Does Health Insurance Reduce Consumption Risk? Evidence from Medicaid Expansions*

Anita Mukherjee Daniel W. Sacks Hoyoung Yoo

March 21, 2024

PRELIMINARY AND INCOMPLETE.
PLEASE DO NOT CITE OR CIRCULATE.

Abstract

We investigate the consumption smoothing effects of Medicaid using staggered, state-level expansion decisions. While Medicaid expansions improve health and financial status, they may not smooth consumption risk because a great deal of uninsured medical spending is financed with bad debt and charity care rather than reduced consumption. Using quantile difference-in-differences and changes-in-changes specifications, we find small effects of Medicaid expansion throughout the consumption distribution. Our estimates, combined with an assumed utility function, imply near-zero insurance value from Medicaid expansion. While our point estimates are uncertain, our confidence intervals nevertheless allow us to rule out the possibility that a large share of Medicaid's value comes from reduced consumption risk.

Keywords: Medicaid expansion, insurance, welfare analysis, risk premium
JEL codes: J21, E24, K32, E61, E71

*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 Wisconsin-Madison, hyoo36@wisc.edu. We are grateful to seminar participants at the Household Finance Seminar, Center for Financial Studies, University of Wisconsin-Madison, and the Finance Department Seminar, Tippie School of Business, University of Iowa.

1 Introduction

In the decade since the implementation of the Affordable Care Act (ACA), the share of non-elderly Americans with health insurance has increased from 79 percent to 89 percent (Kaiser Family Foundation, 2024). The majority of new insurance coverage has been in the Medicaid program (Frean, Gruber and Sommers, 2017), which covers over 80 million people and is the largest means-tested program in the United States (Donohue et al., 2022). The expansion of this behemoth program has spurred many studies documenting an array of positive benefits on the physical and financial health on those newly eligible. Yet little is known about how well Medicaid expansions perform their core insurance function: smoothing the marginal utility of consumption across states of the world (Finkelstein, Mahoney and Notowidigdo, 2018).

In this paper, we provide new evidence on the consumption-smoothing effects of Medicaid expansion, and quantify the insurance value implied by this consumption smoothing. Despite its importance, the consumption smoothing insurance value of Medicaid expansion remains an open question in the literature, for two reasons. First, existing research does not provide direct evidence on how health insurance coverage affects the distribution of consumption in general, and consumption risk in particular. Of the more than four hundred papers studying the effects of Medicaid expansion following the Affordable Care Act reviewed in Guth et al. (2020), only one studies the effects of Medicaid on consumption: Levy, Buchmueller and Nikpay (2019), which examines only average effects, not distributional. The literature instead focuses on the effects of Medicaid expansion on health, health care utilization, and financial outcomes. In an important exception, Finkelstein, Hendren and Luttmer (2019) estimates the value of Medicaid to beneficiaries, including its insurance value. That paper, however, does not estimate the effect of Medicaid on the consumption distribution, instead inferring this impact from the impact of Medicaid on medical out-of-pocket spending.

Second, the impact of health insurance on consumption risk is unresolved because the benefits of health insurance coverage flow not only to households in the form of increased consumption, but also to health care providers, via reductions in uncompensated care—specifically bad debt and charity care. Such uncompensated care is an important source of health care financing for the uninsured, who pay only a small fraction of their health care bills (Mahoney, 2015). Health insurance expansions therefore end up benefiting health care providers, especially hospitals (Dranove, Garthwaite and Ody, 2016; Garthwaite, Gross and Notowidigdo, 2018; Duggan, Gupta and Jackson, 2022). Thus greater health insurance coverage does not automatically imply a reduced exposure to health care risk among households; instead, coverage reduces risk born by providers. Indeed, Finkelstein, Hendren and Shepard (2019) finds that recipients of subsidized health insurance have willingness-to-pay below their own costs, and Finkelstein, Hendren and Luttmer (2019) finds that the value of Medicaid to beneficiaries is well below the program's cost.

To study how health insurance affects consumption risk, we estimate the impact of Medicaid expansion on the distribution of household consumption. Like much work in the literature, our empirical strategy takes advantage of the uneven adoption of Medicaid expansion, comparing states that did and did not expand Medicaid under the ACA. We depart from this literature, however, in estimating not only average but distributional effects, measuring the impact of Medicaid expansion on the entire household consumption distribution, using quantile difference-in-differences and change-in-changes methods (Athey and Imbens, 2006). These methods allow us to recover the entire counterfactual distribution of consumption in the absence of expansion. With that distribution, and an assumed utility function, we can calculate the insurance value of Medicaid expansion, defined as the willingness to pay for the expansion in excess of the increase in average consumption (Finkelstein, Hendren and Luttmer, 2019).

To measure consumption, we use the Consumer Expenditure Survey. A concern of using consumption data, however, is measurement error—individuals remember different categories with varying accuracy, and display an increasing tendency to report "zero" for many categories (Bee, Meyer and Sullivan, 2015). Thus, we use the subset of "well-measured consumption" (Meyer and Sullivan, 2023), which is comprised of regular, recurring expenses rather than one-off purchases, and shows less vulnerability to measurement error. Following Kaestner et al. (2017), we focus on low-education and non-elderly adults who are most likely to be eligible for Medicaid under the ACA expansion, and we limit the analysis to 2008-2019 to avoid complications from COVID-19.

Our main finding is that the estimated impact of Medicaid expansion on consumption risk is small and, although somewhat uncertain, precise enough to rule out large benefits. Specifically, looking across the consumption distribution, we estimate that Medicaid expansion increases consumption by a small, insignificant, and roughly constant effect between the 15th and 85th percentiles of risk protection among likely Medicaid eligible households. We find no increase at the bottom of the distribution and a small, noisy decrease at the top of the distribution. We quantify the implied insurance value under an assumed expected utility function. Our point estimates are generally small and statistically insignificant. The 95 percent confidence intervals rule out two important benchmarks: we find that the insurance value of Medicaid expansion is small relative to the government costs of providing that coverage, and also small relative to the health benefits of that coverage. When we focus on especially low-income groups, we find larger but much noisier estimates.

These findings contribute primarily to two stfands of the literature. First, a large literature examines the impact of the ACA on people and firms. Guth et al. (2020) provide a comprehensive review of the household-side effects of ACA. The ACA and especially the Medicaid expansions increased Medicaid coverage and all-source health insurance (e.g. Miller and Wherry (2017); Kaestner et al. (2017); Frean, Gruber and Sommers (2017). Greater coverage, has in turn, increased health care utilization, including primary care, preventive care, behavioral care, and, possibly, emergency department use (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)). Recent work has found important health improvements, and in particular mortality reductions, from Medicaid expansion (Miller et al., 2021; Wyse and Meyer, 2023) and ACA Marketplace coverage (Goldin, Lurie and McCubbin, 2021). The ACA has also improved financial well-being, reducing bankruptcy, medical debt, and mortgage delinquency, and by improving credit scores and terms of credit (Gross and Notowidigdo, 2011; Mazumder and Miller, 2016; Hu et al., 2018; Brevoort, Grodzicki and Hackmann, 2020; Gallagher et al., 2020; Dodini, 2023). Our primary contribution to this literature is to quantify the insurance value of the ACA's Medicaid expansions, a theoretically important benefit of the ACA, which has not yet been studied. In fact few papers study any consumption impacts, only Levy, Buchmueller and Nikpay (2019), who estimate the impact of Medicaid expansion on average consumption, but do not estimate its distributional impact or quantify its insurance value.

The parallel effect of the ACA on firms is important for the context of our results. The ACA provided a financial windfall for health care providers, especially charity care providers (Duggan, Gupta and Jackson, 2022; Dranove, Garthwaite and Ody, 2016), and even pre-ACA, insurance gains in a county improved hospital finances Garthwaite, Gross and Notowidigdo (2018). That is, even in the absence of Medicaid expansions, many low-income households were receiving emergency care at low or no cost to them; Medicaid changed how this care was financed, shifting the burden from providers to the government. Our results imply, however, that Medicaid expansion did not change the likelihood that a health emergency results in low consumption.

Second, we contribute to a smaller literature measuring the insurance value of health insurance. Finkelstein, Hendren and Luttmer (2019) develops a framework for measuring this value and applies it to the Oregon Health Insurance Experiment, accounting for benefits not only from reduced consumption risk but also improved health and greater consumption levels. That paper does not however directly estimate the impact of Medicaid on consumption risk. The authors find that the value to beneficiaries is smaller than the program's cost, with an insurance value that is sensitive to the exact method. Also closely related is recent work by Lockwood (2022), which argues that the insurance value of health insurance is low or negative, because of the interaction between health insurance and other aspects of the social safety net, especially uncompensated care. While that paper studies a wide variety of insurance programs, not just Medicaid, it does not use exogenous variation in these programs to identify their impact on the consumption distribution. Finally, Dodini (2023) estimates the impact of the ACA's premium and cost-sharing subsidies (for private coverage, not Medicaid) on the distribution of financial outcomes such as debt and bankruptcy, and uses the distributional estimates along with a utility function to obtain an insurance value. That paper does not, however, estimate impacts on consumption. Relative to these papers, our contributions are two-fold. First, we provide the first direct evidence on how exogenous variation in health insurance affects the consumption distribution. Second, we estimate the insurance value of Medicaid expansion in particular.[1]

The rest of this paper proceeds as follows. Section 2 details a theoretical basis for our estimation; Section 3 explain the empirical strategy; Section 4 overviews the data; Section 5 presents the results; and Section 6 concludes.

2 The Insurance Value of Medicaid Expansions

2.1 Model

We develop a simple model that we can take to the data to infer the insurance value of Medicaid. The model also highlights a key channel by which Medicaid may provide low insurance value despite improving financial health: if it crowds out uncompensated care.

[1]Also relevant are papers by Dague (2014) and Finkelstein, Hendren and Shepard (2019), which use enrollment choices to estimate the value of subsidized insurance coverage to enrollees, and find that a sizable fraction of beneficiaries have low willingness-to-pay. Relative to our approach, these papers do not specifically estimate the insurance value of this coverage, but require weaker assumptions, such as revealed preference rather than a specific utility function.

Consider two economies indexed by $m \in \{0, 1\}$, $m = 0$ corresponds to Medicaid non-expansion and $m = 1$ to Medicaid expansion. There is a mass of households indexed by i which would be eligible for Medicaid if $m = 1$. Households earn income y_i and may experience health shocks, resulting in health costs h_i. Health costs may be paid by the households themselves or by a third party. Third party transfers include private insurance, government insurance (in the Medicaid state of the world), and payment by the health care providers themselves, in the form of charity care or bad debt.

We assume households do not borrow or save, so

$$c_i = y_i - h_i + T_m(y_i, h_i).$$

Here $T_m(y_i, h_i)$ is an economy-specific transfer that captures, in a reduced form way, the role of third-party financing of health costs. If a household finances its health care entirely on its own, $T_m(y_i, h_i) = 0$. For example, an economy with full insurance against health costs would have $T_m(y_i, h_i) = h_i$.

We let the transfers depend on both income and health costs to reflect the role of uncompensated care and social insurance. For example, many hospital systems offer cost assistance programs for low income households (Adams et al., 2022), and households with fewer assets discharge more of their medical debts via bankruptcy (Mahoney, 2015), so transfers vary with income.

To focus on the impact of Medicaid on consumption risk, we assume that Medicaid expansion affects the economy exclusively via the transfer function. In particular, we abstract away from the impact that Medicaid expansion has on income and health costs; we discuss these factors below.

This simplification implies that the expected utility of a would-be Medicaid-eligible household in economy m can be written as

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

where $F_m(c)$ is the distribution of consumption in economy m (among Medicaid-eligible households), and u is the von-Neumann-Morgensterm utility function. We take expected utility over the entire consumption distribution (among eligible households), so consumption variability reflects both risk from uninsured medical expenses as well as permanent differences in income. EU_m therefore represents expected utility from behind the veil of ignorance, among the Medicaid-eligible population, and improvements in EU_m represent improvements in "ex ante welfare," in the sense of Hendren (2021).

Following Finkelstein, Hendren and Luttmer (2019), we define the willingness-to-pay for Medicaid expansion γ as the amount of consumption an individual would give up in the economy with Medicaid to be just indifferent to the economy without Medicaid. That is, γ solves

$$\int_c u(c+\gamma) dF_1(c) = \int_c u(c) dF_0(c). \tag{1}$$

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

$$\pi = \gamma - \left(\int_c c dF_1(c) - \int_c c dF_0(c) \right). \tag{2}$$

The second term in Equation (2) is the difference in expected consumption between the two economies, roughly equivalent to the transfer component of the value of Medicaid (Finkelstein, Hendren and Luttmer, 2019) because it represents the net transfer to Medicaid eligible households.[2]

[2]Our model differs from Finkelstein, Hendren and Luttmer (2019) in that we do not account for the health gains from Medicaid. This simplifies our estimation of insurance value, but it also means that the mean consumption gain in our model is not the same as the transfer component in Finkelstein, Hendren and

While the transfer component is not a net welfare gain—it represents a cost borne elsewhere in the economy—the insurance premium is a welfare benefit from Medicaid expansion. This benefit arises for the traditional reason that insurance improves welfare: risk averse households are happy to trade off reduced average consumption for a reduction in consumption risk, i.e., an improvement in the lower tail of the consumption distribution. Thus the insurance value of Medicaid depends on how Medicaid expansion affects the consumption distribution, and in particular on whether expansion increases the lower tail of consumption.

While it is natural to think that Medicaid improves the lower tail of the consumption distribution and so has a large insurance value, this need not be true. For example, even in the absence of Medicaid, would-be Medicaid beneficiaries experiencing health cost shocks may have access to medical debt or charity care, they may have informal insurance through friends and family, and they may have needs that are sufficiently basic that a financial shock moves them to become eligible for other forms of government benefits.

Our empirical goal, therefore, is to estimate the impact of Medicaid on the consumption distribution and recover the implied insurance value of Medicaid. We estimate the distribution of consumption F_1 in Medicaid expansion states, and recover the counterfactual distribution F_0. Armed with the estimates $\hat{F}_1$ and $\hat{F}_0$, we can calculate $\hat{\gamma}$ and $\hat{\pi}$ for any assumed utility function.

2.2 Discussion

Insurance value vs. overall values. Our focus is on estimating the *insurance value* of Medicaid, and we therefore abstract from several potentially important channels whereby Medicaid may generate value. The clearest such channel is improvements in beneficiary health; a growing body of evidence documents improvements in physical and mental health (e.g. Finkelstein et al. (2012)) and, for the Medicaid expansions, reductions in mortal-

Luttmer (2019).

ity (Miller et al., 2021; Wyse and Meyer, 2023). Improvements in financial health may also matter for beneficiaries in ways that are not captured by the consumption distribution alone, for example via reduced stress or hassle cost of interacting with debt collectors. Our abstraction from these outcomes does not suggest that they are unimportant, but rather to estimate in a relatively clear way the importance of one channel: consumption insurance.

Behavioral responses to Medicaid. Our model abstracts from two potentially important behavioral responses to Medicaid: income and health care. The main reason we abstract from these features is to simplify the model and focus on the insurance value. Prior evidence suggests that the ACA's Medicaid expansions have not generated substantial labor supply responses (Kaestner et al., 2017; Leung and Mas, 2018).[3]

The ACA has however, had a large impact on health care utilization. While early evidence found little impact on overnight hospitalizations or doctor visits (Miller and Wherry, 2017), more recent work reviewed by Guth et al. (2020) has found positive impacts on a range of preventive care measures, including for example immunizations and cancer screenings (Simon, Soni and Cawley, 2017) as well as tobacco cessation (Maclean, Pesko and Hill, 2019; Cotti, Nesson and Tefft, 2019). While we do not directly account for these changes in health care utilization in our model, they may affect welfare through improved health, and we view the value of these health improvements as an important benchmark for assessing the insurance value of Medicaid expansions.

Take-up vs. eligibility. Our model focuses on a Medicaid eligible population, but not everyone eligible for Medicaid takes up coverage (e.g. Decker, Abdus and Lipton (2022)). This distinction is important for many outcomes, where we would expect Medicaid to have no impact on people who are eligibile but not enrolled. However, the risk-reducing benefits of Medicaid extends even to these eligible but not-yet-enrolled people, because Medicaid

[3]The ACA has generated substantial responses in *taxable income* as households report income levels to qualify for subsidized coverage (Kucko, Rinz and Solow, 2018; Heim et al., 2021), but this appears to be changes in reporting rather than earnings. Evidence on pre-ACA Medicaid coverage is mixed; for example, Garthwaite, Gross and Notowidigdo (2014) finds large employment effects of Medicaid disenrollment in Tennessee, but Dague, DeLeire and Leininger (2017) finds smaller effects in Wisconsin.

offers *retroactive coverage*.[4] That is, if a person is eligible but not enrolled, and suffers a hospitalization, the hospital typically enrolls that person, covering her hospitalization and reducing her exposure to out-of-pocket payments. Medicaid provides insurance even for people who have not previously signed up. We therefore focus on the consumption distribution among eligibles, rather than among those taking up coverage.

Saving and borrowing. Our model abstracts from saving and borrowing, but one response to health expenditure risk is precautionary savings. Households may want to maintain a buffer stock of assets to finance sudden medical expenses. Indeed this mechanism seems important for post-retirement savings (De Nardi, French and Jones, 2010). If present, this savings represents a welfare loss because it results in permanently lower consumption. If precautionary savings is important, then the transfer component we estimate (i.e. the impact on mean consumption) in part represents a welfare gain. We suspect, however, that precautionary savings is relatively unimportant in our context, because the low-income population that we study generally has limited financial resources and little scope for saving.[5] Indeed, Gallagher et al. (2020) finds no average effect of Medicaid expansion on savings.

Another possible response is that households borrow to finance their medical spending. In fact such a response is unlikely because health care spending is almost never pre-paid; instead, providers bill patients and eventually collect. Uncollected payments become debt, and indeed medical debt is a major liability on household balance sheets (e.g. Kluender et al. (2021)). Although we do not directly model medical debt, we indirectly capture its effects. In particular, to the extent that Medicaid expansion affects medical debt and medical debt affects consumption—because households pay down their debt—then differences in medical

[4] See e.g. `https://www.kff.org/medicaid/issue-brief/medicaid-retroactive-coverage-waivers-implications-for-beneficiaries-providers-and-states/`.

[5] Lusardi, Schneider and Tufano (2011) defines a financial fragility metric of whether a household could come up with $2,000 in 30 days. The paper finds that one-quarter of Americans could not, and an additional 20% would need to sell possessions or obtain payday loans to deal with this level of financial shock that is plausible for an insured medical need. In a separate paper, Pfeffer, Danziger and Schoeni (2013) studies the wealth effects of the Great Recession and finds that the least advantaged households had relatively little savings even prior to the recession.

debt will be reflected in differences in the consumption distribution. If medical debt does not affect the consumption distribution (for example because it is forgiven or discharged in bankruptcy) then our model would correctly say that such debt is not relevant for the insurance value of Medicaid.[6]

3 Empirical Strategy

Our goal is to estimate the insurance value of Medicaid expansions, the empirical analog of π. This parameter is determined by the consumption distribution under Medicaid expansion, $F_1(c)$ and the counterfactual distribution, $F_0(c)$. We recover these objects using data on consumption—described below—and a research design taking advantage of the uneven adoption of Medicaid expansion across states.

3.1 Recovering the Counterfactual Consumption Distribution with a Medicaid Expansion Design

While the Patient Protection and Affordable Care Act (ACA) originally directed all states to expand their Medicaid program to cover all adults with income up to 138 percent 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 (including DC) expanded their Medicaid programs. Five had expanded earlier than 2014, an additional 8 expanded between 2015 and 2019, and 17 have not expanded as of 2019, the end of our sample period. Our empirical strategy takes advantage of this uneven expansion, following a large literature investigating the effects of Medicaid expansion.[7] Figure 1 shows the timing of state expansion adoptions. While the expansion states are spread throughout

[6]Of course it may be welfare relevant for other reasons, for example psychological stress or the hassle of debt collectors.

[7]Within this literature there is some inconsistency on how to classify and date state expansion decisions. We follow the classification in Miller, Johnson and Wherry (2021).

the country, the never-expanders are concentrated in the Southeast and the Great Plains.

While much of this literature focuses on the average impact of Medicaid expansion, our interest is in the distributional impacts. To study average impacts, the prior literature has largely estimated difference-in-differences (DID) models, looking at the difference in mean outcomes between expansion states and non-expansion states, before and after expansion. To measure the distributional impact of Medicaid expansion, we turn to two natural alternatives to DID. The first alternative is quantile regression (in particular, quantile DID or QDID), which applies DID's parallel trends assumption to quantiles rather than means. However, as Athey and Imbens (2006) explains, this application has some unappealing features. They propose an alternative method, change-in-changes (CIC). We use both approaches and find very similar results.

To explain our approach, we introduce some notation. Let $t = 1, 2$ denote time periods ($t = 1$ prior to expansion, $t = 2$ post expansion), 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 without and with expansion.

The usual focus in the Medicaid expansion literature is on the average effect of expansion among the expands 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 potential outcome, $E\left[C_{i,2}(0) \middle| D_{s(i)} = 1\right]$, by invoking a parallel trend 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:

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

Our focus differs from the usual one in two ways. First, rather than mean effects, we are interested in distributional effects. Second, rather than causal contrasts, we are directly interested in the distribution of untreated potential outcomes, because (as we explain below) this distribution is key to recovering the insurance value of Medicaid.

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 percentile of consumption for expansion group g in period t and state of the world 0 (non-expansion) or 1 (expansion). Then for all quantiles q, we assume

$$P_{q12}(0)-P_{q11}(0)=P_{q02}(0)-P_{q01}(0) \tag{3}$$

We call this assumption parallel trends in percentiles. 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, ever expansion, and their interaction, as in the standard DID model; hence the name "quantile" DID. We estimate the QDID model at a set of quantiles denoted by $p_1,\ldots,p_N$. In practice we estimate impacts on 20 evenly spaced quantile bins, starting at the 2.5th percentile. We perform inference via the bootstrap, described below.

The parallel trends in quantiles assumption makes QDID a simple and intuitive application of DID ideas to quantiles rather than means. However, the cost of this simplicity is 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 cross-group heterogeneity.[8] Athey and Imbens (2006). Athey and Imbens (2006) therefore develops an alternative approach to extending DID to distributional impacts, change-in-changes, which avoids these unappealing features of QDID.

We estimate CIC models as well as QDID and, reassuringly, find very similar estimates with both. To implement CIC, we recover 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 ours does. We therefore coarsen our outcome slightly, rounding consumption amounts to the nearest \$1. We use the CIC implementation developed by Kranker (2016). Again we use the bootstrap for inference.

3.2 From consumption distributions to insurance value

Our CIC and QDID estimates give us the factual and counterfactual consumption distribution in the post period in Medicaid expansion states. Denote these distributions $\hat{F}_1$ and $\hat{F}_0$. Our CIC and QDID estimation recover the percentile $p_1, ..., .p_N$. Denote the value of consumption under distribution F_j at these quantiles as c_{im}, so that for example c_{10} is the p_1 percentile of consumption in the non-expansion state of the world.

To go from these percentiles to the insurance value of Medicaid, we assume that the consumption distribution has N points of support, and we assume that utility is of the

[8]An additional unappealing feature of additive separability is that it is also generally not invariant to alternative scalings of outcomes; if the level of consumption satisfies parallel trends in quantiles, which is in additive, the log of consumption does not. Of course DID also suffers from this problem (e.g. Roth and Sant'Anna (2023)).

constant relative risk aversion class with risk aversion ρ (following Finkelstein, Hendren and Luttmer (2019)). We can therefore estimate expected utility in each state of the world:

$$\hat{EU}_m = \sum_{i=1}^{20} (p_i - p_{i-1}) \frac{c_{im}^{1-\rho}}{1-\rho}. \tag{4}$$

To recover $\hat{\gamma}$, we find the consumption offset that makes expected utility equal between both states of the world. That is, $\hat{\gamma}$ solves

$$\sum_{i=1}^{N} (p_i - p_{i-1}) \frac{(c_{i1} - \hat{\gamma})^{1-\rho}}{1-\rho} = \sum_{i=1}^{N} (p_i - p_{i-1}) \frac{c_{i0}^{1-\rho}}{1-\rho}. \tag{5}$$

Finally, to estimate insurance value π, we subtract off the difference in expected consumption:

$$\hat{\pi} = \hat{\gamma} - \sum_{i=1}^{N} (p_i - p_{i-1}) (c_{i1} - c_{i0}). \tag{6}$$

We report estimates of $\hat{\pi}$ for different levels of ρ. In implementation we have $N = 20$ uniformly spaced percentiles.

3.3 Accounting for staggered adoption

Up until now we have assumed that all states expand Medicaid at the same time, or not at all. In fact states expanded at different times. Pooling different timing groups 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), and dropping late expanders reduces power. To account for the 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, then second aggregate the estimates together to obtain an overall average effect.

Our first step is to estimate parameters separately by expansion timing. We divide the data up into timing groups, defined by the year each state adopted the Medicaid expansion. Let g index such groups, and let $\hat{F}_m^g$ refer to timing-group specific estimated distributions (factual and counterfactual). We estimate $\hat{F}_1^g$ directly from the data and we recover $\hat{F}_0^g$ using QDID or CIC, using never expanders (as of 2020) as the control group for timing group g. In all timing groups our data start in 2008 and end in 2019. We use this year range to have the widest symmetric window for 2014 expanders, but avoid too much the influence of the COVID-19. From $\hat{F}_0^g$ and $\hat{F}_1^g$ we recover $\hat{\gamma}^g$ and $\hat{\pi}^g$ using Equations (5) and (6).

Our second step aggregates the timing group specific parameters to overall averages. To do so, we average them together, weighting each timing group by the share of the post-expansion population it contains. That is,

$$w^g = \frac{\sum_i w_{it} \cdot 1\{\text{timing} = g\} \cdot 1\{g \leq t\}}{\sum_{g'} \sum_i w_{it} \cdot 1\{\text{timing} = g'\} \cdot 1\{g' \leq t\}}, \tag{7}$$

$$\hat{F}_m = \sum_g w^g \hat{F}_m^g \tag{8}$$

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

where w_{it} is the survey weight observation i in year t. These weighted averages are analogous to average treatment on the treated parameter. For example, $\hat{\pi}$ is the weighted average insurance value experienced post-expansion, accounting for the uneven length of the post-period and the (potentially) unequal insurance value across timing groups.

3.4 Inference via the block bootstrap

We conduct inference via the block bootstrap, resampling states. We use the bootstrap because our ultimate object of interest, π, is defined as an implicit function of the estimated objects, 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 distributions and insurance values, for the resampled timing groups. Third, we recalculate the timing group weights within the bootstrap iteration, and obtain bootstrap-specific estimates $\hat{\pi}_b$ and $\hat{F}_{m,b}$.[9] We repeat this procedure 1,000 times. We then report bootstrap confidence intervals as the 2.5th and 97.5th percentile of the bootstrap estimates.

4 Data

Our analysis requires well-measured, comprehensive consumption data, with state identifiers, covering much or all of the nation, spanning years before and after 2014. Several data sources meet some of these requirements but few meet all of them. Consumption data derived from credit and debit card transactions, as in e.g. Baker (2018), does not cover all consumption nor all households, and in particular may miss low-income households most likely to benefit from Medicaid expansions. Household scanner data such as the Nielsen Household Panel provides roughly representative, high-quality data, but mainly measures food and grocery consumption. Likewise, the Current Population Survey asks some consumption-related questions in its food security module, but lacks comprehensive consumption measures. The Panel Study on Income Dynamics has national coverage and comprehensive consumption measures, but is too small to estimate state-level models. We therefore turn to the Consumer Expenditure Survey (CEX), a long-running, nationally representative survey of household expenditures, for our primary analysis sample.[10]

The main limitation of the CEX is that overall consumption expenditures are measured with error, primarily because households have difficulty recalling all of their purchases. Such measurement error may be especially problematic in the tails of the distribution, which are

[9] Our bootstrap procedure does not stratify on timing group, so some bootstrap samples exclude some timing groups. Thus 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.

[10] Although the CEX is nationally representative, its sampling scheme does not result in every state being included in the sample.

our focus. To minimize the role of measurement error, our analysis focuses on "well-measured consumption" (Meyer and Sullivan, 2022), developed by Meyer and Sullivan (2023), although we also investigate the effects of Medicaid expansion on overall consumption. To construct well-measured consumption, Meyer and Sullivan (2023) take the subcategories of consumption for which under-reporting seems lowest (in the sense that for these subcategories, the implied aggregate spending aligns with aggregate spending in the national accounts (Bee, Meyer and Sullivan, 2015)). They exclude expenditures that are closer to investments than consumption—health-related ones, education, and pension and outlays for retirement. To translate spending to consumption, Meyer and Sullivan (2023) converts expenditures such as new vehicle purchases and housing expenditures into flows of vehicle services and housing services.

The resulting well-measured consumption consists of food at home, the flow value of housing (rent and the rental value of owner-occupied housing), utilities, the flow value of owned vehicles, gas and motor oil expenses, and communications. Intuitively, well-measured consumption is common, recurring consumption, plus the flow value of very large durable expenditures (housing and car purchases). The components of well-measured consumption are measured in the Quarterly Interview Survey, so our consumption measures are quarterly. We construct inflation-adjusted measures using the CPI-U and, following Meyer and Sullivan (2023), we adjust for differences in household composition using the equivalence scale suggested by Constance F. Citro and Robert T. Michael (1995), $(A + 0.7K)^{0.7}$ where A, K are the number of adults and children in the household. With this adjustment, we interpret all effects on a per-person basis (rather than per household).[11]

Because the components of well-measured consumption are either difficult to adjust or inferior goods—food at home—it might seem unlikely that well-measured consumption would respond to uninsured health expenditure shocks or to Medicaid expansion. Our results

[11] Despite using Meyer and Sullivan's data, our implementation of well-measured consumption differs from theirs in that we include communication expenses, which they drop because of changes in communication technology over the 60-year period they study.

using all consumption are reassuring on this point. We note that this concern assumes that households pay for unanticipated health care shocks exclusively through a large one-time decrease in consumption. However, medical debt is an important financing mechanism. Paying down the principal and interest on medical debt can reduce consumption for several years, giving households time to adjust not only luxuries but also their food, housing, or transportation consumption.

We limit our sample to exclude households unlikely to benefit from Medicaid expansion. Our primary analysis sample is constructed using the inclusion criteria of Kaestner et al. (2017). We restrict the sample to people age 22-64, without any college education. Restricting the sample based on education rather than income *per se* is helpful for avoiding potential sample selection bias that arises because income may respond to Medicaid expansion. Throughout, we use 2008-2019, stopping our analysis before 2020 to avoid complexities in the analysis arising from the COVID-19 pandemic, which greatly affected both health needs and consumption patterns. The analysis sample consists of 58,314 observations.

Table 1 reports summary statistics on income and consumption for our sample, separately by state expansion status. Income and consumption are higher in expansion states than non-expansion states. In our sample of less-educated households, real annual income averages about $25,000 (after adjusting by the equivalence scale). Quarterly consumption is about $4,250, and well-measured consumption is about 80 percent of overall consumption. The most important component of well-measured consumption is housing, followed by food at home and gas and motor oil.

5 Results

We begin by showing trends in the distribution of well-measured consumption in Figure 2. The figure plots percentiles of well-measured consumption, from 2008 to 2019, for non-expansion states and for 2014 expansion states.[12] Several trends are apparent in the figure. First, even in our relatively homogeneous sample, consumption is fairly dispersed, with a 90/10 ratio of about 3. Second, consumption is higher at each percentile in expansion states than in non-expansion states. Third, consumption fell throughout the distribution in Great Recession, although it fell more sharply at higher percentiles. Consumption does not reach 2008 levels until 2019.

Most critically, prior to Medicaid expansion, at each percentile, consumption moves roughly in parallel for expansion and non-expansion states. High percentiles of consumption experienced a sharp decline in consumption from 2008 to 2009 in both expansion and non-expansion states. Low percentiles experienced more modest declines. The figure also shows some divergence in consumption percentiles after Medicaid expansion. The median, 75th, and 90th percentiles all appear to diverge somewhat, falling in non-expansion states but holding steadier in expansion states. These results suggest that Medicaid expansion increases consumption more in the middle of the distribution than in the lower tail.

The estimates in Figure 3 confirm this impression: Medicaid expansion increases consumption in the middle of the distribution but not in the tails. The figure plots the estimated factual and counterfactual distributions (for expansion states, post expansion), as well as the estimated quantile treatment effects. The estimation approach underlying this uses all timing groups: we estimate distributional effects separately by timing group, and average the estimates. The figure shows positive impacts of expansion from about the 15th percentile to the 80th or 85th percentile. The estimates in the left tail are positive but smaller, and the estimate at the top of the distribution is slightly negative. Overall the average impact is an increase in quarterly consumption of $97. None of the estimates is statistically significant and neither is the overall mean effect. While this figure reports results obtained with QDID, CIC produces very similar results, as Appendix Figure A.5 shows.

[12] We focus on 2014 expansion states here to avoid the problems from pooling later adopters. We show analogous plots for later expanders (2015, 2016, 2017, and 2019) in Appendix Figures A.1-A.4. Note that no states is expanded in 2018.

We use the estimated consumption distributions, along with an assumed utility function, to estimate the insurance value of Medicaid expansion. We report these estimates in Table 2, along with their 95 percent and 50 percent confidence intervals. With a risk aversion parameter of 3, we estimate a small insurance value: \$22 per quarter or \$89 per year. With risk aversion of 1, we have essentially zero insurance value, but with a high enough risk aversion, we find a large and highly uncertain insurance value. We focus on the estimates for lower values of risk aversion because Chetty (2006) reports that labor supply choices rule out values of risk aversion above 2, and point towards values closer to 1.[13]

Overall we estimate low insurance values. Of course, this low insurance value does not necessarily imply that Medicaid has low or no value to recipients. Recall that this value reflects the *insurance* component of the value of Medicaid, not the transfer component and not non-consumption components (such as health improvements or avoided stress from bankruptcy), which are outside our model. Indeed, Medicaid expansion increases consumption throughout much of the distribution, resulting in a (quarterly) transfer value of about \$100 (the mean consumption effect reported in Figure 3). Intuitively, the reason we estimate a low *insurance* value is that the consumption increase is larger in the middle of the distribution and not in the lower tail. The insurance value is nonetheless positive because the consumption effect is actually slightly negative at the top of the distribution. Like the individual percentile estimates, our estimated insurance value is statistically insignificant, with an annual 95 percent confidence interval of (-160, 391).

Despite this wide confidence interval, our estimated insurance value is informative, and in particular rules out the possibility that the insurance value of Medicaid represents a large share of its gross costs. Kaiser Family Foundation (2019) report annual costs of \$5,225 per newly insured beneficiary under the ACA. This estimate is not directly comparable to ours, however, because we estimate the insurance value per eligible individual, rather than the

[13] We report insurance values for risk aversion parameters of 3, 1, and 5 so that our results can be compared to Finkelstein, Hendren and Luttmer (2019), the mostly closely related prior work.

annual cost per covered individual. To make them comparable, we need to scale down the costs by the Medicaid take-up rate, which Decker, Abdus and Lipton (2022) estimate to be 44 to 46 percent. Putting these numbers together, our estimates imply that the insurance value of Medicaid is not more than 18 percent of the direct costs.[14]

Our estimates also imply that the insurance value of Medicaid is small relative to its mortality-reducing benefits. Miller et al. (2021) and Wyse and Meyer (2023) find substantial mortality reductions from Medicaid expansions. The estimates in Wyse and Meyer (2023) imply that Medicaid expansions increase life expectancy by 0.03 years per beneficiary.[15] Valuing a statistical life-year, conservatively, at $100,000 implies a mortality-reducing benefit of roughly $1,300 per eligible—about four times as large as the upper bound of our confidence interval.

A natural benchmark for our estimates is the insurance value reported by Finkelstein, Hendren and Luttmer (2019). Using different assumptions about household optimization, Finkelstein, Hendren and Luttmer (2019) estimates annual insurance values that range from $133 to $1,106 per beneficiary. Our estimates and confidence intervals are consistent with the smaller but not the larger of these estimates (even after scaling by a take-up rate of roughly 50 percent and of course the quarter-to-annual adjustment). Finkelstein, Hendren and Luttmer (2019) finds that the transfer component of Medicaid—the annual value to beneficiaries of the extra consumption, health care, and health—is $570-$863.[16] Our estimated transfer component is consistent with this number, and we cannot rule out that the insurance value to beneficiaries is equal to this transfer component. Importantly, this transfer component is a small share of the total expenditures on Medicaid, because much of these expenditures go to recipients other than beneficiaries (including providers), and because some spending is

[14] That is, the direct costs are on the order of $.44 \times 5225$ or $2,300 per eligible per year.

[15] Wyse and Meyer (2023) reports that expansions increased coverage by 28.7 million coverage-years and reduced mortality by 831,890 life-years.

[16] This transfer component includes the health benefits of Medicaid but it does not include any mortality benefits, because Finkelstein, Hendren and Luttmer (2019) bases their calculations on data from the Oregon Health Insurance experiment, which did not find a mortality impact (Finkelstein et al., 2012).

valued below cost (because of moral hazard).

Our finding of a small insurance value is not especially sensitive to the estimator or consumption measure, as the results in Table 2 show. Working with overall consumption instead of well-measured consumption (panel B) produces similar point estimates but wider confidence intervals, consistent with the greater noise in overall consumption than well-measured consumption. Change-in-changes instead of QDID produces very similar results (panel C).

In the results so far, we have considered a sample that is likely to benefit from Medicaid expansion, but not everyone in our sample is necessarily eligible for Medicaid. In Panels D and E, we turn to low-income samples, where most or all people would be eligible for Medicaid in expansion states. In these samples, we find larger point estimates, with a (quarterly) insurance value of about \$100. However the estimates are much less precise, ranging from -\$20 to \$300 (or wider). Thus these estimates are largely uninformative: they are consistent with our main estimates, with zero or negative insurance value, and with a substantial insurance value. Although our main sample includes people not currently eligible, we nonetheless believe our main estimates are informative about the insurance value of Medicaid. This is because Medicaid eligibility is fluid as well as retroactive, and indeed people regularly churn in and out of Medicaid (Einav and Finkelstein, 2023). Thus people with moderately high current income may have previously been eligible for Medicaid, and should they be hospitalized, we might expect a large fall in income (Dobkin et al., 2018), triggering Medicaid eligibility.

6 Conclusion

Using household-level data on well-measured expenditures, we investigate the impact of Medicaid expansion on the consumption distribution. While our point estimates suggest that Medicaid expansion increased consumption in the middle of the distribution, this estimate is imprecise, and we find no evidence of increased consumption at the bottom of the distribution. Combining these distributional estimates with a utility function, we estimate an insurance value of Medicaid expansion of \$89 per eligible person per year. This point estimate is small relative to the overall government expenditures on the expansion population, roughly \$2,500 per eligible person per year, and it is small relative to other benefits of the Medicaid expansion, especially its mortality-reducing benefits. Although our estimated insurance value is uncertain, our 95 percent confidence intervals nonetheless let us rule out that the insurance value is a substantial fraction of either the overall outlays or the value from reduced mortality. Overall, our results suggest that the consumption risk-reducing benefits of Medicaid expansion are not the main source of its value to beneficiaries.

References

Adams, Alyce, Raymond Kluender, Neale Mahoney, Jinglin Wang, Francis Wong, and Wesley Yin. 2022. "The Impact of Financial Assistance Programs on Health Care Utilization: Evidence from Kaiser Permanente." *American Economic Review: Insights*, 4(3): 389–407.

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.

Bee, Adam, Bruce D Meyer, and James X Sullivan. 2015. "The Validity of Consumption Data." *Improving the Measurement of Consumer Expenditures*, 74: 204.

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.

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

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

Cotti, Chad, Erik Nesson, and Nathan Tefft. 2019. "Impacts of the ACA Medicaid expansion on health behaviors: Evidence from household panel data." *Health Economics*, 28(2): 219–244.

Dague, Laura. 2014. "The effect of Medicaid premiums on enrollment: A regression discontinuity approach." *Journal of Health Economics*, 37: 1–12.

Dague, Laura, Thomas DeLeire, and Lindsey Leininger. 2017. "The effect of public insurance coverage for childless adults on labor supply." *American Economic Journal: Economic Policy*, 9(2): 124–154.

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.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J Notowidigdo. 2018. "The economic consequences of hospital admissions." *American Economic Review*, 108(2): 308–352.

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

Donohue, Julie M, Evan S Cole, Cara V James, Marian Jarlenski, Jamila D Michener, and Eric T Roberts. 2022. "The US Medicaid program: coverage, financing, reforms, and implications for health equity." *JAMA*, 328(11): 1085–1099.

Dranove, David, Craig Garthwaite, and Christopher Ody. 2016. "Uncompensated care decreased at hospitals in Medicaid expansion states but not at hospitals in nonexpansion states." *Health Affairs*, 35(8): 1471–1479.

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.

Einav, Liran, and Amy Finkelstein. 2023. "The risk of losing health insurance in the United States is large, and remained so after the Affordable Care Act." *Proceedings of the National Academy of Sciences*, 120(18): e2222100120.

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.

Finkelstein, Amy, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P Newhouse, Heidi Allen, Katherine Baicker, and Oregon Health Study Group. 2012. "The Oregon health insurance experiment: evidence from the first year." *The Quarterly Journal of Economics*, 127(3): 1057–1106.

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.

Gallagher, Emily A, Radhakrishnan Gopalan, Michal Grinstein-Weiss, and Jorge Sabat. 2020. "Medicaid and household savings behavior: New evidence from tax refunds." *Journal of Financial Economics*, 136(2): 523–546.

Garthwaite, Craig, Tal Gross, and Matthew J Notowidigdo. 2014. "Public health insurance, labor supply, and employment lock." *The Quarterly Journal of Economics*, 129(2): 653–696.

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, Rachel Garfield, Robin Rudowitz, et al. 2020. "The effects of Medicaid expansion under the ACA: updated findings from a literature review." *Kaiser Family Foundation*, 37(6): 944–50.

Heim, Bradley T, Gillian Hunter, Adam Isen, Ithai Z Lurie, and Shanthi P Ramnath. 2021. "Income responses to the affordable care act: Evidence from a premium tax credit notch." *Journal of Health Economics*, 76: 102396.

Hendren, Nathaniel. 2021. "Measuring ex ante welfare in insurance markets." *The Review of Economic Studies*, 88(3): 1193–1223.

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.

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.

Kaiser Family Foundation. 2019. "Medicaid Spending per Enrollee (Full or Partial)." https://www.kff.org/medicaid/state-indicator/medicaid-spending-per-enrollee.

Kaiser Family Foundation. 2024. "Health Insurance Coverage of Adults 19-64, 2008-2022." https://www.kff.org/other/state-indicator/adults-19-64. Last accessed January 30, 2024.

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

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.

Kucko, Kavan, Kevin Rinz, and Benjamin Solow. 2018. "Labor market effects of the Affordable Care Act: Evidence from a tax notch." *Available at SSRN 3161753*.

Leung, Pauline, and Alexandre Mas. 2018. "Employment effects of the affordable care act medicaid expansions." *Industrial Relations: A Journal of Economy and Society*, 57(2): 206–234.

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

Lockwood, Lee M. 2022. "Health Insurance and Consumption Risk."

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

Maclean, Johanna Catherine, Michael F Pesko, and Steven C Hill. 2019. "Public insurance expansions and smoking cessation medications." *Economic inquiry*, 57(4): 1798–1820.

Mahoney, Neale. 2015. "Bankruptcy as implicit health insurance." *American Economic Review*, 105(2): 710–746.

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. 2022. "Replication Data for: Consumption and Income Inequality in the US since the 1960s." Harvard Dataverse, https://doi.org/10.7910/DVN/587F9Z.

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.

Pfeffer, Fabian T, Sheldon Danziger, and Robert F Schoeni. 2013. "Wealth disparities before and after the Great Recession." *The Annals of the American Academy of Political and Social Science*, 650(1): 98–123.

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

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."

Figure 1: Timing of state Medicaid expansion adoptions

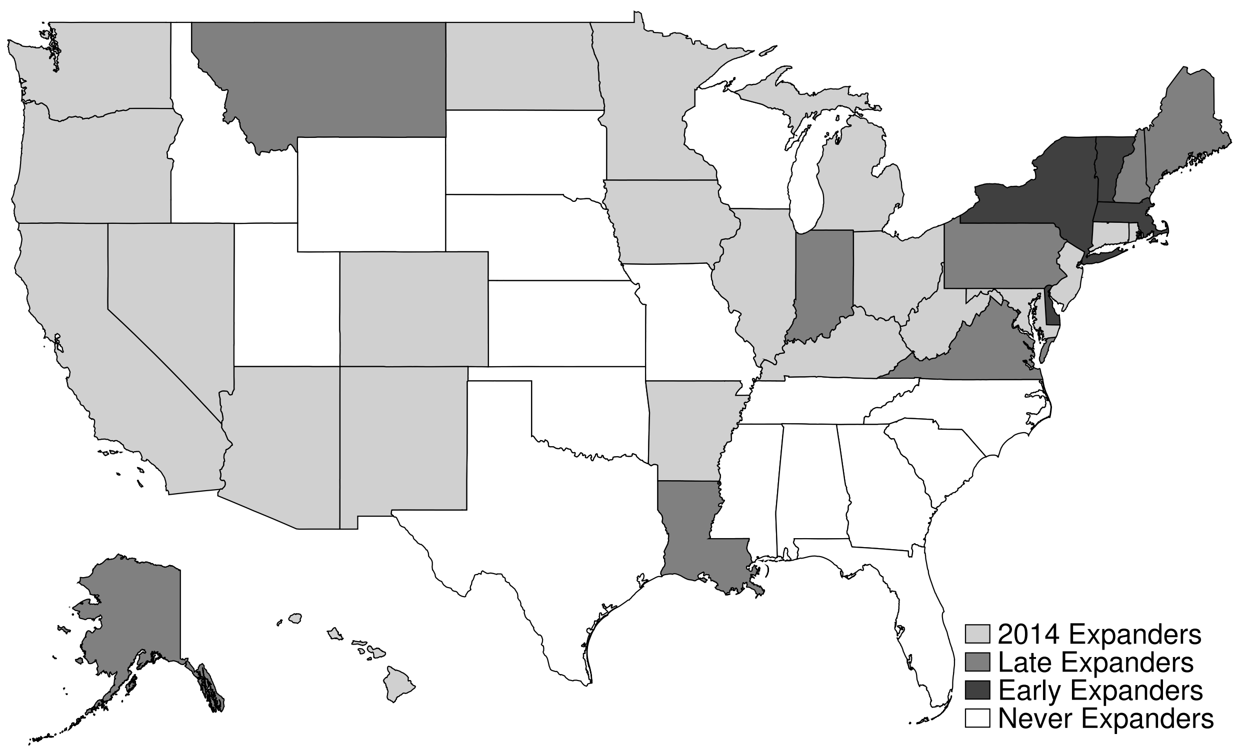

Notes: Figure indicates the timing of state Medicaid expansion decisions as of 2019. We use the classification in Miller, Johnson and Wherry (2021). Among the late expanders, Indiana, New Hampshire, and Pennsylvania adopted the expansion in 2015, Alaska and Montana in 2016, Louisiana in 2017, and Maine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

Figure 2: Parallel trends in percentiles of well-measured consumption

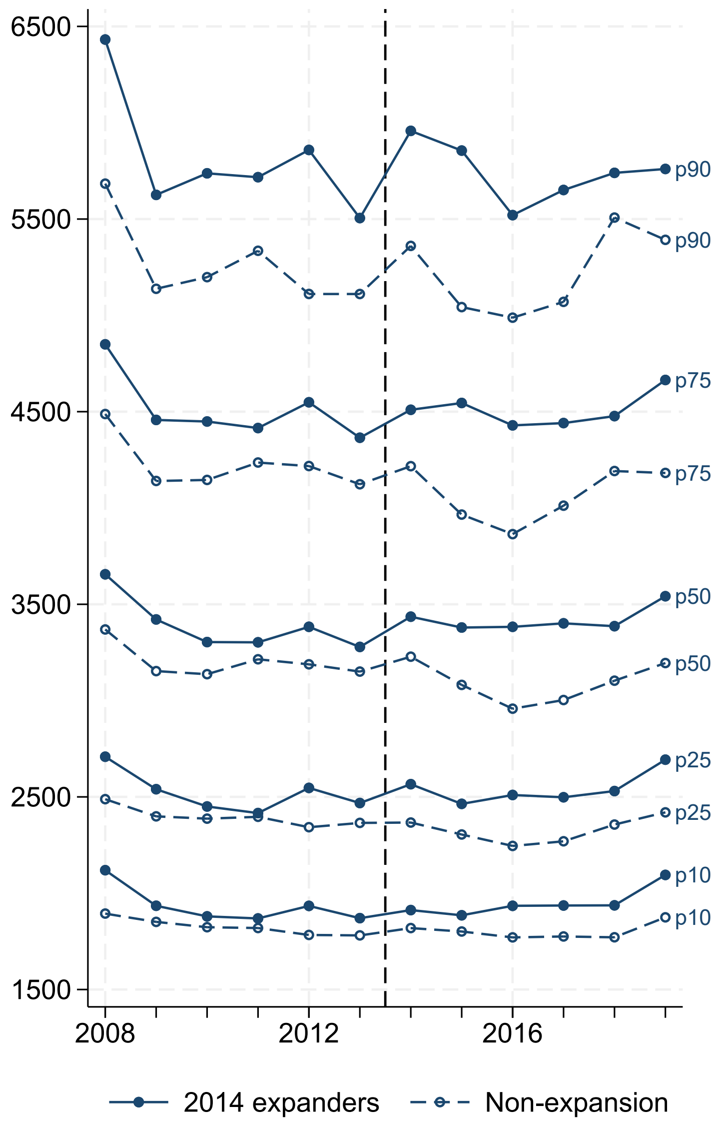

Notes: Figure plots the indicated percentiles of well-measured consumption, for 2014 expansion states (solid line) and non-expansion states (dashed line). Sample is defined in the notes to Table 1.

Figure 3: Impact of Medicaid expansion on the consumption distribution

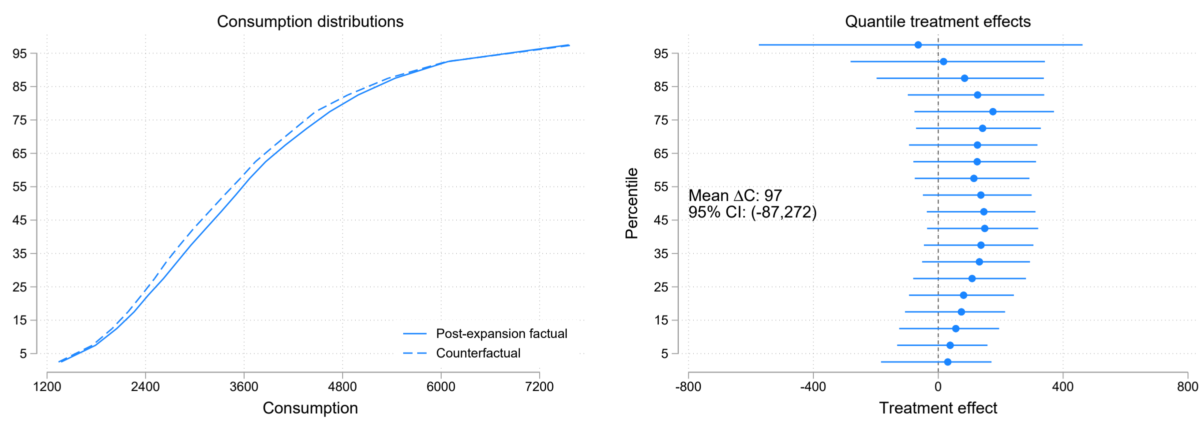

Notes: Figure plots the estimated factual and counterfactual consumption distribution (left panel) and the quantile treatment effects (right panel), as well as 95% confidence intervals, calculated via the bootstrap (resampling states). We estimate separate effects for each timing group and the figure reports the average effect. Distributional effects are estimated with QDID. Sample is defined in the notes to Table 1.

Table 1: Summary statistics on consumption

	Expansion states		Non-expansion states	
	Mean	(SD)	Mean	(SD)
A. Overall income and consumption				
Income (annual, before tax)	25,734	(22,556)	23,558	(20,420)
Consumption (excl. health ins, quarterly)	4,503	(2,504)	4,160	(2,307)
Well-measured consumption	3,647	(1,622)	3,393	(1,494)
B. Consumption by category				
Flow value of housing	1,620	(1,032)	1,376	(878)
Food at home	642	(367)	627	(360)
Gas and motor oil	560	(532)	557	(542)
Utilities	457	(276)	464	(250)
Flow value of new vehicles	212	(274)	215	(267)
Communication	157	(121)	154	(123)
# Observations		36,387		21,927

Notes: The sample consists of households in the Consumer Expenditure Survey with no college education, 2008-2019. Expansion states include 2014-2019 expanders, and non-expansion states include never expanders (as of 2020). Well-measured consumption consists of the sum of the categories in panel B..

Table 2: The insurance value of Medicaid expansion is uncertain but small

Risk aversion	$\rho = 3$	$\rho = 1$	$\rho = 5$
A. Main estimates: Well-measured consumption, QDID, sample: $\leq$ HS, age 22-64			
Insurance value	22.2	0.1	191.9
95% CI	(-39.9, 97.8)	(-25.8, 30.9)	(-171.1, 541.8)
50% CI	(2.2, 47.5)	(-8.7, 10.7)	(54.8, 291.9)
B. Vary consumption measure: all consumption excluding health insurance			
Insurance value	25.3	8.1	164.1
95% CI	(-74.1, 136.4)	(-40.7, 62.3)	(-310.8, 612.0)
50% CI	(-9.9, 63.7)	(-9.7, 26.3)	(-4.0, 297.7)
C. Vary estimator: CIC			
Insurance value	17.9	-1.8	197.4
95% CI	(-55.5, 95.6)	(-32.3, 30.7)	(-183.4, 543.1)
50% CI	(-3.3, 44.9)	(-11.0, 10.2)	(52.5, 299.5)
D. Vary sample: $\leq$ 200% FPL, age 18-64			
Insurance value	99.9	19.2	92.9
95% CI	(-233.0, 283.2)	(-37.1, 71.3)	(-197.0, 397.2)
50% CI	(-5.0, 141.4)	(-1.7, 34.0)	(-14.7, 191.2)
E. Vary sample: $\leq$ 138% FPL, age 18-64			
Insurance value	102.7	21.4	89.9
95% CI	(-1932.0, 325.7)	(-91.5, 88.2)	(-252.2, 423.3)
50% CI	(-49.6, 158.4)	(-14.0, 40.8)	(-34.2, 193.2)

Notes: Table reports the implied quarterly insurance value of Medicaid expansion, for the indicated risk aversion, consumption measure, estimation approach, and sample. QDID is quantile differnce-in-differnences, and CIC is change-in-changes. 95% and 50% confidence intervals, calculated via the bootstrap (resampling states), reported in parentheses.

Figure A.1: Percentiles of well-measured consumption, 2015 expanders

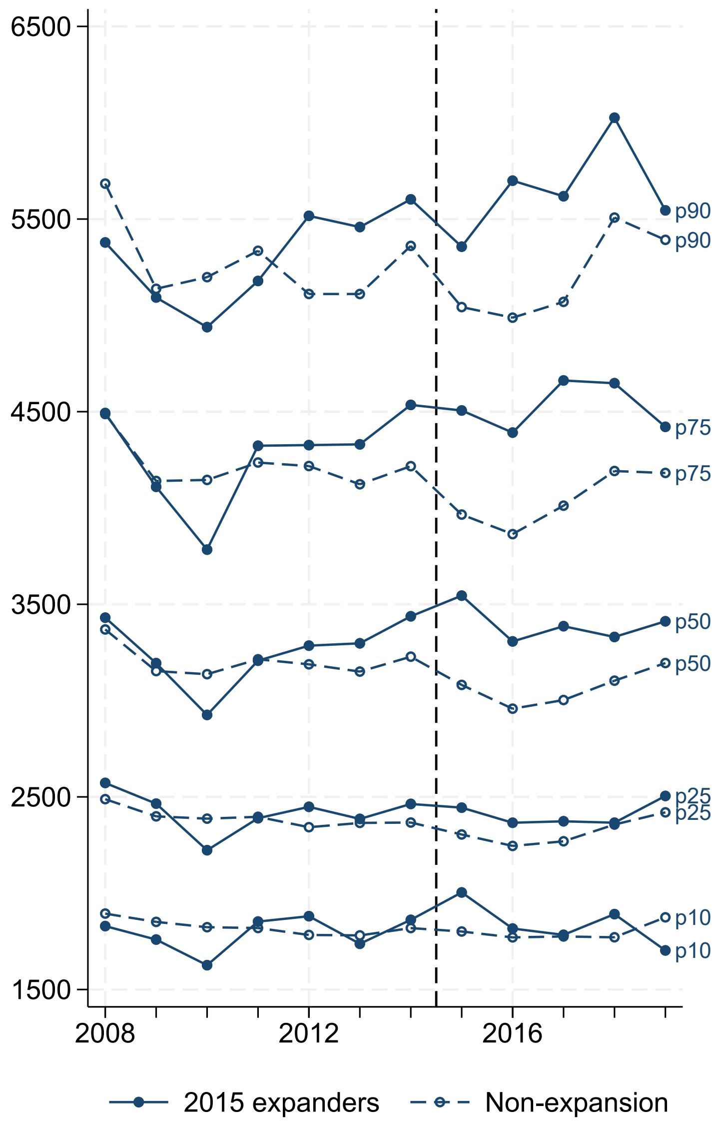

Notes: Figure plots the indicated percentiles of well-measured consumption, for 2015 expansion states (solid line) and non-expansion states (dashed line). Sample is defined in the notes to Table 1.

Figure A.2: Percentiles of well-measured consumption, 2016 expanders

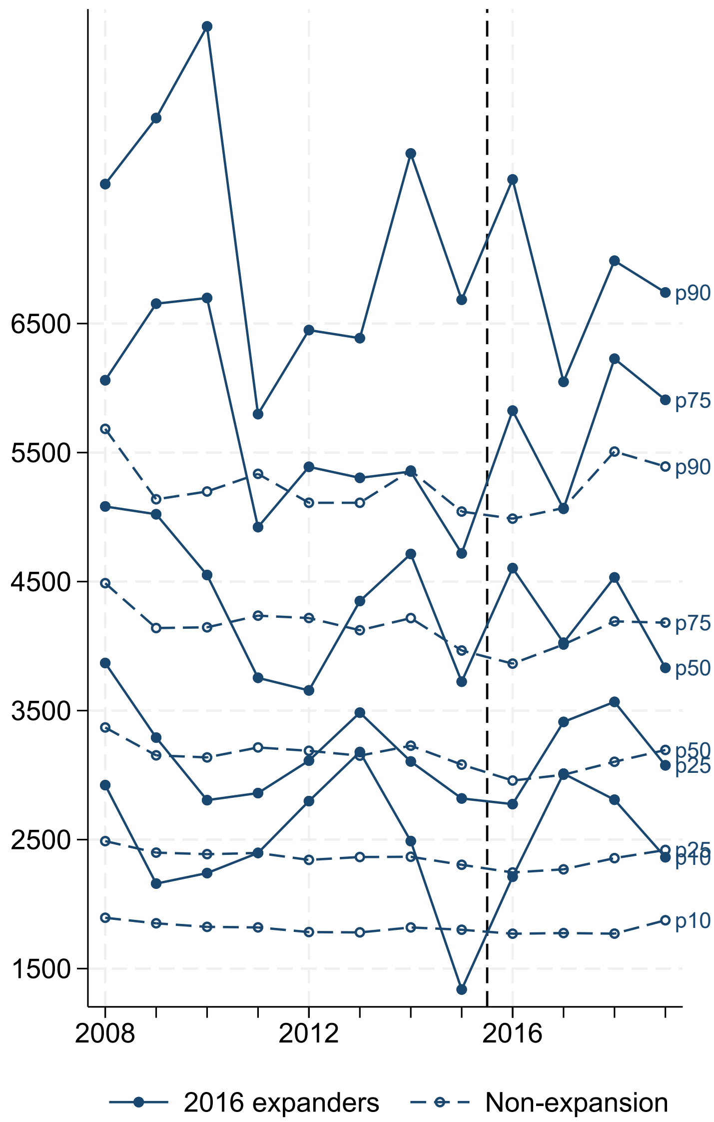

Notes: Figure plots the indicated percentiles of well-measured consumption, for 2016 expansion states (solid line) and non-expansion states (dashed line). Sample is defined in the notes to Table 1.

Figure A.3: Percentiles of well-measured consumption, 2017 expanders

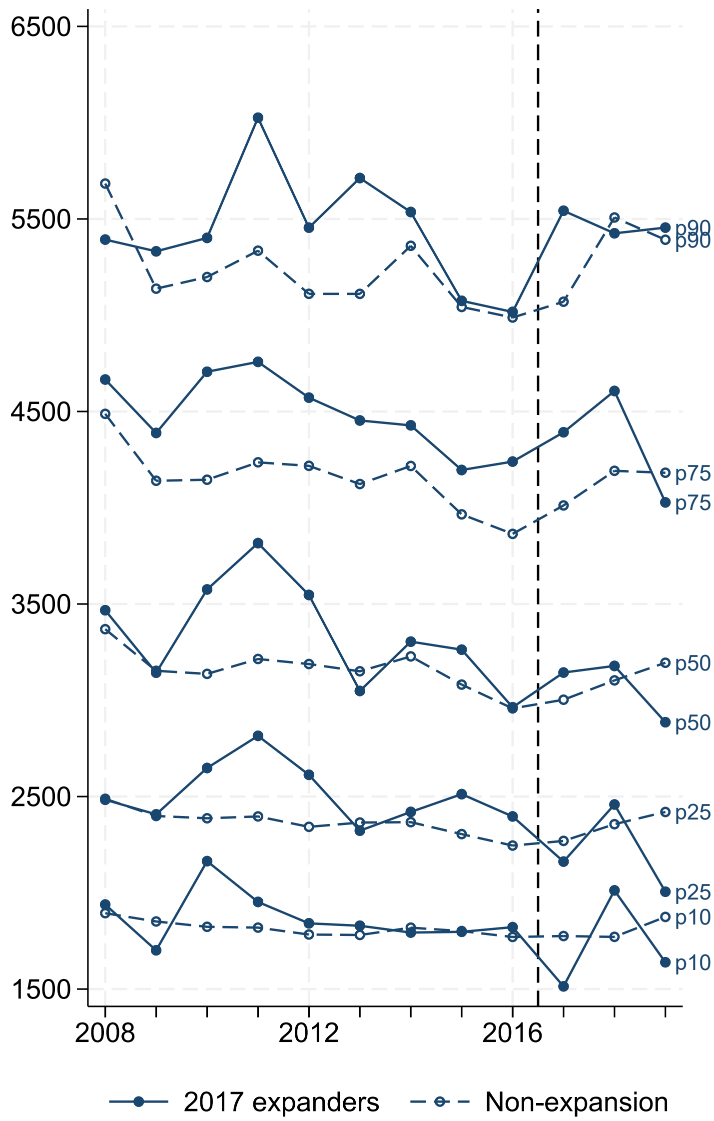

Notes: Figure plots the indicated percentiles of well-measured consumption, for 2017 expansion states (solid line) and non-expansion states (dashed line). Sample is defined in the notes to Table 1.

Figure A.4: Percentiles of well-measured consumption, 2019 expanders

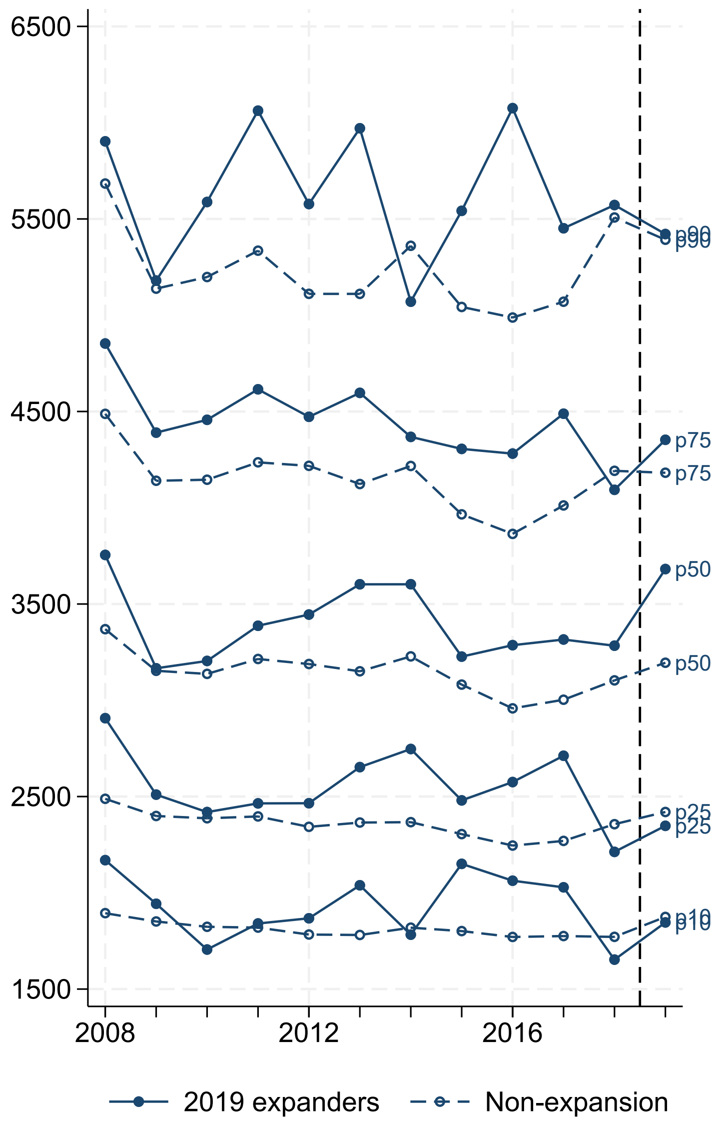

Notes: Figure plots the indicated percentiles of well-measured consumption, for 2019 expansion states (solid line) and non-expansion states (dashed line). Sample is defined in the notes to Table 1.

Figure A.5: CIC estimates of impact of Medicaid expansion on the consumption distribution

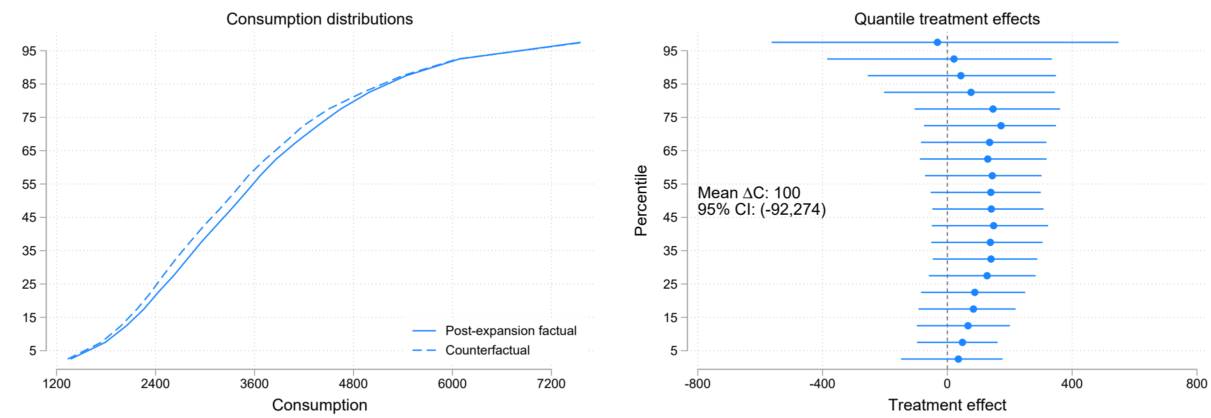

Notes: Figure is identical to Figure 3, except we use CIC rather than QDID to estimate distributional effects.