

The Consumption Smoothing Effects of Health Insurance*

RTS 2023 Submission, Extended Abstract

Anita Mukherjee Daniel W. Sacks Hoyoung Yoo

December 15, 2023

Abstract

We investigate the value of health insurance as consumption insurance by leveraging state-level Medicaid expansion decisions. While Medicaid expansions improve health and financial status, they may not reduce consumption risk because a great deal of uninsured medical spending is financed with bad debt and charity care rather than reduced consumption. Using a combination of difference-in-differences and changes-in-changes specifications, we find small effects of Medicaid expansion throughout the consumption distribution. Our estimates imply near-zero insurance value from Medicaid expansion, and our confidence intervals let us rule out the possibility that a large share of Medicaid's value comes from reduced consumption risk.

*Mukherjee: Risk and Insurance Department, Wisconsin School of Business, University of Wisconsin-Madison, anita.mukherjee@wisc.edu. Sacks: Risk and Insurance Department, Wisconsin School of Business, University of Wisconsin-Madison, daniel.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. The research reported herein was performed pursuant to a grant from the US Social Security Administration (SSA) funded as part of the Retirement and Disability Consortium. The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA or any agency of the Federal Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation, or favoring by the United States Government or any agency thereof.

Extended Abstract

The decade since the Affordable Care Act has seen a great increase in the share of Americans with health insurance. This increased insurance coverage has brought myriad health benefits for the newly insured. It has also provided financial security, as health insurance expansions reduce bankruptcy and medical debt, and increases credit scores. Do these *financial* improvements also imply a reduction in *consumption* risk?

The impact of health insurance on consumption risk is an important but unresolved empirical question. It is important because reducing consumption risk—decreasing the incidence of very low consumption—is the neoclassical reason that health insurance is welfare-improving despite leading to over-consumption via moral hazard. It is unresolved for two reasons. First, existing research does not provide any direct evidence on how health insurance coverage affects the distribution of consumption. Instead it focuses on health, health care utilization, and financial outcomes such as bankruptcy. Second, as prior research has found that health insurance coverage reduces bankruptcy, the benefits of health insurance coverage flow not only to households (in the form of increased consumption), but also to health care providers. Indeed, hospitals providers are a major financial beneficiary of insurance expansions. Thus greater health insurance coverage does not automatically imply a reduced exposure to health care risk, as that risk may have already been financed via uncompensated care.

To study how health insurance affects consumption risk, we estimate the impact of Medicaid expansion on the distribution of consumption. Our analysis uses difference-in-differences (DID) and changes-in-changes (CIC) models. While our basic identification approach is similar to the large literature on Medicaid expansions, we depart from the literature by estimating effects across the entire distribution of consumption, rather than focusing on the mean. This broader view is critical for assessing the risk-reducing role of Medicaid expansions. Our

empirical approach yields estimates of the (actual) consumption distribution with Medicaid and the (counterfactual) distribution absent Medicaid. We use these two distributions, as well as assumed expected utility functions, to estimate Medicaid’s “insurance premium”, i.e. the value of the reduction in consumption risk generated by Medicaid expansion, over and above any mean resource transfer.

We use two sources of consumption data for our analysis. Both sources of data provide information at the household-level, so we use an equivalence scale following Meyer and Sullivan (2023) and Citro et al. (1995) to obtain per-capita outcomes that account for differences in family size and composition. To focus our analysis on those who may be affected by Medicaid expansion, we restrict our sample in each dataset to those aged 22–64 with education less than or equal to high school. We also restrict our analysis to states that expanded Medicaid in 2014 or never expanded Medicaid.¹ In what follows, we do not extend our data to 2020 or further to avoid complexities in the analysis arising from the COVID-19 pandemic, which greatly affected both health needs and consumption patterns. Our data contain 48,471 households in the CEX data and 45,678 households in the Nielsen data.

As a preview of results, Table 1 shows the DID and CIC results for well-measured consumption from the CE data. Here, we find that Medicaid expansions do not exert a protective effect over consumption even at the left tail of the distribution. In column (1), the mean DID estimate is \$50, though the estimate is noisy. We observe that the CIC estimates are not statistically significant at all points of the consumption distribution. The same is true for the logged specification in column (2).

To place our results in improved context of the literature and the Medicaid program, we calculate the implied consumption risk premium for the health insurance expansion. The risk premium is a measure of insurance or risk-reducing value of Medicaid expansion. Below, we outline the theory and estimation for this object of interest.

¹Kaestner et al. (2017) shows that those with low education/low income were more likely to gain insurance coverage following the 2014 Medicaid expansions.

Theory: Consider two states of the world, state 1 with Medicaid expansion and state 0 without. Let C_1 and C_0 , random variables, be consumption in each state of the world. We assume that both have finite support taking on k values, $c_1^1, \dots, c_1^K$ and $c_0^1, \dots, c_0^K$. Expected utility in state of the world $j \in 0, 1$ is

$$EU_j = \sum_k p_k u(c_k^j), \quad (1)$$

where u is the von Neuman-Morgenstern utility function.

As in Finkelstein et al. (2019), define the certainty equivalent γ of the Medicaid expansion as the amount of consumption a person would have to give up in the expansion state of the world to be indifferent between expansion and non-expansion. γ is implicitly defined by

$$\sum_k p_k u(c_k^1 - \gamma) = \sum_k p_k u(c_k^0). \quad (2)$$

Finally, define the risk premium π as the difference between the certainty equivalent and the expected value (in consumption) of expansion:

$$\pi \equiv \gamma - E[C_1 - C_0]. \quad (3)$$

Estimation: In the data, we observe consumption under expansion for the expansion states in the post period. We don't know the distribution of consumption had these states not expanded. Let $\hat{c}_1^1, \dots, \hat{c}_K^1$ be the empirical percentiles of the consumption distribution, and percentiles $P_1, \dots, P_K$. We assume that the CIC method yields estimates of percentiles of the counterfactual distribution, i.e., $\hat{c}_1^0, \dots, \hat{c}_K^0$.

Given these estimates, estimate $E[C_j]$ in a straightforward manner:

$$E[\hat{C}_j] = \sum_{k=1}^K (P_k - P_{k-1}) \hat{c}_k^j, \quad (4)$$

with $p_0 = 0$.

To find γ , we need to assume a utility function. For this report, we follow Finkelstein et al. (2019) and use constant relative risk aversion (CRRA) utility with coefficient of relative risk aversion $\rho = 3$.

Given $u(\cdot)$ and the estimated distributions, the only challenge is recovering γ . Estimating γ requires solving one equation in one unknown (i.e., equation 2), which is straightforward. Thus the approach is:

1. Approximate the factual consumption distribution with a, say, 19 point distribution, evenly spaced from 5th to 95th percentile.
2. Use CIC to recover counterfactual consumption percentiles.
3. Solve for $\hat{\gamma}$ by implementing the sample analog of equation 2.
4. Solve for $\hat{\pi}$ by subtracting off $E[\hat{C}_1] - E[\hat{C}_0]$.

Inference: Steps 1–4 yield a single estimate of π . We obtain confidence intervals via the bootstrap, re-estimating π in each bootstrap iteration.

The estimated risk premia are presented in Table 2. Following Finkelstein et al. (2019) we present our estimates for a range of risk aversion parameters. We estimate the risk premium when $\rho = 3$ to be \$12.7; the 95 percent confidence interval is (-480, 196). The risk premium we estimate is below the lower end of the range of estimated consumption welfare benefit from Medicaid as estimated in Finkelstein et al. (2019). That paper uses the 2008 Oregon Health Insurance Experiment to estimate an insurance value ranging from \$112 to \$883 per recipient-year (Table 2 in that paper). The \$12.7 benefit is also small relative to the per-capita cost of Medicaid, which that paper estimates to be about \$3,600 (average across states). Another set of useful benchmarks for the risk premium of \$12.7 come from the breakdown of Medicaid in other aspects: the program’s insurance premium is inferred

from spending effects to be \$133 to \$1,106; the value of the transfer component is \$600 to \$900; the reduction in out-of-pocket spending is about \$600.

Taken together, we find that health insurance expansion does not exert a protective consumption effect for individuals even at the left tail of the consumption distribution. This finding is supported using data on household grocery expenditures as well as total well-measured consumption from the Consumer Expenditure Interview Surveys as studied in Meyer and Sullivan (2023). Our estimate of the risk premium implies a low *consumption insurance* value to Medicaid, though certainly the program offers large insurance benefits on health, financial shocks, and other dimensions of wellness such as access to more regular (versus emergency) care (Ellis and Esson, 2021). Strength in social insurance in other forms, such as food stamps or payment assistance plans that can be activated on hardship, could be other reasons that we do not observe strong consumption insurance benefits to Medicaid.

Cited References

- Citro, C. F., R. T. Michael, et al. (1995). *Measuring poverty: A new approach*. National Academy Press.
- Ellis, C. M. and M. I. Esson (2021). Crowd-out and emergency department utilization. *Journal of Health Economics* 80, 102542.
- Finkelstein, A., N. Hendren, and E. F. Luttmer (2019). The value of medicaid: Interpreting results from the Oregon health insurance experiment. *Journal of Political Economy* 127(6), 2836–2874.
- Kaestner, R., B. Garrett, J. Chen, A. Gangopadhyaya, and C. Fleming (2017). Effects of ACA Medicaid expansions on health insurance coverage and labor supply. *Journal of Policy Analysis and Management* 36(3), 608–642.
- Meyer, B. D. and J. X. Sullivan (2023). Consumption and income inequality in the united states since the 1960s. *Journal of Political Economy* 131(2), 247–284.

Table 1: DID and CIC: Well-measured Consumption

	(1)	(2)
A. Difference-in-Differences		
	Well-measured Consumption	Log
β	50.37 (84.50)	0.01 (0.03)
B. Changes-in-Changes		
	Well-measured Consumption	Log
mean	45.00 [-181.07, 272.27]	0.01 [-0.05, 0.08]
p5	109.65 [-25.69, 248.98]	0.06 [-0.02, 0.16]
p10	81.65 [-85.00, 222.36]	0.04 [-0.05, 0.12]
p25	55.28 [-159.98, 230.42]	0.02 [-0.06, 0.09]
p50	80.97 [-149.90, 285.56]	0.02 [-0.04, 0.08]
p75	98.74 [-164.85, 331.10]	0.02 [-0.03, 0.07]

Notes: State-clustered standard errors (95 percent) in parentheses; these are only provided for the DID model. Bootstrapped and state-clustered confidence intervals (95 percent) in brackets are provided for the CIC models.

Table 2: Risk Premium Estimates (Annualized)

Risk aversion:	$\rho=3$	$\rho=1.01$	$\rho=5$
Risk premium:	12.7 (-480,196)	17.7 (-148,800)	-195 (-976,148)

Notes: Estimation assumes a CRRA utility function and with risk aversion parameter ρ ; standard errors are bootstrapped. Bootstrapped and state-clustered confidence intervals (95 percent) in parentheses.