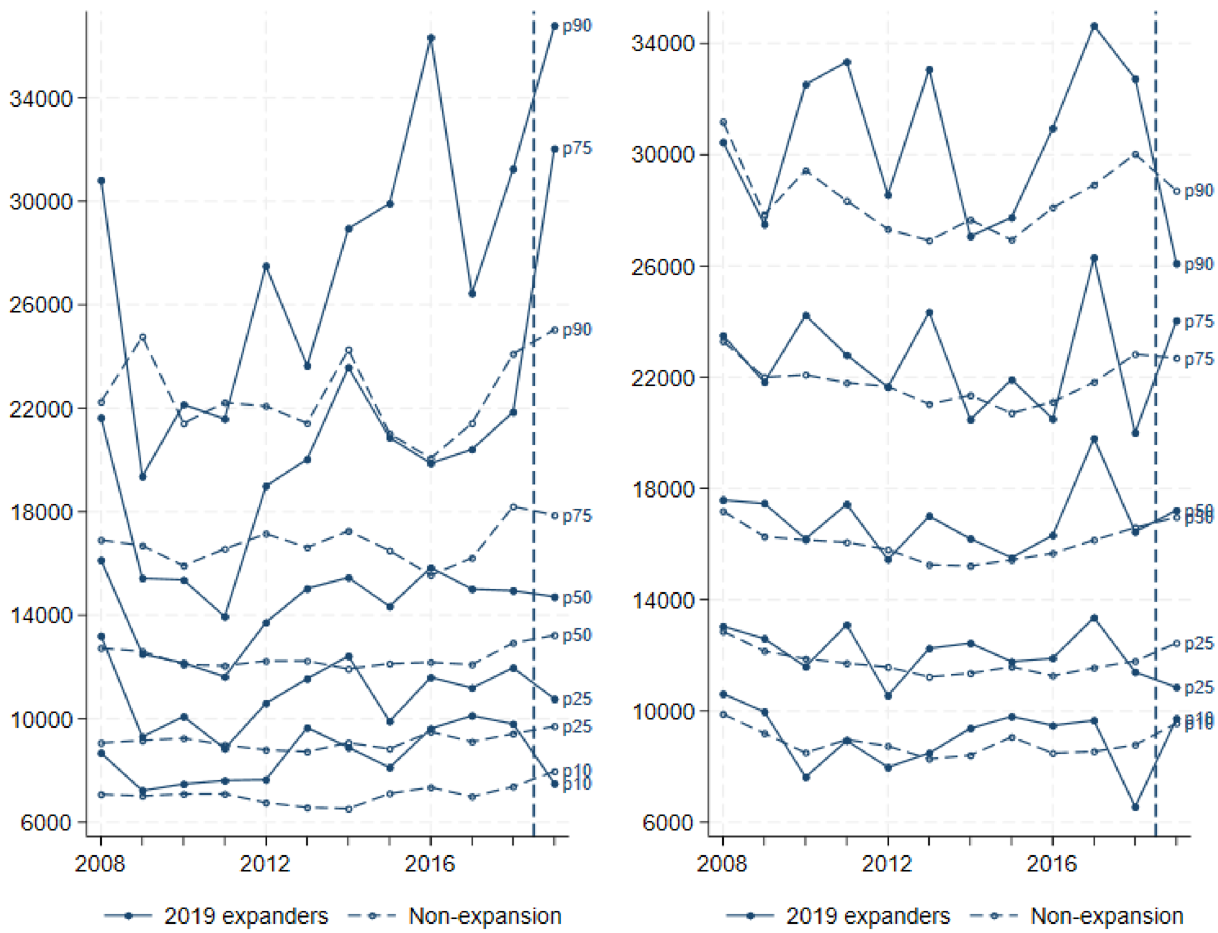

Notes: The figure plots the indicated percentiles of total consumption (excluding health insurance) for the states that expanded Medicaid in 2019 (solid line) and the states that never expanded (dashed line). The sample consists of people aged 22–64, with Panel (a) restricted to those with less than a high school education and Panel (b) to high school graduates (Kaestner et al., 2017). We multiply quarterly consumption from the Consumer Expenditure Survey by four to convert it to an annual basis for consistency.

Figure A.6: CIC estimates of the impact of Medicaid expansion on the consumption distribution among low-education groups

(a) Less than High School

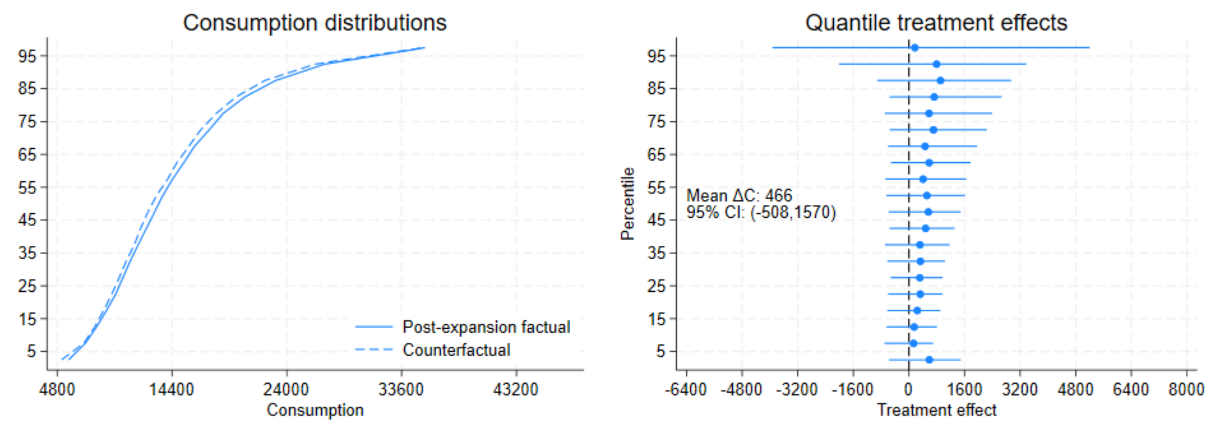

(b) High School Graduates

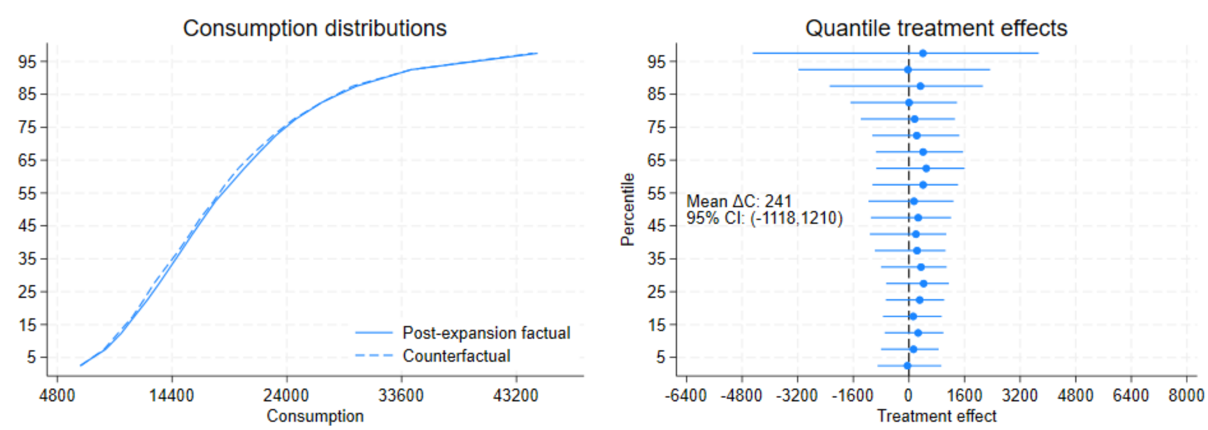

Notes: The figure is identical to Figure 3 except that we use change-in-changes (CIC) without covariates, rather than quantile difference-in-differences (QDID), to estimate the distributional effects.

Figure A.7: Impact of Medicaid expansion on the consumption distribution among higher-education groups

(a) Some college

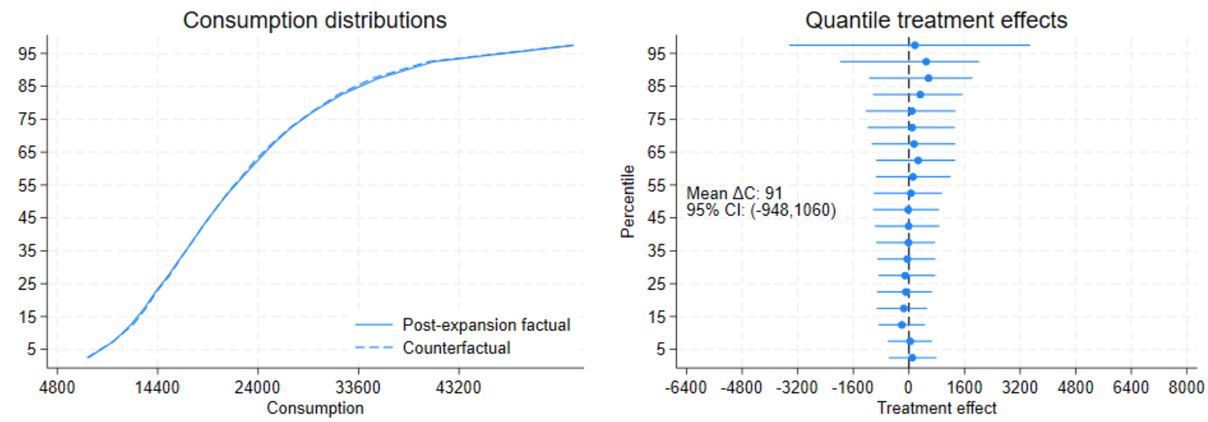

(b) College graduates or higher

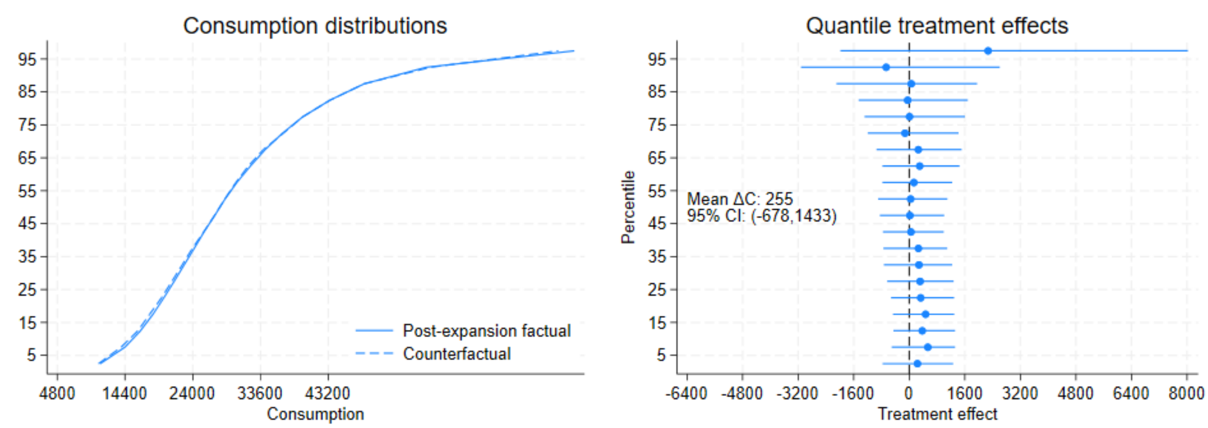

Notes: The figure plots the estimated factual and counterfactual consumption distributions (left panel) and the quantile treatment effects (right panel) for two higher-education groups: (a) some college and (b) college graduates or higher. The quantile treatment effects are estimated with quantile difference-in-differences with covariates (dummies for education, dummies for household size, top-coded at 6, dummies for family type, and age) with 95% confidence intervals calculated via the bootstrap (resampling states). We use the `qrprocess` command in Stata for quantile regression (Chernozhukov, Fernández-Val and Melly, 2022). We estimate effects separately for each timing group, and the figure reports the weighted average effect. The sample is defined in the notes to Table 1.

Table A.1: Summary statistics on major consumption categories by education level

	p(10)	p(25)	p(50)	Mean	p(75)	p(90)
			A. Less than high school			
a. Expansion states						
Consumption (all)	8,141	10,628	14,516	16,522	20,073	27,279
Food at home	1,047	1,538	2,187	2,435	2,938	4,053
Gas and motor oil	0	624	1,557	2,047	2,819	4,587
Flow value of housing	2,276	3,300	4,822	5,583	6,960	9,637
Utilities	455	821	1,349	1,524	2,029	2,778
Flow value of new vehicles	0	58	253	551	743	1,441
Communication	0	223	440	528	737	1,072
b. Nonexpansion states						
Consumption (all)	7,797	10,186	13,881	15,533	19,001	25,068
Food at home	1,064	1,568	2,198	2,407	2,946	3,978
Gas and motor oil	0	651	1,613	2,040	2,768	4,466
Flow value of housing	1,922	2,705	3,951	4,556	5,712	7,712
Utilities	621	1,030	1,513	1,658	2,126	2,819
Flow value of new vehicles	0	77	285	593	754	1,534
Communication	0	230	440	527	713	1,087
			B. High school graduates			
a. Expansion states						
Consumption (all)	10,806	14,460	20,311	22,321	27,627	35,796
Food at home	1,190	1,736	2,391	2,625	3,182	4,296
Gas and motor oil	0	950	1,835	2,325	3,210	4,811
Flow value of housing	2,759	4,063	5,989	6,891	8,675	12,086
Utilities	713	1,214	1,823	1,963	2,541	3,315
Flow value of new vehicles	0	190	626	978	1,381	2,339
Communication	111	314	592	673	940	1,302
b. Nonexpansion states						
Consumption (all)	9,856	13,424	18,423	20,352	25,303	32,233
Food at home	1,115	1,664	2,338	2,554	3,101	4,149
Gas and motor oil	318	997	1,866	2,314	3,117	4,684
Flow value of housing	2,529	3,607	5,197	5,960	7,410	10,113
Utilities	833	1,279	1,815	1,947	2,461	3,215
Flow value of new vehicles	0	204	630	986	1,395	2,367
Communication	40	305	577	660	925	1,294

	p(10)	p(25)	p(50)	Mean	p(75)	p(90)
C. Some college						
a. Expansion states						
Consumption (all)	12,526	17,266	23,844	26,195	31,940	41,563
Food at home	1,218	1,792	2,492	2,750	3,335	4,499
Gas and motor oil	425	1,063	2,034	2,469	3,302	4,983
Flow value of housing	3,272	4,811	7,042	8,236	10,198	14,303
Utilities	796	1,309	1,922	2,072	2,639	3,460
Flow value of new vehicles	0	298	821	1,167	1,626	2,674
Communication	132	365	669	749	1,029	1,424
b. Nonexpansion states						
Consumption (all)	11,655	15,512	21,411	23,392	28,496	36,730
Food at home	1,167	1,718	2,379	2,609	3,191	4,267
Gas and motor oil	538	1,070	1,950	2,386	3,231	4,702
Flow value of housing	2,846	4,075	5,880	6,741	8,287	11,539
Utilities	862	1,337	1,879	2,018	2,536	3,290
Flow value of new vehicles	93	339	832	1,188	1,616	2,645
Communication	81	349	647	726	1,005	1,399
D. College graduates or higher						
a. Expansion states						
Consumption (all)	17,931	23,957	31,694	35,359	41,946	55,530
Food at home	1,416	2,027	2,784	3,131	3,789	5,146
Gas and motor oil	593	1,151	2,079	2,478	3,283	4,840
Flow value of housing	4,915	6,952	10,013	11,722	14,255	20,535
Utilities	896	1,404	2,033	2,236	2,809	3,760
Flow value of new vehicles	137	539	1,155	1,500	2,051	3,169
Communication	138	395	713	787	1,073	1,488
b. Nonexpansion states						
Consumption (all)	15,851	21,098	28,105	31,318	37,166	48,720
Food at home	1,356	1,910	2,645	2,927	3,552	4,726
Gas and motor oil	644	1,186	2,079	2,498	3,283	4,832
Flow value of housing	4,182	5,726	8,061	9,442	11,501	15,954
Utilities	966	1,443	2,042	2,215	2,755	3,615
Flow value of new vehicles	222	597	1,204	1,550	2,088	3,226
Communication	72	384	703	786	1,079	1,500

Notes: The sample consists of households from the 2008–2019 Consumer Expenditure Survey, classified by education level: (a) less than high school, (b) high school graduates, (c) some college, and (d) college graduates or higher. Expansion states include 2014–2019 expanders, and nonexpansion states include never-expanders (as of 2020). Consumption (all) mainly consists of food at home, gas and motor oil, flow value of housing, utilities, flow value of new vehicles, and communication.

B Details on MEPS Data

In some supplemental analyses, we use data from the Medical Expenditure Panel Survey (MEPS) to provide evidence on uninsured health spending risk. MEPS is a high-quality, government-run survey designed to measure medical spending. We use data from 2008 to 2013, the period immediately prior to most states' adoption of their Medicaid expansions. (We use these years rather than post-2014 years in nonexpansion states because the publicly available MEPS data do not contain state identifiers.) Households participate in five semiannual interviews covering two years, in which they report the insurance coverage and health care utilization of each member. MEPS officials then contact health care providers to determine amounts paid for services and amounts charged for services.

We use these data in two ways. First, in Figure 1, we show the distribution of spending and charges among uninsured, potential Medicaid eligibles prior to the Affordable Care Act (ACA). Second, we use the distribution of spending and charges among a broader population to develop benchmarks for the insurance value of Medicaid expansion. Here, we provide more detail.

For both approaches, we begin by shifting the data from the individual level to the household level. We use MEPS's constructed "CPS family ID" to identify individuals sharing consumption resources. We use family weights in calculating all statistics and distributions. We then construct the total paid amount for each household (h) and year (t) by summing out-of-pocket medical payments for each household member, m_{ht}. To construct billed amounts, we sum the charged amount for uninsured members and the paid amount for insured members, b_{ht}. (We do not include charges for insured members because insurers typically negotiate prices well below charges.)

The MEPS carefully records medical expenditures and billed amounts but does not record household consumption. We therefore impute three measures of consumption, starting from family income y and then subtracting different measures of medical spending. All measures of imputed consumption are subject to a consumption floor $\underline{c}$. We begin with two measures of baseline consumption, in the absence of Medicaid expansion:

1. c_0^{paid}: Consumption in the absence of Medicaid expansion (state 0), assuming that households pay off charges in the year they accrue: $c_0^{paid} = \max(y - b, \underline{c})$.

2. c_0^{unpaid}: Consumption in the absence of Medicaid expansion (state 0), assuming that households never pay charges beyond their immediate expenditures m: $c_0^{unpaid} = \max(y - m, \underline{c})$.

For insured households, we set consumption equal to c_0^{unpaid}, assuming that providers accept the payment offered by insurers and do not balance bill. For uninsured households, we simulate both c_0^{paid} and c_0^{unpaid} to show the role of uncompensated care, i.e., billed amounts that are not collected.

Finally, our third measure of consumption is counterfactual consumption under Medicaid expansion, c_1. For individuals who are uninsured at baseline and live in households with income below 138% of the poverty line, $c_1 = \max(y, \underline{c})$, i.e., we assume that Medicaid covers all medical expenses. For other households, $c_1 = c_0^{unpaid}$, meaning expansion does not affect their consumption.

To measure medical spending and consumption of the uninsured (Figure 1), we limit the sample to household-years with at least one uninsured member and income below 138% of the poverty line. These are the households most likely to have benefited from the Medicaid expansion. We then plot the distribution of paid amounts and charged amounts, as well as c_0^{paid}, c_0^{unpaid}, and c_1.

To create our benchmarks for the insurance value, we focus on a MEPS sample analogous to our CEX estimation sample: households with at least one member aged 22–64 with no college education. We then make the consumption measures analogous to our CEX measures by dividing by our household equivalence scale. Finally, to evaluate utility functions, we impose a floor on consumption of $\underline{c}$. We impose $\underline{c} = \$5{,}000$ per year, a conservative choice that produces lower insurance value benchmarks.[13] We construct two benchmarks: the insurance value assuming that the uninsured do or do not pay for their uncompensated care. For both benchmarks, we calculate the insurance value using the main approach in the text, and we take as the consumption distribution with Medicaid the empirical cumulative distribution function (CDF) of c_1. To obtain the insurance value of Medicaid expansion if the uninsured do pay for their uncompensated care, we set the consumption distribution without Medicaid equal to the empirical CDF of c_0^{paid}. To obtain the insurance value if the uninsured do not pay for their uncompensated care, we set the consumption distribution without Medicaid equal to the empirical CDF of c_0^{unpaid}. In Figure 4, we plot these insurance values (as a function of risk aversion) as gray lines.

[13] This value is high relative to what others use. For example, De Nardi, French and Jones (2010) estimate $\underline{c} = \$2{,}700$, and Finkelstein, Hendren and Luttmer (2019) set $\underline{c} = \$2{,}000$.

C Details on Benchmarks

In Figure 4, we compare our consumption insurance value estimates to several benchmark values. We describe the construction of the Medical Expenditure Panel Survey–related benchmarks in Appendix B. The other benchmarks are taken from the literature, which typically reports estimates per Medicaid enrollee. Because our sample includes both unenrolled eligible and ineligible individuals, the average benefit and the average cost of expansion are lower in our sample than the corresponding enrollee-based benchmarks. We therefore make simple adjustments to ensure comparability. Specifically, we calculate (income-based) eligibility for all sample members and assume that half of eligible individuals in the expansion population take up Medicaid (Decker, Abdus and Lipton, 2022). In our sample, the shares of individuals eligible for Medicaid are 42 percent for those with less than a high school education and 26 percent for high school graduates. Therefore, we scale down all enrollee-based benchmarks by 21 percent and 13 percent, respectively.

Below, we describe the benchmarks that we take from the literature and how we adjust them for comparability.

Outlays: Guth et al. (2021) report that government outlays on the expansion population averaged $6,110 per enrollee. Scaling by 21 percent and 13 percent for each low-education group to account for both noneligibility and nonenrollment in our population yields costs of $1,283 and $794 per person, respectively.

Moral hazard cost: Finkelstein, Hendren and Luttmer (2019) estimate the moral hazard cost of Medicaid coverage using data from the Oregon Health Insurance Experiment and multiple approaches. They find annual moral hazard costs that range from $585 to $879 per Medicaid beneficiary. Using $800 as the relevant benchmark in this range, scaling by 21 percent and 13 percent to reflect limited eligibility and take-up in our low-education sample yields moral hazard costs of about $168 and $104, respectively.

Mortality reduction benefit: Wyse and Meyer (2023) estimate the impact of Medicaid expansion on mortality. They find that expansions increased coverage by 28.7 million coverage-years and reduced mortality by 831,890 life-years, or 0.03 life-years per enrollee-year. We value this increase at $3,000 per enrollee-year, assuming a value of $100,000 per life-year. This value is low relative to government guidance for valuing mortality risk reductions, which is closer to $500,000 per life-year (Kearsley, 2024). This low value is conservative with respect to our conclusion that the consumption-smoothing component of insurance value is low relative to the mortality benefit. Scaling by 21 percent and 13 percent yields mortality reduction benefits of $630 and $390, respectively.

Other estimates of insurance value: Finkelstein, Hendren and Luttmer (2019) and Shupe (2023) use medical spending data to estimate the insurance value of Medicaid. Finkelstein, Hendren and Luttmer (2019) estimate the insurance value of Medicaid to enrollees in the Oregon Health Insurance Experiment using multiple approaches. Their estimates range from $133 to $1,106 per enrollee-year. Scaling down yields benchmarks of $28 and $232 for non-high school graduates and $17 and $144 for high school graduates. Shupe (2023) estimate the insurance value of Medicaid expansion among the Medicaid-eligible (but not necessarily enrolled) population, finding a value of $0–$70, depending on risk aversion. Scaling down by our assumed take-up rates yields benchmarks of \$0–\$20.

D Insurance Value for Two Higher Education Groups

The purpose of this section is to report the effects of Medicaid expansion on insurance value among individuals with higher educational attainment—those with some college education or a college degree or higher—for completeness.

To ensure comparability with our estimated values, we linearly scale all benchmark estimates by (i) an assumed Medicaid take-up rate conditional on eligibility of 50 percent and (ii) the share of our sample that is income-eligible for Medicaid. Specifically, the shares of individuals eligible for Medicaid in our sample are 42 percent for those with less than a high school education, 26 percent for high school graduates, 19 percent for those with some college education, and 7 percent for college graduates or above. This implies scale-down factors of 21 percent, 13 percent, 9 percent, and 3 percent, respectively.

This procedure mechanically reduces the benchmark magnitudes substantially. Given limited statistical power among individuals with education beyond high school, we focus our primary analysis on the low-education groups, while also presenting results for those with some college education and for college graduates or higher. Moreover, Appendix Figure A.7 indicates that not only is the number of affected individuals small, but the conditional effects on consumption risk are also modest relative to the low-education groups. Accordingly, we center our discussion on the low-education population most directly affected by Medicaid expansion. The analogue to Figure 4 is presented in Figure D.1.

Figure D.1: Consumption insurance value of Medicaid expansion and some benchmarks

(a) Some college

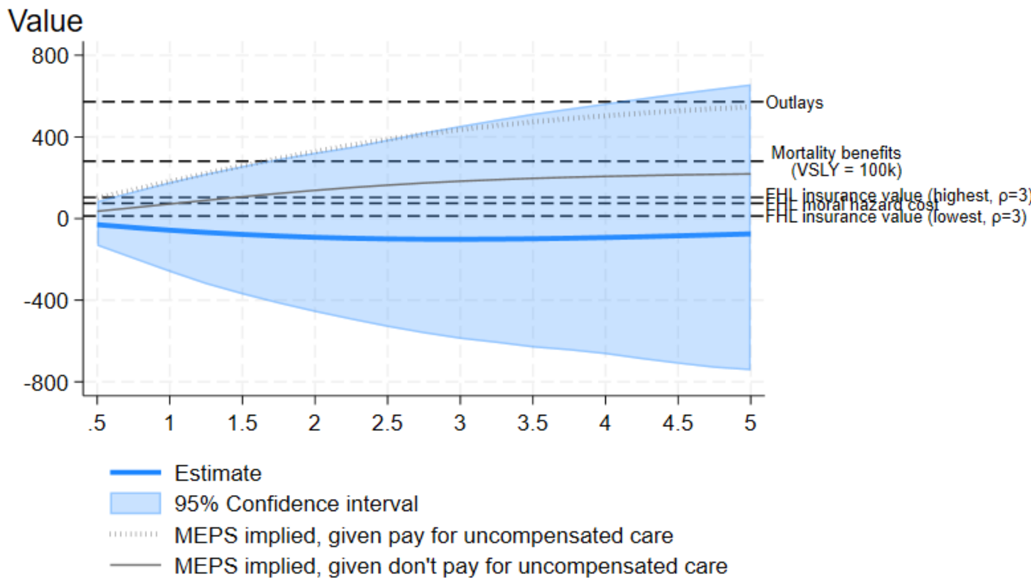

(b) College graduates or higher

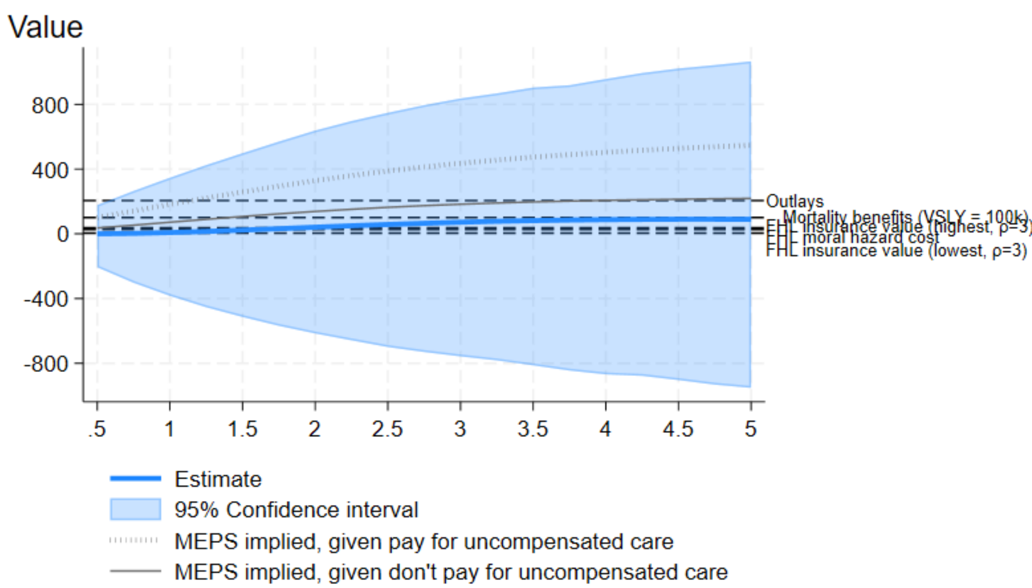

Notes: The figure plots our estimated average consumption insurance value with its 95% confidence interval for two higher-education groups. The estimates are weighted using survey weights. Confidence intervals are based on the bootstrap, resampling states. This figure is analogous to Figure 4, but for higher-education groups. VSLY denotes the Value of Statistical Life Year, FHL denotes Finkelstein, Hendren and Luttmer (2019), and MEPS is the Medical Expenditure Panel Survey.