Dynamic Complementarity in Public Insurance Provision: Evidence from the Medicaid Expansion*

Spencer Perry†

November 25, 2023

Click here for the latest version

Abstract

Current healthcare consumption may yield future health benefits that lower longer-run expenditures. In this study I leverage exogenous variation in the expansion of Medicaid under the Affordable Care Act to assess how comprehensive insurance affects immediate healthcare spending for the near-elderly (aged 60-64) and future healthcare spending once they receive Medicare at age 65. Using panel data from the Health and Retirement Study, I provide evidence that Medicaid coverage increases immediate expenditures by roughly 112 percent; but, after receiving Medicare, the same individuals are healthier and consume 77 percent less care measured by total expenditures. I then develop and estimate a life-cycle model of near-elderly individuals that incorporates the Medicaid expansions as well as dynamic health investment and endogenous mortality. I find that the Medicaid expansions were valued by the near-elderly population at slightly above net government costs, and that modest reforms to increase Medicare's generosity could generate considerable welfare improvements relative to program costs.

*I am indebted to my advisors, Dan Silverman, Nicolai Kuminoff, and Esteban Aucejo, for their invaluable mentorship and advice. This paper also benefited greatly from my conversations with Tomas Larroucau, Jacob French, Nirman Saha, Andreas Kostøl, Domenico Ferraro, Raphael dos Santos, Hector Cardozo, Nicolas Bozzo, Matt Millington, Jasdeep Mandia, Matthew Merkle, Ignacio Amaral, Riley League, and Barton Willage. For helpful comments, I thank Hunt Allcott, Lindsay Allen, David Arnold, Joseph Benitez, Doug Bernheim, Peter Blair, Sophie Calder-Wang, Murat Çelik, Jeffrey Clemens, Jackson Dorsey, Matt Freedman, George-Levi Gayle, Ben Handel, Peter Hull, Ammanda Kreider, Anthony LoSasso, Joana Naritomi, Ivan Rudik, Jane Ruseski, Immo Schott, Paulo Somaini, Dimitry Taubinsky, Alisa Tazhitdinova, Yulya Truskinovsky, Dan Waldinger, Laura Wherry, seminar participants at Arizona State University, the 2021 Association for Public Policy and Management Fall Conference, the 2023 American Society of Health Economists Annual Conference, the 2023 Western Economics Association Graduate Student Workshop, and the 2023 NBER Behavioral Public Economics Bootcamp. All errors are my own.

†Department of Economics, Arizona State University, Email: asperry5@asu.edu

1 Introduction

Current health investments have the potential to influence future demand for medical care; however, much of the economics literature examining health insurance omits this mechanism. In particular, if incremental healthcare consumed due to lower out-of-pocket prices (i.e., moral hazard) improves health outcomes, then moral hazard spending may reduce future healthcare consumption. These dynamics are particularly relevant in the context of public health insurance in the United States, where nearly all individuals aged 65 and older receive government financed insurance through Medicare. Moreover, high uninsurance rates due to market incompleteness or information frictions may lead some individuals to underinvest in health prior to age 65. In this case, earlier provision of public health insurance may generate a positive dynamic fiscal externality on the Medicare system in addition to providing value to consumers.

The goal of this paper is to quantify the role of health dynamics in shaping the value and cost-effectiveness of expansions in public health insurance. To do so, I investigate how Medicaid enrollees' moral hazard spending affects their future healthcare consumption, accounting for potential selection. My research design leverages state-level variation in Medicaid expansions under the Affordable Care Act (ACA). These expansions led to plausibly exogenous variation in *whether* and *when* low-income individuals gained Medicaid eligibility. Through this quasi-experimental variation, I estimate the causal effect of a large-scale increase in publicly provided health insurance on current and future healthcare spending and health. I focus on individuals who are near the threshold age for Medicare eligibility, which allows me to better compare future spending and utilization across individuals since insurance options after age 65 are relatively homogeneous. To evaluate the welfare implications of expanding public health insurance, in the second part of this paper I develop and estimate a life-cycle model that incorporates the Medicaid expansions, dynamic health investment, and endogenous mortality. In estimation, I use the reduced form causal estimates as targeted moments to credibly identify key model parameters.

Focusing on the potential complementarity between the Medicaid expansions and Medicare program is both policy relevant and technically appealing. First, both programs are large—together, they insure over 120 million Americans and account for nearly one quarter of all annual federal outlays (i.e., \$1.3 trillion). Furthermore, these two programs specifically cover the low-income (Medicaid) and the elderly (Medicare), which are some of the most vulnerable members of society. Second, many people covered by Medicaid would otherwise be uninsured. Third, the federal government pays for nearly all of the costs of the Medicaid expansions; meaning that any dynamic fiscal externalities generated are internalized by the same entity paying the upfront costs. Finally, the expansions changed insurance eligibility rules differentially across states and across time, whereas the discontinuity in access to Medicare (i.e., turning 65) is constant. This allows me to exploit the exogenous change in insurance eligibility for both programs without concerns that one policy affects the rules of the other.

I begin my empirical analysis by employing an instrumented triple-differences (TDIV) strategy to estimate the causal effect of Medicaid coverage on current and future healthcare expenditures using panel data from the Health and Retirement Study (HRS).[1] Consistent with prior literature, I find that the expansions increased Medicaid enrollment and access to healthcare for individuals aged 60 to 64.[2] Specifically, I find that new enrollees immediately purchased roughly twice as much healthcare as measured by annual expenditures—an increase of approximately $1,900 per enrollee, on average. Next, I examine how exposure to Medicaid affected later healthcare consumption for new Medicare enrollees. My results indicate that, all else equal, those who received Medicaid due to the expansions were healthier during their early Medicare tenure, spent less out-of-pocket, and amassed roughly 60-80 percent less in annual gross medical expenditures—an average annual decrease of $1,300 per person. Additionally, I find that the reduction in expenditures and improvements in health persist for several years after turning 65. A back-of-the-envelope calculation suggests that the dynamic effect of Medicaid coverage recovers roughly half of the initial Medicaid costs. Taken together, these results suggest that insurance expansions allow individuals to afford timely medical care that has a dynamic benefit of lowering future healthcare needs and expenditures.

Though these results indicate that the Medicaid expansions imposed a positive fiscal externality onto Medicare, recent work has shown that insurance coverage (Goldin, Lurie, and McCubbin, 2020 and Abaluck et al., 2020) and the Medicaid expansions specifically (Miller, Johnson, and Wherry 2021) lead to significant reductions in mortality. Therefore, a robust accounting of the costs and benefits of expanding needs to account for endogenous survival. In the second part of this paper I develop and estimate a finite-horizon life-cycle model for individuals over age 55 that incorporates the primary features of my empirical setting along with the key insight of dynamic health investment introduced by Grossman (1972). I use my TDIV estimates as targeted moments to identify the parameters of the health production function and preference parameters governing the complementarity of health and consumption in utility. Then, I evaluate how cost-effective the Medicaid expansions were for the near-elderly population, finding that the expansions were valued at $1.05 for every dollar of net cost to the government.

The cost-benefit analysis yields several additional insights: (1) a large share of the welfare generated by the program accrues to people who are unlikely to ever be eligible for Medicaid. However, they benefit considerably from the risk protection the program offers because the state of the world in which they want to enroll in Medicaid (i.e., low-income, uninsured, and unhealthy) is exactly when the marginal utility from consumption and health is highest.[3] (2)

[1] Intuitively, identification comes from cross-state and within-state variation in eligibility over time, where Medicaid enrollment is instrumented by Medicaid eligibility based on year, state of residence, and low-income status. Formally, I derive identification of the TDIV estimator in Appendix D.

[2] I also find no evidence of anticipatory behavior prior to the expansions.

[3] This result is intuitively similar to the key finding in Braun, Kopecky, and Koreshkova (2017) that means-tested

Though my reduced form results suggest a positive fiscal externality of the Medicaid expansions onto Medicare in the short term, I find that accounting for endogenous survival is likely to cause net Medicare costs to increase following the expansions. The main reason being that the Medicaid expansions decrease mortality (particularly among individuals with poor health), increasing both the number of individuals covered by Medicare and their longevity.[4] (3) Roughly one third of the value of the Medicaid program arises specifically due to individuals' ability to increase their future health through health investments. I interpret this result as suggestive evidence that, for this population, additional healthcare spending induced by lower out-of-pocket prices is considerably valuable and unlikely far out on the "flat of the curve." This also reflects the fact that these households face considerable liquidity constraints, and public insurance drastically shifts their access to (potentially large) health investments.

Finally, I use the model to analyze a counterfactual reform to the Medicare program that allows lower income individuals up to 250 percent of the Federal Poverty Level (FPL) to "buy-in" to Medicare at age 62.[5] I find that this reform would generate roughly \$1.15 of welfare for every dollar of net cost to the government, which is more cost-effective than the Medicaid expansions for individuals in the same age range. However, the total welfare gains are smaller and more dispersed among the population. This is primarily due to the less generous cost-sharing of Medicare, which mitigates the initial moral hazard response but limits the immediate and dynamic benefits of coverage. Overall, the model results suggest that there is considerable scope to cost-effectively reform public health insurance for the near-elderly population.

My findings have both broad and specific policy implications. My TDIV estimates suggest that provision of public health insurance does induce additional expenditures, but that these costs are somewhat mitigated by the dynamic benefit of coverage. However, the structural model results suggest that future cost reductions are short lived due to increased longevity associated with expanding public health insurance. Broadly, I find that increasing access to public insurance for the low-income near-elderly is cost-effective, but that the exact level of efficacy is dependent on policy design. A stricter means-test coupled with more generous insurance can provide larger and more targeted benefits but at much higher costs. Alternatively, broader provision with leaner benefits delivers more benefits per net cost. Additionally, combining my evidence on the dynamic health benefits of the Medicaid expansions for older adults with prior evidence of smaller effects for younger adults (e.g., Finkelstein et al., 2012) suggests that there is likely an age at which these benefits become relevant and that the most efficient age at which to provide a universal form of public health insurance (i.e., Medicare) should take advantage of these complementarities.

---

insurance programs for the elderly produce large welfare gains.

[4]In counterfactual simulations of the implementation of Medicare Part D, Yang, Gilleskie, and Norton (2009) also find that longevity is an important determinant of program costs.

[5]This is motivated by recent proposals from U.S. lawmakers to lower the Medicare eligibility age to age 60 (Luhby, 2021).

This study is closely related to a small literature that examines the dynamic role that health insurance plays in determining the trajectory of future health and healthcare demand (e.g., Fang and Gavazza, 2011; Black et al., 2017; McWilliams et al., 2007; and Yang, Gilleskie, and Norton, 2009).[6] While previous studies have provided some insight into the dynamic effects of insurance coverage, the results have been somewhat contradictory, likely reflecting the inherent challenge in isolating causal effects due to potential unobserved differences between the uninsured and insured populations. My research adds to this literature by utilizing a research design that leverages plausibly exogenous variation in Medicaid eligibility, which helps to overcome bias due to selection.

This paper is also related to a larger literature on the effect of insurance on contemporaneous healthcare consumption and health. In particular the Rand Health Insurance Experiment from the 1970s (Newhouse and The Insurance Experiment Group, 1993) and the Oregon Health Insurance Experiment in 2008 (Finkelstein et al., 2012 and Baiker et al., 2013), which document that healthcare consumption increases in response to more generous insurance coverage. Additional work, however, has demonstrated that insurance coverage substantially reduces mortality, especially for the low-income near-elderly population (Card, Dobkin, and Maestas, 2008; Miller, Johnson, and Wherry, 2021; and Goldin, Lurie, and McCubbin, 2020). I contribute to this literature by providing novel evidence unifying these results. Namely, that additional healthcare consumption induced by insurance coverage improves future health, which can explain reductions in mortality. Moreover, the model results reflect each of these margins and provide valuable insight into their implications for consumer welfare and insurer costs.

The structural model developed in this study relates to a separate literature on life-cylce models with health uncertainty (e.g., Palumbo, 1999; De Nardi et al., 2010; French and Jones, 2011; De Nardi et al., 2016; and Lockwood, 2018). I contribute to this literature by modeling medical spending as a choice variable and allowing for endogenous substitution between consumption, health, and mortality risk. Several studies have also treated medical expenditures and health as endogenous choices (e.g., Yang, Gilleskie, and Norton, 2009; Khwaja, 2010; Ozkan, 2017; and Saha, 2021). I extend this literature primarily through my strategy for identifying the structural parameters. I take an indirect inference approach that allows my causal TDIV estimates to help identify my structural parameters of interest, while remaining consistent with life-cycle profiles of health, medical expenditures, and mortality rates. Additionally, I model endogenous enrollment into public insurance (i.e., Medicaid). Therefore my results reflect the fact that not all eligible individuals enroll and that hassle costs may exacerbate selection.

The rest of this paper is organized as follows. Section 2 outlines key institutional features

[6] A number of studies that examine the shorter-run dynamics in health insurance within a contract period, typically a year (e.g., Diaz-Campo 2022; Aron-Dine et al., 2015; Einav, Finkelstein, and Schrimpf, 2015; Cronin, 2019, Lin and Sacks, 2019), but these are less related. There is also a large literature examining the long-run of insurance coverage for children (e.g., Boudreaux, Golberstein, and McAlpine, 2016 and Goodman-Bacon, 2021), which is also distinct.

of the Medicaid expansions under the ACA and summarizes prior research on the impacts. Section 3 describes the data that I use, and discusses how the Medicaid expansions affected insurance coverage for individuals in the HRS. Section 4 details my reduced-form models and their identifying assumptions. I present the results from my reduced form analysis in Section 5. Section 6 presents a stylized late in life-cycle model of healthcare demand. Section 7 outlines how I estimate the structural model, discusses identification, and presents the model results and fit. In Sections 8 and 9 I use the model to evaluate the welfare gains and associated net government costs of the Medicaid expansions for the near-elderly and a counterfactual reform to the Medicare program. Finally, Section 10 concludes.

2 The Medicaid Expansion Under the Affordable Care Act

A key component of the ACA was to expand Medicaid to adults in families with incomes up to 138 percent of the FPL.[7] Prior to the ACA, Medicaid was fairly restrictive, providing insurance mainly to pregnant women and children. Originally, the ACA intended for the expansions to apply to all states, with states given some flexibility in implementation. However, in 2012 the Supreme Court ruled that the expansions needed to be voluntary at the state level. As a result, only twenty-nine states and Washington D.C. adopted the expansions upon initial implementation in 2014. As of summer 2023, ten states have yet to expand and North Carolina has adopted, but not yet implemented, the expansion (Kaiser Family Foundation 2023). Figure 1 presents states by expansion status as of 2019.[8]

In addition to increasing eligibility, the expansions aimed to simplify the application process and the eligibility rules. Before the ACA, Medicaid applicants had to pass an asset test to qualify. The ACA abandoned the asset test for most people under 65. Therefore, to be eligible in an expansion state, an individual needed to only demonstrate that their Modified Adjusted Gross Income (MAGI) was below the new FPL threshold.[9] In total the expansions are estimated to have affected roughly 13.6 million people across the U.S. (Centers for Medicare and Medicaid Services 2020).

Since passing in 2010, the ACA and the Medicaid expansions have been studied extensively. Early analyses concluded that the Medicaid expansion was particularly effective in lowering the

[7]In addition to expanding Medicaid the ACA introduced several additional policies aimed aimed at moving the United States towards universal health insurance coverage. The first implemented a new regulator regime for private insurers, which, among other changes, famously required insurers to provide health insurance to all individuals regardless of pre-existing conditions. The second was an individual mandate for all legal residents of the U.S. to purchase health insurance. Individuals violating the mandate are subject to a fine in the form of a tax penalty.

[8]I chose this version of the map because my data only extend through 2018. Therein, this version of the map provides the relevant information on expansion status to help visualize the variation underlying my identification strategy.

[9]MAGI includes labor income, Social Security income, and other sources of capital income but excludes Supplemental Security Income (SSI).

Figure 1. Map of Medicaid expansion status as of 2019

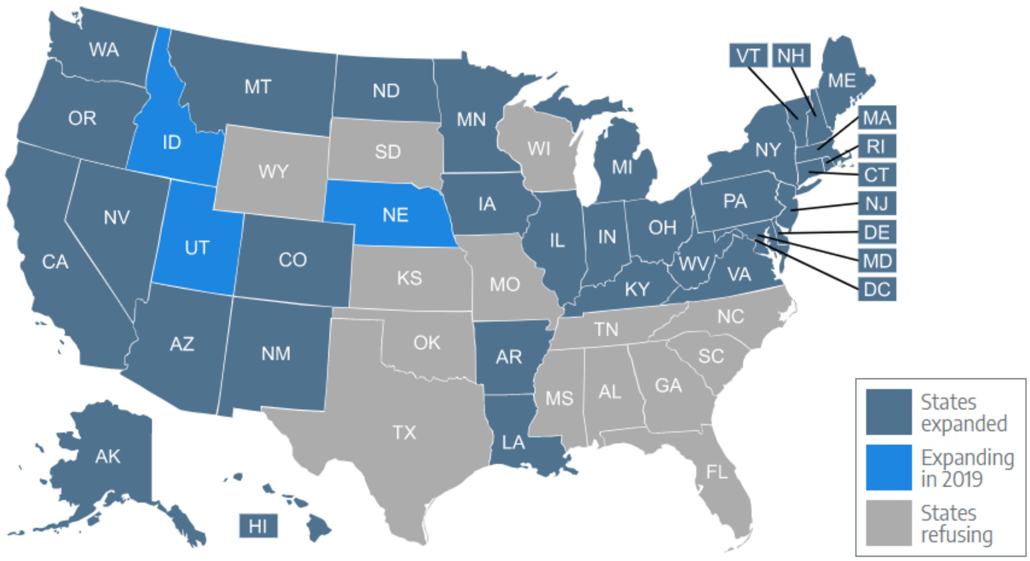

Notes.—Figure from Kaiser Family Foundation's "Status of State Medicaid Expansion Decisions: Interactive Map".

non-elderly uninsured rate (e.g., Courtemanche et al. 2017 and Mazurenko et al. 2018). Consistent with the moral hazard hypothesis, the expansion has consistently been associated with increased healthcare utilization.[10] Other studies have focused on health effects of the expansions. For instance Miller, Johnson, and Wherry (2021) find that adults most likely to benefit from the policy saw a 9.4 percent reduction in annual mortality, and Khatana et al. (2019) find that rates of cardiovascular disease among adults aged 45-65 declined in states that adopted the Medicaid expansion relative to those that did not. Tipirneni et al. (2020) looks at how the Medicaid expansion affected access to and utilization of care for the sample of individuals that I am most interested in. Their results suggest that the expansions led to an increase in Medicaid coverage for low-income individuals across the country, but coverage increased substantially more in expansion states. In terms of consumption, the authors find that expansion states saw increases in hospitalization rates but not along any other measures of healthcare use. The authors conclude that this result may suggest "poor access to chronic disease management and pent-up demand for hospital services" for the low-income uninsured.

In addition to finding that expansion states saw increases in coverage and access to care for the target population, Simon, Soni, and Cawley (2017) investigate the potential "ex-ante moral hazard" response to the expansions.[11] Overall they find that expansions increased the use of

[10] Mazurenko et al. (2018) provide a comprehensive review.

[11] "Ex-ante moral hazard" refers to changes in lifestyle behaviors such as smoking habits or diet that occur due to the fact that medical care is relatively cheaper with insurance. This differs from "ex-post moral hazard", or the responsiveness of demand for healthcare to the price an individual has to pay for it, which is the focus of this paper.

some preventative care, but did not seem to drive any additional risky health behaviors. Further, they conclude that the expansions led to a slight increase in self-reported health amongst low-income individuals. Also of note is the effect that the expansions had on non-health outcomes. For instance, Hu et al. (2018) explore the effect of the expansion on financial outcomes for low-income populations. Their results suggest that the number of unpaid bills and the amount of debt sent to collection agencies declined substantially for the low-income, highly uninsured communities in expansion states.

The reduced form exercise in this paper is also closely related to recent work by Shupe (2023) that uses the Medical Expenditure Panel Survey (MEPS) to examine the effect of the Medicaid expansions on healthcare consumption. Somewhat contradictory to my results, she finds that healthcare spending did not increase for the target population in expansion states. However, the expansions dramatically shifted the incidence of payment from private insurers to public sources and, therefore, taxpayers. There are a number of reasons to expect that I would find different results. In particular, their study examines the entire adult population rather than just the near-elderly. If the near elderly derive a larger benefit from consumption of medical care, then it is likely that they will be more price sensitive than the average adult.

3 Data

3.1 Data Sources and Construction of Estimating Samples

To estimate the extent to which the expansion of Medicaid affected medical consumption I use the Health and Retirement Study and the Medical Expenditure Panel Survey. In this section I briefly describe how use each of these data sets and my strategy for imputing medical expenditures in the HRS using MEPS. I then discuss how I construct the main samples for my analysis. I conclude the section with a discussion of how I identify individuals affected by the Medicaid expansions in the data and how important the expansion was in increasing actual Medicaid enrollment.

Data Sources.—The HRS provides a nationally representative panel of individuals near and in retirement, spanning from 1992 to 2018. Individuals are interviewed every two years from their first interview wave until they die. The survey records individual demographics and economic variables of interest such as educational attainment, income, employment history, and wealth.[12] Important for my study, the HRS also has detailed information on health status, healthcare utilization, insurance coverage, and out-of-pocket spending. Measures of health include self-reported health and objective indicators such as smoking habits, disease diagnoses, and if the respondent has difficulties completing daily tasks. Utilization information includes measures

[12]I deflate all monetary values in my analysis using the GDP implicit price deflator, with 2012 as the base year.

such as the number of doctor visits and nights spent in the hospital.

There is one important limitation with the HRS data for this study. In earlier years, the survey included questions about total medical expenditures. This variable aimed to report the total cost of all medical services the individual received, rather than just their out-of-pocket payments. To estimate the extent of moral hazard induced spending due to the Medicaid expansion and its effect on future expenditures, I would ideally observe this variable for the entire length of the study. Unfortunately, the HRS removed this question from the survey in the mid 2000s, well before the Medicaid expansion took place. I address this limitation by adopting the approach of Fang, Keane, and Silverman (2008). Specifically, I complement the HRS with another data set—the MEPS—which is also nationally representative, contains many of the same survey measures for health and utilization, but also includes detailed expenditure information. Using information that is similar across the data sets, I develop a regression based imputation strategy which allows me to predict an individual's total medical expenditures in the HRS. I provide an additional discussion of both the HRS and the MEPS in Appendix A and I provide extensive details on the imputation process and results in Appendix B.

HRS Sample Construction.—The empirical strategy that I present in Section 4 requires that I construct two primary samples for analysis. The first is the pre-Medicare sample which consists of individuals who I can observe enrolling in Medicare at age 65 at some point during the panel. The second is the sample of individuals who are on their first or second waves of Medicare and are at least age 65.[13] In order to do this I determine the second sample first. I then construct the first sample such that it includes only individuals that I observe in the second sample. This way I am sure that any effects that I detect in the pre-Medicare and Medicare groups are occurring within the same set of individuals. I provide additional details and descriptive statistics for the main analysis samples in Appendix A.

3.2 Medicaid Expansion in the HRS

To help validate the main empirical approach that I outline in Section 4, I first look to the data to see how the expansions affect Medicaid enrollment in the HRS. Table 1 presents estimates from a series of triple-difference-in-differences (TD) models of the following form:

$$Medicaid_{i,t} = \alpha_0 + \alpha_1 Expand_{s,t} \times LI_{i,t} + \alpha_2 Expand_{s,t} + \alpha_3 LI_{i,t} + X'_{i,t}\Phi + \Psi + \epsilon_{i,t} \qquad (1)$$

where $Medicaid_{i,t}$ is an indicator for whether individual i claims to have Medicaid coverage in time t; $Expand_{s,t}$ is a dummy equal to one if individual i is living in a state s that has expanded Medicaid by time t; $LI_{i,t}$ is an indicator for if individual i is in the target population of the expansions at time t; X is a vector of time-variant individual characteristics like health status

[13]A wave of Medicare corresponds to one to two years of Medicare coverage. So individuals in their first wave are either 65 or 66 and on Medicare for the first time in the sample.

and income; and Ψ denotes fixed effects for state, year, state by year, state by low-income, and year by low-income. In this TD equation, the parameter of interest is α_1 which captures the average effect of the expansion on Medicaid enrollment for the eligible low-income population due to the expansions.[14] Table 1 reports estimates for the pre-Medicare sample of individuals aged 60-64. Consistent with previous literature (e.g., Miller, Johnson, and Wherry 2021 and Frean, Gruber, and Sommers 2017), the results in Table 1 demonstrate that the expansion had a large positive affect on enrollment for the targeted population.

The coefficient estimate on *Medicaid Expansion*×*Low-income* in column 1 suggests that enrollment into Medicaid for low-income individuals in expansion states increased by roughly 17 percentage points after the expansions. Recent work by Goodman-Bacon (2021a) and others shows that the results estimated via Equation (1) can be biased in the presence of time-varying treatment effects, since $\hat{\alpha}_1$ is estimated using variation in treatment (i.e. expansion) timing across states. Given that most states expanded in 2014 and that the post period window is fairly short, I do not expect this to be a major issue. However, to address this potential bias, column 2 presents results where I limit my main estimating sample to individuals living in states that expanded Medicaid in 2014 or at no point throughout the sample period.[15] This estimate is somewhat larger in magnitude, but qualitatively similar to the coefficient estimate in column 1.

Columns 3 and 4 present additional results where I include interactions between prior insurance coverage and the expansion and low-income terms. The results in these columns show that the majority of low-income individuals gaining Medicaid coverage through the expansions were previously uninsured. For example, the coefficient on the triple interaction of *Expansion* × *Low-income* × *Uninsured*$_{t-1}$ in column 4 suggests that, following the expansions, enrollment in Medicaid for low-income previously uninsured individuals increased by 28 percentage points. Alternatively, the coefficient on the triple interaction with an indicator for being previously privately insured is much smaller in magnitude and statistically insignificant. This suggests that the scope of private insurance crowd out (i.e., privately insured people switching to Medicaid) for the near-elderly population is relatively small. This is somewhat contradictory to prior work that has documented substantial crowd-out of private coverage due to the introduction of public insurance (Cutler and Gruber 1996). There are a number of reasons that could explain this difference. First, my sample is different from those previously studied, and the near-elderly may be less likely to actively drop employer sponsored coverage for Medicaid than the sample

[14]Defining eligibility according to the law is difficult in the HRS. As a result, I take a number of approaches. The first is to come up with my best guess of eligibility according to the poverty level threshold. The second is to separate out individuals who are in the lower three deciles of the sample income distribution and flag them as "low-income". For the main text of the paper I present results using the latter definition, however the results are robust to either. The primary reason I adopt the latter as my indicator for eligibility, is that it better predicts actual Medicaid enrollment. I take this as suggestive that the latter definition, while less grounded in the law, is the more empirically relevant definition of the eligible population due to potential errors in my eligibility determination process. Ultimately, these measures are highly correlated (correlation coefficient of 0.77) and the main results are robust to either definition, as shown in Appendix F.

[15]27 states expanded in 2014 out of 31 states that expanded during the entire sample period.

Table 1. Effect of expansion status on Medicaid enrollment

	(1)	(2)	(3)	(4)
Medicaid Expansion × Low-income	0.171***	0.192***	0.039	0.040
	(0.041)	(0.043)	(0.089)	(0.089)
Expansion × Low-income × Uninsured$_{t-1}$			0.264***	0.281***
			(0.097)	(0.097)
Expansion × Low-income × Private$_{t-1}$			0.056	0.071
			(0.089)	(0.088)
Sample	Whole sample	2014 expansions only	Whole sample	2014 expansions only
Controls included	✓	✓	✓	✓
Individual FEs	✓	✓	✓	✓
Obs.	18871	17096	17962	16273

Notes.—Table reports regression coefficient derived from separate triple-difference regressions of Medicaid enrollment onto exposure to the Medicaid expansions (Equation (1)). The models in columns 3 and 4 also include interactions between the triple-difference terms and indicators for prior insurance status (private insurance or uninsured). Columns 1 and 3 report results when including the entire pre-Medciare sample. Columns 2 and 4 present results when the estimating sample is limited to states that either expanded Medicaid in 2014 or at no point throughout the sample period. Each regression equation also includes controls for demographics (gender, age, age squared, race, educational attainment, marital status, smoker status, and employment status), health controls (self-reported health and counts of ADLs and IADLs), indicators for household income and household asset deciles, individual fixed-effects, and a set of fully-interacted indicators for survey wave, state of residence, and low-income eligibility status. Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

of low-income pregnant women studied in Cutler and Gruber (1996). Second, my measure of past insurance coverage is only a noisy measure of the type of coverage individuals would have available to them in the current period. Specifically, it is possible that the previously uninsured would have chosen private insurance in the current period had the Medicaid expansions not occurred. In this sense, the results in columns 3 and 4 of Table 1 are likely a lower bound on the amount of crowd-out.

To further highlight the role that the expansions played in Medicaid coverage for the near-elderly, Figure 2 reports results from the following event study specification on the pre-Medicare sample of individuals in states that expanded Medicaid in 2014 or did not expand during the sample period:

$$Medicaid_{i,t} = \rho + LI_{i,t} \times \sum_{y=2008}^{2016} \theta_y \mathbb{1}(t - t_s^E = y) + X'_{i,t}\Gamma + \chi + \epsilon_{i,t}, \tag{2}$$

where, as before, $LI_{i,t}$ is equal to one if individual i is in the low-income group in year t, and zero otherwise. The indicator variables $\mathbb{1}(t - t_s^E = y)$ denote the year of the wave for individuals in expansion states, whereas for non-expansion states these indicators are always zero.[16] I omit the

[16]I am able to treat 2014 as the treatment year for all observations because I limit the sample to individuals in states that either expanded in 2014, or at no point throughout the sample.

Figure 2. Variation in Medicaid coverage around Medicaid expansions

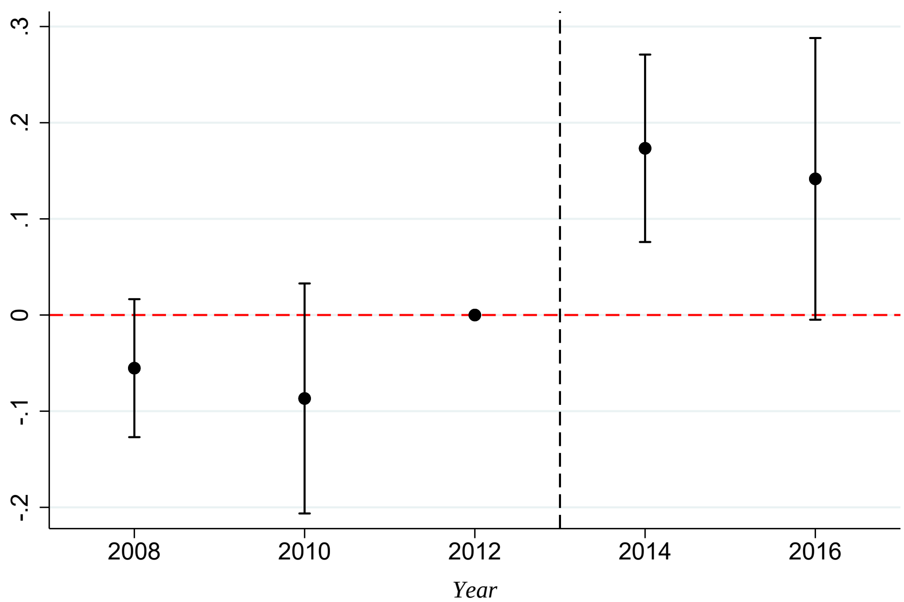

Notes.—Figure displays coefficient estimates and 95% state-level clustered confidence intervals from an event study regression of Medicaid coverage onto indicators for the survey wave relative to the Medicaid expansions (Equation (2)). The estimating sample is limited to the pre-Medicare sample of individuals in states that expanded Medicaid in 2014 or did not expand during the sample period. The regression model includes a rich set of observable controls for demographics and health status.

wave before the expansions take place ($y = 2012$) as the base year. Each $\hat{\theta}_y$ can be interpreted as the change in Medicaid coverage in wave y relative to the year before an expansion took place. χ denotes a fully interacted set of indicator variables for state, wave, and low-income status. These indicators difference out differences in outcomes within low-income groups across states and over time, and differences within states but across income groups over time.

As with Table 1, the results from the event study specification demonstrate that the Medicaid expansions had a economically (and statistically) significant effect on Medicaid enrollment for the near-elderly population that has persisted across time. Importantly, the figure also demonstrates that prior to the expansion periods, there were likely common trends across in Medicaid enrollment across the treatment and control groups and that there were no obvious anticipation effects.

4 Empirical Strategy

In this section I present my main reduced form strategy. First, I outline the ways in which I empirically test how the Medicaid expansions affected the near-elderly's healthcare consumption. Second, I present a strategy to test how past insurance coverage, due to the Medicaid expansions, affected future healthcare consumption once on Medicare.

4.1 Estimating the Utilization Response to the Expansions

Local Average Treatment Effect of Medicaid Coverage.— My first task is to determine how the Medicaid expansions affected healthcare use and expenditures. To do this I estimate a LATE of Medicaid coverage by using the policy change as an instrument for Medicaid enrollment with an instrumented triple-difference (TDIV) design.[17] I model this as follows:

$$Y_{i,t} = \pi_0 + \pi_1 \widehat{Medicaid}_{i,t} + \pi_2 Expand_{s,t} + \pi_3 LI_{i,t} + X'_{i,t}\Gamma + \Psi + \epsilon_{i,t} \quad (3)$$

where $Y_{i,t}$ is individual i's utilization outcome in survey-wave t; $\widehat{Medicaid}$ is the predicted measure of Medicaid coverage; $Expand_{s,t}$ is an indicator for whether state s expanded Medicaid by year t; $LI_{i,t}$ is a dummy variable indicating if individual i is in the target low-income group in year t; X is a vector of time-variant and invariant controls including flexible functions of household income and assets, self-reported health status, a third-order polynomial of age, indicators for the number of limitations to activities of daily living (ADLs) and instrumental activities of daily living (IADLs), gender, race, education level, marital status, an indicator for employment status, and an indicator for whether the individual was ever a smoker; and Ψ is a vector of fixed effects.[18] I estimate equation (3) by two-stage least squares (2SLS), where the first-stage is the triple difference-in-differences estimating equation shown in Equation (1).

The excluded instrument, $Expand_{s,t} \times LI_{i,t}$, is the interaction between the Medicaid expansion and target income group indicators.[19] In this model I interpret the coefficient of interest, π_1, as the causal effect of Medicaid coverage on Y for individuals who gain Medicaid through the expansion, and would not have had access to Medicaid coverage without the policy change.[20]

Triple Difference and Event Study Analyses.— In addition to estimating the LATE of Medicaid enrollment, I also estimate the effects of the Medicaid expansions on medical expenditures for the near elderly directly using triple difference-in-differences and event study designs. These strategies are common in the public health and Medicaid specific literatures, and help ground

[17] I discuss identification of the TDIV estimator later in this section and formally derive the estimator in Appendix D.

[18] For the majority of my analysis, the dependent variable of interest is the imputed natural log of total medical expenditures, $\widehat{\ln(E_{i,t})}$. In additional analysis I also include other measures of utilization such as the number of doctor's office visits. I exclude health variables from the X vector that could be a function of insurance status, such as medical diagnoses. For all specifications, Ψ includes state, year, state by year, state by low-income, and year by low-income fixed effects. In some specifications I also include individual level fixed effects.

[19] Note that the additional indicators $Expand_{s,t}$ and $LI_{i,t}$ appear in equation (3). This ensures that the variation Medicaid coverage that is driving identification arises directly from the policy change.

[20] In other words, π_1 identifies the LATE for low-income individuals in expansion states. With stronger assumptions on treatment effect homogeneity–namely that the effect of the expansions on Medicaid coverage *and* outcomes for low-income groups in expansion and non-expansion states are the same–it is also possible to interpret π_1 as the LATE of Medicaid coverage for the whole low-income population.

my LATE estimates of Medicaid coverage by providing information on the intent-to-treat (ITT) effect of the expansions. For the TD analysis, I estimate the reduced form equation:

$$Y_{i,t} = \beta_0 + \beta_1 Expand_{s,t} \times LI_{i,t} + \beta_2 Expand_{s,t} + \beta_3 LI_{i,t} + X'_{i,t}\Omega + \Psi + \varepsilon_{i,t} \tag{4}$$

where each element is similar to those in Equation (1). Under the standard TD identifying assumptions, the estimand β_1 summarizes the effect of the Medicaid expansion on medical expenditures for individuals in the target income group across all expansion years. As previously discussed, in my main analysis I limit my sample to the states that expanded in 2014 or did not expand at any point throughout the sample period.

In order to further explore how treatment may evolve over time, I also use an event-study design to see how the Medicaid expansions affected medical expenditures for the low-income near elderly. Specifically, I estimate a version of Equation (2), but with $Y_{i,t}$ as the dependent variable.

4.2 Long-Run Consequences of the Expansions

The Effect of the Expansions on Healthcare Expenditures in Retirement.—Next I explore how exposure to Medicaid due to the expansion affects future expenditures. To do so, I first estimate how first year Medicare enrollees' expenditures vary based on their previous (pre-Medicare) Medicaid status using an TDIV strategy analogous to (3):

$$Y_{i,t} = \alpha_0 + \alpha_1 \widehat{Medicaid}_{i,t-1} + \alpha_2 Expand_{s,t-1} + \alpha_3 LI_{i,t-1} + X'_{i,t}\Lambda + \psi + \varepsilon_{i,t}, \tag{5}$$

where $X_{i,t}$ is a vector of individual time-variant and invariant characteristics including sets of lagged health indicators (diagnoses and self-reported measures), gender, race, education level, marital status, an indicator for smoking history, and flexible functions of household income and non-housing assets; ψ is a set of indicator variables for the state that i was living in the previous wave, wave, lagged state by wave, lagged state by lagged low-income status, and wave by lagged low-income status; and $\widehat{Medicaid}_{i,t-1}$ is the predicted Medicaid status in the previous period generated from the following first stage equation:

$$Medicaid_{i,t-1} = \pi_0 + \pi_1 Expand_{s,t-1} \times LI_{i,t-1} + \pi_2 Expand_{s,t-1} + \pi_3 LI_{i,t-1} + X'_{i,t}\Gamma + \psi + \epsilon_{i,t}. \tag{6}$$

Unlike Equation (3), in this formulation I only observe each individual in their first wave of Medicare once; however, because the data are longitudinal I include controls for past health-status in X which helps mitigate concerns of selection into treatment. That said, if negative selection into Medicaid prior to the first Medicare wave does occur (i.e. that less healthy individuals are more likely to be enroll in Medicaid due to the expansions), I would expect $\hat{\alpha}_1$ to be upward biased. I discuss the implications of this potential bias further in Section 5.

Equation (5) only takes into account Medicaid coverage in the year right before going onto Medicare. Given that the data are available through 2018, there are two waves in which first-year Medicare beneficiaries could have been exposed to the expansions. To explore how additional years of exposure affect future spending I estimate the following equation:

$$Y_{i,t} = \beta_0 + \beta_1 \sum_{j=1}^{2} \widehat{Medicaid_{i,t-j}} + X'_{i,t}\Lambda + \psi + \varepsilon_{i,t}, \tag{7}$$

with a first stage equation that is analogous to Equation (6). The estimate of β_1 describes the effect of having an additional year of Medicaid coverage on an individual's healthcare spending in their first period with Medicare coverage.

Additionally, since the majority of expansions occurred in 2014 and the data go through 2018, I can also explore how previous Medicaid coverage affects expenditures for individuals in their second wave of Medicare. To this end, I estimate a model analogous to Equation (5) but on the sample of individuals in their second period of Medicare coverage. In this case, the estimate will tell me how persistent the effects of previous Medicaid coverage are on future spending. For example, if receiving earlier insurance coverage through the Medicaid expansions lowers spending on the first year of Medicare, this may simply due to shifting pent-up demand from one period to another. However, if receiving care earlier is more beneficial for future health, it is possible that the effects of that spending could persist and future spending may remain lower for a prolonged period.

As with the pre-Medicare analysis, I also employ more standard TD and event-study strategies for the Medicare periods. Each of these approaches are analogous to those presented in (4) and (2); however, as with the rest of the Medicare analysis, I limit the sample to individuals who are enrolled in Medicare for the first (or second) time and my explanatory variables of interest are the interactions between lagged (or twice-lagged) indicators of expansion status and low-income status or event time in the case of TD or event-study equations, respectively.

4.3 Identification

Identifying Assumptions for TD.—Causal identification of the impact of the Medicaid expansions on medical expenditures comes from cross-state and within-state variation in eligibility over time. In this sense, the control groups in my analysis are (1) low-income individuals in non-expansion states and (2) high-income individuals in expansion states. The rationale for adopting a TD specification is that healthcare demand of low-income individuals across expansion and non-expansion may differ systematically due to time-variant differences in the distribution of health or other public services provided to low-income households.[21] By including the within-state control group, I am able to better control for differential trends in healthcare demand across expansion and non-expansion states.

More formally, identification of the TD estimator requires that the path of expenditures, conditional on covariates, across the target (i.e., low-income) and control populations trend similarly in non-expansion states and expansion states had they not expanded.[22] Ex-ante I have no reason to believe that this assumption would be violated. In particular, once controlling for potentially time-varying attributes of each state, it seems unlikely that states that chose whether to expand or not would have had different trends in expenditures or Medicaid enrollment given that the decision not to expand Medicaid was primarily partisan. Ultimately, it is impossible to know for certain if this assumption is met, but I do provide suggestive evidence its validity in Section 5.[23]

Identifying Assumptions for TDIV.—As discussed in Hudson, Hull, and Liebersohn (2017), DDIV models require assumptions that are standard across both the differences-in-differences and IV literatures. However, given that my first stage is actually a triple difference, the presentation in Hudson, Hull, and Liebersohn (2017) is incomplete. In addition to the standard IV assumptions of exogeniety, relevance, and monotonicity; identification for the DDIV estimator requires several additional assumptions, the most important being parallel trends for both treatment and outcomes. In contrast, my setting requires the weaker TD parallel trend assumption for both expenditures (i.e., outcomes) and Medicaid enrollment (i.e., treatment). In Appendix D I formally derive identification for the TDIV estimator for the just-identified case (which is the relevant case for Equations (3) and (5)).[24]

In addition to parallel trends, the TDIV estimator also requires two exclusion restrictions on the instrument, which are common with the DDIV estimator. First, the instrument (i.e., the interaction between expansion status and income group) should not affect outcomes or treatment prior to the introduction of the policy. This would be violated if individuals decide to consume less healthcare leading up to the expansion since they anticipate receiving more generous coverage in the future. Similarly, if the years leading up to the expansions resulted in a surge in Medicaid enrollment because of increased discussion of Medicaid in the popular press, this assumption could be violated.[25] Violations of this assumption could also take the form of individuals moving across states and low-income status in response to the expansions. However, as discussed above, there is little evidence that individuals were acting strategically in either way.

[21] For example, Vermont (an expansion state) provides much larger public transfers to low-income households than Georgia (a non-expansion state), and it is possible that these differences evolve over time and differentially affect healthcare demand.

[22] For more details, see Olden and Møen (2022).

[23] I provide further discussion of the TD identifying assumptions in Appendix C.

[24] Though I believe I am the first to formally demonstrate this (and possibly to use this estimator), the derivation closely follows the work in Hudson, Hull, and Liebersohn (2017).

[25] In my event-study analysis, I find neither of these trends in the data for states that expanded Medicaid.

The second exclusion restriction is that, as in the standard IV setting, the instrument must only affect the outcome of interest through Medicaid enrollment. This assumption could be violated if other state-level policies change around the same time as the expansions that effect expenditures. For example the assumption would be violated if a state decided to expand Medicaid while also increasing income subsidies to near-elderly individuals. In this case, the unobserved policy change could increase the household budget constraints for the near-elderly in that state allowing them to increase their medical expenditures. Though such a policy would violate this exclusion restriction, I am unaware of any such state policy. Another potential concern is that the ACA changed other aspects of the insurance landscape for the working-age population. However, most of the policy changes that the ACA brought about were at the federal level, suggesting that, conditional on observables (including state-by-wave dummies), all states experienced similar changes from the ACA at around the same time. The Medicaid expansions were a unique aspect of the ACA which allowed states to decide if and when the policy would be enacted. With that in mind, and given my rich set of controls, I find this assumption plausible.

The final key assumptions for TDIV in this setting are that the effect of the instrument on Medicaid enrollment is monotone and that the instrument is relevant in determining Medicaid enrollment. Though monotonicity is not directly testable, I believe it is a reasonable assumption in this context. Specifically, since the expansions made enrolling in Medicaid easier for anyone who was previously eligible as well as expanded eligibility to more people, it is unlikely that expanding Medicaid reduced the likelihood of Medicaid enrollment for low income individuals. Further, the relevance of the instrument is testable and the results presented in Section 3 suggest that this assumption is also satisfied.

Overall, I take multiple econometric approaches to answering my two questions of interest, with the goal of presenting a complete story of how the Medicaid expansions affected healthcare spending for individuals over time. One of the benefits this holistic approach is that the different econometric have different limitations and, in fact, some help provide insight into the validity of the identifying assumptions of the others. For example, the common pre-trends assumption necessary for the TDIV and TD strategies is visually tested in the event-study estimates.

5 Empirical Analysis

5.1 Initial Response to the Expansions

Table 2 presents my estimates for the effects of the Medicaid expansions and Medicaid coverage on total imputed medical expenditures. Columns 1-3 report estimates of the triple-difference term in Equation (4) (with corresponding state-level clustered standard errors reported in parentheses and the implied Kennedy corrected percentage change in brackets) for different model specifications and estimating samples.[26] The coefficient estimate in column 1, implies an ITT of the Medicaid expansions on medical expenditures of roughly 30 percent for the full pre-Medicare sample. In order to mitigate the concerns related to the staggered roll-out of the Medicaid expansions, column 2 presents results when the sample is restricted to individuals living in states that expanded in 2014 or at no point during the sample period. The coefficient estimate implies that total expenditures for low-income individuals increased roughly 33 percent following the Medicaid expansions. Finally, column 3 presents results from a model that includes individual fixed-effects estimated on the same sample as column 2, and implies an increase in total expenditures of about 18 percent. To help contextualize these findings, consider that the average level of total imputed expenditures in the pre-Medicare sample is roughly \$3,400 per wave, which implies that total expenditures increased by roughly \$620-\$1,100 for low-income individuals in expansion states. Comparing results across the columns, it appears that including the entire sample introduces little if any bias due to the staggered nature of the expansions. Additionally, when controlling for time-invariant individual level observables, the estimated effect is somewhat muted. This may be suggestive that, even conditional on observable health, there is still adverse selection into the Medicaid program. However, all three of the coefficient estimates are within the others' confidence interval.

Columns 4-6 present the TDIV coefficient estimates on Medicaid coverage in Equation (3) (with Equation (1) as the first-stage) across various model specifications on the same estimating samples as columns 1-3. The first-stage Kleibergen-Paap cluster-robust F-statistics are presented in the bottom of the table and demonstrate that the instrument is sufficiently powered to explain Medicaid enrollment (the values of the F-stats range from 18.2-25.5, depending on the specification).[27] Unsurprisingly given that these are estimates of the LATE, the results in columns 4-6 are qualitatively consistent with the ITT estimates in columns 1-3 but larger in magnitude. For example, the coefficient estimate in columns 5 and 6 suggest that total medical expenditures increased by 230 and 112 percent, respectively, for individuals who gained Medicaid coverage through the expansions. These estimates imply increases in expenditures of roughly \$7,820 and \$3,800, respectively. These numbers are substantially larger than those found in previous work (e.g., Finklestein et al. 2012).[28] There are a number of reasons why these estimates, though larger, are plausible. First, the population I am looking at is older and likely sicker than the sample of individuals in their setting. Therefore, if individuals are increasing their healthcare consumption

[26] The bracketed terms ($\hat{p}$) are calculated according to:

$$\hat{p} = exp(\hat{\beta} - .5Var(\hat{\beta})) - 1$$

.

[27] Given that the model is just-identified, the Kleibergen-Paap F-statistic is the relevant measure of instrument strength. For details see Andrews and Stock (2018).

[28] Though the point estimates are relatively large, the confidence intervals overlap with estimates in previous work.

TABLE 2. EFFECT OF EXPANSION STATUS ON MEDICAL SPENDING

	TD (1)	TD (2)	TD (3)	TDIV (4)	TDIV (5)	TDIV (6)
Medicaid expansion × Low-income	0.265^{**}	0.292^{**}	0.172^{*}			
	(0.106)	(0.098)	(0.100)			
	[0.296]	[0.333]	[0.182]			
Medicaid coverage				1.329^{**}	1.228^{***}	0.886^{*}
				(0.528)	(0.449)	(0.515)
				[2.286]	[2.087]	[1.124]
Sample	Whole sample	2014 expansions only	2014 expansions only	Whole sample	2014 expansions only	2014 expansions only
Controls included	✓	✓	✓	✓	✓	✓
Individual FEs			✓			✓
First-stage F-stat				18.2	25.51	20.82
Obs.	20050	18197	17081	20050	18197	17081

Notes.—Table reports regression coefficient derived from separate triple-difference (TD) and instrumented triple-difference (TDIV) regressions of imputed log total medical expenditures onto exposure to the Medicaid expansions. The results in columns 1-3, correspond to Equation (1). The results in columns 4-6 correspond to Equation (3)). Columns 1 and 4 report results when including the entire pre-Medciare sample. All other columns present results when the estimating sample is limited to states that either expanded Medicaid in 2014 or at no point throughout the sample period. Each regression equation also includes controls for demographics (gender, age, age squared, race, educational attainment, marital status, smoker status, and employment status), health controls (self-reported health and counts of ADLs and IADLs), indicators for household income and household asset deciles, and a set of fully-interacted indicators for survey wave, state of residence, and low-income eligibility status. Columns 3 and 6 include individual fixed-effects as well. Standard errors clustered at the state level in parentheses and Kennedy corrected percentage change estimates are presented in brackets. Kleinbergen-Paap first-stage F-statistics clustered at the state-level reported. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

on the extensive margin, they are probably more likely to discover additional health problems or experience complications, which is likely to increase the cost of additional care relative to a younger and healthier population. Second, since this change in access to Medicaid is only for a short period of time, the individuals in my sample may be more price responsive due to the fact that current Medicaid out-of-pocket prices are lower than what they will experience once turning 65 and enrolling in Medicare. In fact, prior evidence indicates that the relative difference between current and future prices plays a significant role in determining current utilization decisions (Lin and Sacks 2019).

Figure 3 presents results from the event study specification in Equation (2). The point estimates in the figure are consistent with the results presented in Table 2 and demonstrate that the effect of the expansions on spending for low-income individuals seems to be consistent over time (with all three post-period point estimates hovering at around 0.3). The figure also lends credibility to the identification strategy given the lack of any noticeable pre-trends. Importantly, there does not appear to be any noticeable decrease in expenditures in the years prior to the expansions in anticipation of the access to future Medicaid coverage, which is a critical assumption for the identification of the TDIV and TD estimators.

Figure 3. Variation in total expenditures due to Medicaid expansions

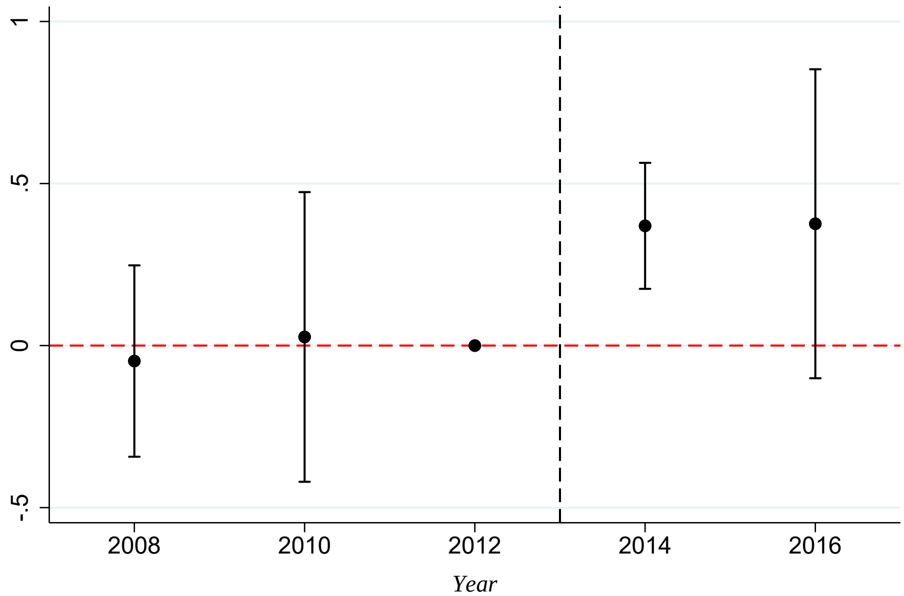

Notes.—Figure displays coefficient estimates and 95% state-level clustered confidence intervals from an event study regression of imputed log total medical expenditures onto indicators for the survey wave relative to the Medicaid expansions interacted with low-income status (Equation (2)). The estimating sample is limited to individuals in states that expanded Medicaid in 2014 or did not expand during the sample period. The regression model includes a rich set of observable controls for demographics and health status.

In Appendix E, I demonstrate that the spending response documented in Table 2 appears to be driven by increases in the likelihood of going to the doctor, receiving any home healthcare, and for previously uninsured individuals, an increase in the likelihood of using prescription drugs regularly. Interestingly, I do not observe any immediate effects on health outcomes following the Medicaid expansions. In Appendix F, I also demonstrate the robustness of the results in Table 2 to alternative definitions of the eligibility criteria and an alternative version of imputed expenditures. Overall the results are robust to these various specifications, with point estimates that are quantitatively similar and not statistically different from each other. I take all of these results as providing consistent evidence that the Medicaid expansions lead to an economically significant increase in medical expenditures for the low-income near-elderly population.

5.2 Effect of Medicaid Coverage on Future Outcomes

Overall, the results presented so far indicate that there was a large positive consumption response to the increased Medicaid coverage generated by the expansions. This is consistent with prior literature on health insurance provision and with the moral hazard hypothesis—when the relative price of healthcare that consumers face falls due to increased insurance coverage, they consume more. However, the primary aim of this study is to examine how this past insurance

TABLE 3. EFFECT OF PRIOR EXPANSION STATUS ON FIRST WAVE OF MEDICARE IMPUTED EXPENDITURES

	(1)	(2)	(3)
	A. Triple-difference model		
Medicaid expansion$_{t-1}$ × Low-income$_{t-1}$	-0.312**	-0.326**	-0.250**
	(0.138)	(0.145)	(0.113)
	[-0.275]	[-0.286]	[-0.226]
Sample	Whole sample	2014 expansions only	2014 expansions only
Health$_{t-2}$ controls			✓
Obs.	9847	8949	8032
	B. Instrumented triple-difference model		
Medicaid coverage$_{t-1}$	-2.078*	-1.757*	-1.278**
	(1.192)	(0.933)	(0.644)
	[-0.938]	[-0.888]	[-0.774]
Sample	Whole sample	2014 expansions only	2014 expansions only
Health$_{t-2}$ controls			✓
First-stage F-stat	7.9	12.75	15.72
Obs.	9847	8949	8032

Notes.—Table reports regression coefficient derived from separate TD and TDIV regressions of imputed log total medical expenditures during the first wave of Medicare onto prior exposure to the Medicaid expansions. The results in panel A, correspond to Equation (6). The results in panel B correspond to Equation (5)). Column 1 reports results when including the entire first wave of Medciare sample, and columns 2 and 3 present results when the estimating sample is limited to states that either expanded Medicaid in 2014 or at no point throughout the sample period. Each regression equation also includes controls for demographics (gender, age, age squared, race, educational attainment, marital status, and smoker status), indicators for household income and household asset deciles, and a set of fully-interacted indicators for prior survey wave, prior state of residence, and prior low-income eligibility status. Column 3 includes all of the same controls, but adds additional controls for twice lagged health status (self-reported health, BMI, and indicators for high-blood pressure, diabetes, cancer, heart disease, lung problems, and stroke diagnoses). Standard errors clustered at the state level in parentheses and Kennedy corrected percentage change estimates are presented in brackets. Kleinbergen-Paap first-stage F-statistics clustered at the state-level reported. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

coverage affects future outcomes. To do this, I begin by estimating Equation (5) as well as the analogous TD model with a number of dependent variables. I first focus on individuals in their first wave of Medicare coverage. Then I move on to explore whether the effects persist over time. For each of these samples I estimate the effect of past Medicaid coverage on medical expenditures as well as other measures of utilization and health.

Medical Expenditures.—Table 3 presents results on the effect of past exposure to the Medicaid expansions on medical expenditures while newly enrolled in Medicare.[29] Columns 1 and 2 of panel A present coefficient estimates for the triple-difference term in the reduced form TD model estimated on the whole sample and the sample limited to either 2014 expansion states or non-expanding states. Column 3 of the panel presents results from a specification that includes a

[29]In Table 3 I limit attention to the version of imputed medical expenditures that includes insurance status, but the results are extremely similar in terms of magnitudes and precision across both measures. I present the results for the measure that excludes insurance in the imputation as well as alternative specifications in Table 22 of Appendix F.

FIGURE 4. VARIATION IN TOTAL EXPENDITURES ONCE ON MEDICARE

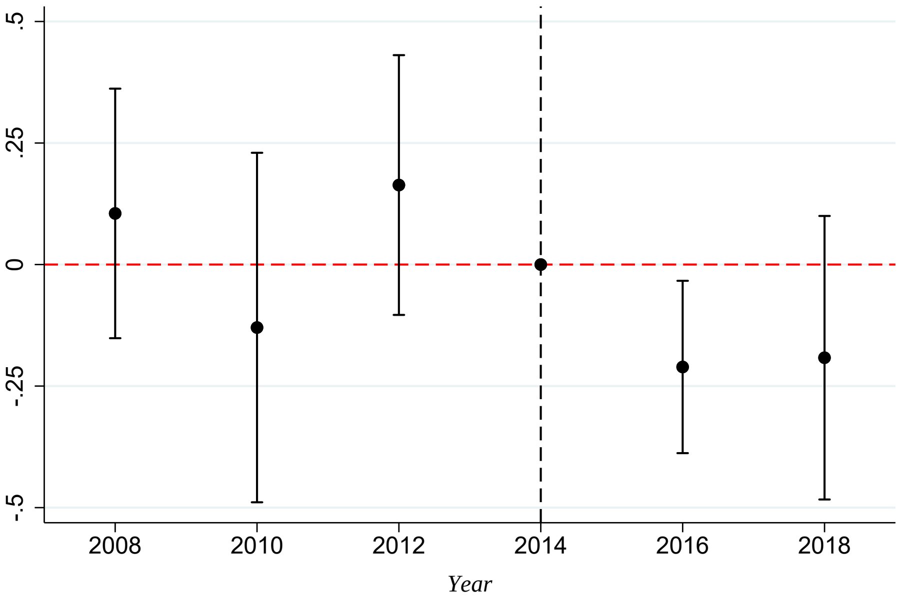

Notes.—Figure displays coefficient estimates and 95% state-level clustered confidence intervals from an event study regression of imputed log total medical expenditures during the first wave of Medicare coverage onto the indicators for the survey wave relative to the Medicaid expansion interacted with previous low-income status. The estimating sample is limited individuals in states that expanded in 2014 or did not expand during the sample period. The regression model includes a rich set of observable controls for demographics and lagged health status.

rich set of controls for an individual's past health estimated on the 2014 expansion sample.[30] The results presented in panel A show that prior exposure to the Medicaid expansions caused a statistically and economically significant decline in medical expenditures for first time Medicare beneficiaries. Specifically, the estimates in columns 1, 2, and 3 indicate that total expenditures fell by 27.5, 28.6, and 22.6 percent, respectively. Given that the mean imputed expenditure for the first wave of Medicare sample is roughly \$3,600, these coefficient estimates imply that total expenditures declined by between \$815 and \$1,030.

Panel B of Table 3 presents the coefficient estimates on (instrumented) Medicaid coverage in the previous wave estimated across the same samples and analogous TDIV models as in panel A. As with the estimates in panel A, all three of these estimates are negative, economically larger, and statistically significant. Columns 1-3 reflect that imputed expenditures decreased by 93, 89, and 77 percent, respectively. This corresponds to declines in total expenditures between roughly \$2,600 and \$3,150. The cluster-robust first-stage Kleibergen-Paap F-statistics presented at the bottom of the panel show a slightly weak first-stage for the whole sample model, which may explain the relatively large point estimate in column 1. Comparing the results in columns

[30]The rational for including twice lagged health measures is that these are unlikely to be affected by the expansions. Table 24 in Appendix F presents additional specifications with thrice lagged health measures and the results are very similar to those presented in Table 3.

Table 4. Effect of cumulative Medicaid coverage on first wave Medicare spending

	(1)	(2)
	A. Triple-difference model	
One wave of Medicaid expansion		
× Low-income	-0.168	-0.166
	(0.119)	(0.125)
Two waves of Medicaid expansion		
× Low-income	-0.469**	-0.481**
	(0.213)	(0.213)
Sample	Whole sample	2014 expansions only
Health_{t-3} controls	✓	✓
Obs.	7683	6971
	B. Instrumented triple-difference model	
Sum of past Medicaid coverage	-0.666**	-0.614**
	(0.303)	(0.277)
Sample	Whole Sample	2014 expansions only
Health_{t-3} controls	✓	✓
First-stage F	105.57	119.74
Obs.	7683	6971

Notes.—Table reports regression coefficient derived from separate TD and TDIV regressions of imputed log total medical expenditures during the first wave of Medicare coverage onto the prior exposure to the Medicaid expansions. In panel A, the coefficient estimates of interest indicate the intent to treat of different years of exposure to the Medicaid expansions. Panel B reports the LATE of an additional year of Medicaid coverage (Equation (7)). Column 1 reports results when including the entire first wave of Medciare sample, and column 2 presents results when the estimating sample is limited to states that either expanded Medicaid in 2014 or at no point throughout the sample period. Each regression equation also includes controls for demographics (gender, age, age squared, race, educational attainment, marital status, and smoker status), indicators for household income and household asset deciles, and thrice lagged health controls (self-reported health, BMI, and indicators for high-blood pressure, diabetes, cancer, heart disease, lung problems, and stroke diagnoses) and a set of fully-interacted indicators for prior survey wave, prior state of residence, and prior low-income eligibility status. Standard errors clustered at the state level in parentheses. Kleinbergen-Paap first-stage F-statistics clustered at the state-level reported. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2 and 3 demonstrates that conditioning on observable health is important given that doing so leads to an attenuation of the estimated decline in expenditures due to past Medicaid coverage.

As with my previous findings, I also examine the robustness of these results by estimating an event-study specification. Figure 4 presents the results. Overall the figure corroborates the findings presented in Table 3—individuals who were exposed to the Medicaid expansions prior to turning 65 had lower initial medical expenditures once on Medicare. Moreover, the fact that the point estimates in the post-period are consistent in magnitude suggests that this effect was similar for both cohorts of new Medicare enrollees that were affected by the expansions in my sample (i.e., individuals who were 63-64 and 61-62 in 2014). Finally, though the pre-treatment coefficients are somewhat noisy, there is no discernible pre-trend which (at least visually) validates my identifying assumptions.

Dose Response to Past Medicaid Coverage.—Given the panel nature of the data, I can also observe the extent to which the response of first wave Medicare expenditures vary based on years of exposure to the expansions. Table 4 presents results where the independent variables of interest capture the cumulative exposure to the expansions as presented in Equation (7). Columns 1 and 2 of panel A, respectively, present results from the reduced form TD model where the effects of each level of exposure (e.g., one wave or two waves) are separately estimated on the whole pre-Medicare sample and the sample including only 2014 expansion and non-expansion states.[31] Regardless of the estimating sample, the results in the panel show that more years of expansion exposure lead to larger declines in Medicare expenditures. Columns 1 and 2 of panel B present estimates for the effect of (instrumented) cumulative Medicaid coverage prior to Medicare estimated on the same samples. These coefficient estimates also suggest that the decrease in Medicare expenditures is increasing in the duration of past Medicaid coverage. Specifically, the point estimate of -0.614 in column 2 of panel B suggests that each additional year of Medicaid coverage reduces first wave Medicare expenditures by roughly 46 percent (which translates to approximately $1,600 based on the mean of predicted expenditures in the sample). These findings affirm those previously discussed and provide the additional insight that the potential dynamic benefits of insurance coverage (in terms of decreased expenditures) depends on the length of coverage. This is consistent with a hypothesis that prolonged spells of uninsurance increase the likelihood of forgoing preventative care or care that aids in the management of chronic diseases.[32]

Utilization in the First Wave of Medicare.—As with the pre-retirement analysis, I also explore how other measures of utilization respond to prior Medicaid coverage in the sample of newly eligible Medicare enrollees. Columns 1 and 2 of panel A in Table 5 present results from TD and TDIV models, respectively, where I estimate the effects of Medicaid coverage in the previous wave on current utilization measures. All of the coefficient estimates are negative and larger in magnitude, which is consistent with my imputed expenditure findings above. However, only the coefficients from the models with more than five doctor visits, any at home healthcare, and (marginally) total out-of-pocket spending are statistically significant at traditional significance levels.[33] The fact that there is a strong effect on doctor visits on the intensive but not extensive margin is encouraging (nor is there a detectable effect on prescription drug use). This may suggest that having prior coverage doesn't decrease the consumption of future preventative care, but–instead–reduces the need for more complex and costly treatment regimens. The signs on the coefficients for hospital use are also consistent with this hypothesis, though the estimates are

[31] In the table I define exposure as being in the low-income group in an expansion state in an expansion year. Models also include controls for thrice lagged health status.

[32] For recent evidence on this, see Chandra, Flack, and Obermeyer (2013) and Harris et al. (2022).

[33] I show in Table 24 in Appendix F that the point estimates on the hospitalization measures are more precise when excluding lagged health controls or including thrice (instead of twice) lagged health controls. I take these result, as well as those in Table 5 as jointly indicating a decline in the intensity of hospitalizations due to past Medicaid coverage.

Table 5. Effect of expansions on other outcomes during the first wave of Medicare

	TD (1)	TDIV (2)
A. ***Other utilization***		
Any doctor visits	-0.044	-0.225
	(0.052)	(0.266)
Doctor visits ≥ 5	-0.162**	-0.826**
	(0.069)	(0.389)
Any hospitalizations	-0.077	-0.393
	(0.059)	(0.305)
Hospitalizations (counts)	-0.137	-0.698
	(0.128)	(0.688)
Nights in hospital	-1.716	-8.763
	(1.322)	(7.342)
Any Rxs	-0.015	-0.076
	(0.051)	(0.261)
Any at home care	-0.050**	-0.255*
	(0.025)	(0.149)
Total out-of-pocket spending	-816†	-4165†
	(551)	(2797)
B. ***Health outcomes***		
Self-reported health	0.230**	1.174*
	(0.099)	(0.623)
Health index	0.014**	0.071*
	(0.006)	(0.038)
Sample	2014 expansions only	2014 expansions only
Health$_{t-2}$ controls	✓	✓
First-stage F-stat		15.84
Obs.	8028	8028

Notes.—Table reports coefficient estimates from regressions of different dependent variables onto exposure to the Medicaid expansions prior to enrolling in Medicare. Triple-difference (TD) and instrumented triple-difference (TDIV) regression models in columns 1 and 2, respectively. The coefficient estimates in each row correspond to models with different dependent variables. All utilization measures correspond to utilization over the previous two years. "Self-reported health" is a subjective measure of health ordered from 1 ("Fair") to 5 ("Excellent"). "Health index" is a measure of health constructed from the first principal component of a number of subjective and diagnostic health measures. See Appendix A for more details. Each regression equation also includes controls for demographics (gender, age, age squared, race, educational attainment, marital status, and smoker status), indicators for household income and household asset deciles, and twice lagged health controls (self-reported health, BMI, and indicators for high-blood pressure, diabetes, cancer, heart disease, lung problems, and stroke diagnoses) and a set of fully-interacted indicators for prior survey wave, prior state of residence, and prior low-income eligibility status. Standard errors clustered at the state level in parentheses. Kleinbergen-Paap first-stage F-statistic clustered at the state-level is reported. † $p < 0.15$, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

imprecise. Overall, these results help validate the previous findings for imputed expenditures and demonstrate that, broadly speaking, healthcare consumption across a number of measures appears to fall when individuals have more generous insurance coverage before going on Medicare, all else equal.

Health in the First Wave of Medicare.—In addition to understanding how Medicare expenditures are affected by previous insurance coverage, I am also interested in what drives these effects. As proposed in Fang and Gavazza (2011), a likely mechanism leading to lower future healthcare consumption is that individuals who get more generous insurance coverage before retiring may enter the retirement phase with better health, due to larger or more frequent health investments before turning 65. To test if this mechanism is important, I also analyze the effect of past insurance coverage on self-reported health and a general health index in the first period of Medicare.[34] Panel B of Table 5 presents the results.

Remarkably, the point estimates across both measures of health and both estimation approaches show that prior exposure to the Medicaid expansions and Medicaid coverage has large positive and statistically significant effect on health outcomes when on Medicare. The point estimates in column 1 indicate an ITT of roughly one fifth and one quarter of a standard deviation increase in self-reported health and the health index, respectively. And the point estimates in column 2 suggest that past Medicaid coverage increased self-reported health and the health index by approximately 1 standard deviation and 1.8 standard deviations, respectively. Together these results demonstrate that individuals who received insurance through the expansions did, in fact, enter the Medicare phase healthier than those unaffected, all else equal. These findings suggest that the primary channel for lower future expenditures is through increased health, which is brought about—in part—from increased healthcare consumption prior to turning 65.[35]

Persistence.—Given that the data go through 2018, there is a cohort of individuals who had been affected by the expansions and that have been on Medicare for at least two waves. In order to see how persistent the effects of past insurance coverage are on future spending and health, I conduct additional analysis on this sample of Medicare enrollees who are further into their Medicare tenure. Table 6 reports results from the TD and TDIV models estimated on this sample, where the independent variables of interest are twice lagged expansion exposure and Medicaid coverage. Columns 1 and 2 report the TD and TDIV results, respectively, and each row of the table corresponds to a regression model with a different dependent variable. Given the lack of explanatory power of the instrument in the first-stage, as indicated by the first-stage F-statistics presented in the table, I am hesitant to interpret the TDIV coefficient estimates as causal. That said, I think that for both expenditures and health status, taken together with the results from the TD models, Table 6 strongly suggests that the dynamic effects of generous insurance coverage are persistent. In other words, the increased healthcare consumption in pre-retirement that was induced by the Medicaid expansions depress spending and increase health for several years into future.

[34] I construct a general health index in the spirit of Poterba, Venti, and Wise (2013) and Poterba, Venti, and Wise (2017). The index is given by the first principal component a large list of diagnostic and subjective measures of health and well-being.

[35] Figure 14 in Appendix E presents results from an event-study specification where the health index is the dependent variable. These results are consistent with those in Table 5 and reflect common trends in the dependent variable prior to the expansion treatment.

TABLE 6. EFFECT OF EXPANSIONS ON EXPENDITURES AND HEALTH DURING THE SECOND WAVE OF MEDICARE

	TD (1)	TDIV (2)
A. ***Expenditures***		
$\widehat{\ln(E_{i,t})}$	-0.322**	-3.118
	(0.146)	(2.727)
B. ***Health outcomes***		
Self-reported health	0.307**	2.983**
	(0.104)	(1.158)
Health index	0.016**	0.158***
	(0.005)	(0.045)
Sample	2014 expansions only	2014 expansions only
Health$_{t-3}$ controls	✓	✓
First-stage F-stat		2.06
Obs.	6937	6937

Notes.—Table reports coefficient estimates from regressions of different dependent variables onto various measures of exposure to the Medicaid expansions prior to enrolling in Medicare. Triple-difference (TD) and instrumented triple-difference (TDIV) regression models in columns 1 and 2, respectively. The coefficient estimates in each row correspond to models with different dependent variables. "$\widehat{\ln(E_{i,t})}$" is an imputed measure of log expenditures that includes insurance coverage in the imputation. "Self-reported health" is a subjective measure of health ordered from 1 ("Fair") to 5 ("Excellent"). "Health index" is a measure of health constructed from the first principal component of a number of subjective and diagnostic health measures. See Appendix A for more details. Each regression equation also includes controls for demographics (gender, age, age squared, race, educational attainment, marital status, and smoker status), indicators for household income and household asset deciles, and thrice lagged health controls (self-reported health, BMI, and indicators for high-blood pressure, diabetes, cancer, heart disease, lung problems, and stroke diagnoses) and a set of fully-interacted indicators for prior survey wave, prior state of residence, and prior low-income eligibility status. Standard errors clustered at the state level in parentheses. Kleinbergen-Paap first-stage F-statistic clustered at the state-level is reported. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 Further Discussion of Empirical Findings

The results presented in this section point to two important phenomena. First, the exogenous enrollment in Medicaid spurred by the expansions under the ACA led to an increase in total medical expenditures for those affected. Second, individuals who were exposed to the Medicaid expansions prior to going onto Medicare ended up consuming fewer healthcare services, spent less out-of-pocket, accumulated fewer total medical expenditures, and reported better health outcomes once on Medicare compared to similar individuals where were unaffected by the policy change. However, the appropriate way to interpret these results is unclear. My findings are consistent with are two non-mutually exclusive hypotheses.

The first hypothesis is that additional healthcare consumption is productive in improving future health which lowers the demand for healthcare services in future periods. The second hypothesis is that the change in future expenditures that I detect is due to an inter-temporal substitution of care; where, upon receiving Medicaid, individuals front load healthcare consumption that would have been realized on Medicare in the absence of the expansions. While the data do not allow me to definitively distinguish between these two channels without further assumptions, the cumulative evidence of health improvements, effect persistence, and that more periods of coverage leads to larger declines in expenditures all provide suggestive evidence for the former.

Additional evidence presented in Appendix E suggests that Medicaid coverage increased utilization of care that improved information about chronic conditions and increased access to chronic disease management, all of which lead to lower expenditures and improved health once enrolled in Medicare.[36]

6 A Stylized Model of the Medicaid Expansion

In this section, I develop a stylized "late in life"-cycle model of the Medicaid expansion. The goals of which are to provide insight into how to interpret the evidence presented in Section 5, to estimate the welfare effects of the expansions, and to evaluate a counterfactual reform to the Medicare program. Features of the model are drawn from a number of different sources including De Nardi, French, and Jones (2016), Finkelstein, Hendren, and Luttmer (2018), Lockwood (2018), Low and Pisteferri (2015), and Ozkan (2017).[37] There are also several innovations in the model which I describe in detail in the proceeding section.

Individuals enter the model as workers at $t_0 = 55$, retire at $t_R = 65$, and live for up to $T = 95$ periods.[38] In this way, the model is divided into a pre-Medicare (working) phase and a Medicare (retirement) phase.[39] Knowing that Medicare coverage is in their future, working individuals choose insurance coverage, how much healthcare and all other goods to consume, and how much wealth to accumulate in each period before retiring.[40] At $t = t_R$ they no longer work, but they receive Social Security payments and Medicare coverage which changes the out-of-pocket price they face to varying degrees depending on their pre-retirement insurance. In order to replicate my empirical setting of the Medicaid expansions, I use the model to simulate a treatment and control group of individuals who start at age 55 with the same set of initial conditions. The treatment group is provided the opportunity to enroll in Medicaid starting at age 62 if they meet the income eligibility criteria.

[36]Table 20 in Appendix E shows that though the expansions improved future health (as measured by self-reported health and the health index), the share of people diagnosed with diabetes also rose. I interpret these seemingly conflicting pieces of evidence as reflective of an increase in *diagnoses* but not an increase in *prevalence*. In this sense, I suspect that the expansions led to more diagnoses and subsequent treatment for chronic conditions such as diabetes, which likely plays a role in reducing future expenditures.

[37]An alternative modeling approach implemented by Yang, Gilleskie, and Norton (2009) and Khwaja (2010) is to treat all choices within the model as discrete.

[38]By modeling retirement as exogenous, I am implicitly assuming that the Medicaid expansions do not affect retirement choices. In the HRS I do not detect any relationship between the expansions and retirement decisions.

[39]Henceforth, I will use the terms working phase and pre-Medicare phase interchangeably. I will also use the terms retirement phase and Medicare phase interchangeably.

[40]Each period in the model is a year.

6.1 Environment

Preferences.—In each period t, individuals derive utility from the consumption of non-medical goods and services c_t, a continuous measure of health h_t, and Medicaid enrollment $M_t \in \{0,1\}$ according to:

$$u(c_t, h_t, M_t),$$

and they discount future utility by β. I do not assume that utility is additive separable in consumption and health. Instead I assume that utility takes the following form:

$$u(c_t, h_t, M_t) = A + \frac{(\gamma c_t^{\theta} + (1-\gamma) h_t^{\theta})^{\frac{1-\sigma}{\theta}}}{1-\sigma} - \Psi M_t \tag{8}$$

where $\sigma > 0$ is the coefficient of relative risk aversion, θ is the substitution parameter between consumption and health, γ is the share of utility loaded onto consumption, A is a positive constant that ensures the value of being alive is non-negative, and Ψ denotes a utility penalty that individuals pay if enrolled in Medicaid ($M_t = 1$).[41] The primary motivation for allowing health and consumption to be non-separable comes from Finkelstein, Luttmer, and Notowidigo (2013), who find that the marginal utility of consumption is likely higher when individuals are in good health.

Health and Death.—In the spirit of Grossman (1972), I treat health as a stock that evolves according to a partially stochastic technology. For simplicity, I assume that health is continuous on the unit interval (i.e., $h_t \in [0,1]$).[42] Therefore, health in each future period (t) is given by:

$$h_t = \max\left\{\min\left\{f(h_{t-1}, m_{t-1}, \varepsilon_t), 1\right\}, 0\right\}$$

where h_{t-1} is the health stock inherited from the end of the previous period, m_{t-1} is the health care services (in dollars) consumed in the previous period, and ε_t is a health shock that is realized at the beginning of the current period. I assume that the function, $f(\cdot,\cdot,\cdot)$, takes the following form:

$$f(h_{t-1}, m_{t-1}, \varepsilon_t) = \delta h_{t-1} + \lambda m_{t-1}^{\alpha} - \varepsilon_t. \tag{9}$$

Here, δ denotes how much the health stock and medical investments depreciate from period to period and λ and α are, respectively, scale and curvature parameters governing the productivity of medical investments in generating future health. The health shock, $\varepsilon_t \geq 0$, enters the production function by (weakly) decreasing the inherited stock of health augmented by prior

[41]The A term follows from Hall and Jones (2007). Specifically, since estimates of σ are typically greater than one, the second term in utility is negative. Since I will be normalizing the value of death to be zero, I include A so that agents are motivated to stay alive in the model. I discuss the Medicaid enrollment choice in detail later on.

[42]Given this assumption, I standardize the constructed health index to lie within this range as well.

investments. For ease of exposition, denote the inherited stock of health augmented by prior investments as $\bar{h}_t$ (where $\bar{h}_t = \delta h_{t-1} m_{t-1}^{\alpha}$). The distribution of health shocks depend on age (t)

$$\varepsilon_t \sim F_{\varepsilon}(\varepsilon; t),$$

which I assume is log-normal with parameters μ_t and σ_t $(log(\varepsilon_t) \sim \mathbb{N}(\mu_t, \sigma_t))$.

As shown above, individuals may choose to consume healthcare because it increases future health, increasing future utility. I also assume that investments in health decrease future mortality risk. Specifically, I map the end-of-period health stock into a function that determines the likelihood of surviving to the next period, $\Phi(h_t, t)$. I follow Chetty et al. (2016) by adopting a Gompertz specification where the log of the hazard rate, $1 - \Phi(h_t, t)$ is assumed to be a linear function of the health stock:

$$\begin{aligned} \log[1 - \Phi(h_t, t)] &= \phi_{1,t} + \phi_{2,t} h_t \\ \Phi(h_t, t) &= 1 - \exp(\phi_{1,t} + \phi_{2,t} h_t). \end{aligned} \tag{10}$$

Therefore, any investments in health also reduced the risk of death in future periods which is a key driver of the demand for healthcare.

Insurance.—Before becoming eligible for Medicare, there are three possible insurance plans: a private plan (P), no insurance (U), or Medicaid (M). Each plan $i \in \{P, M, U\}$ is summarized by a premium, p_i, and a cost sharing structure, $\pi_i(\cdot)$, which assigns the out-of-pocket payments associated with each level of total expenditures:

$$\pi_i(m) = \begin{cases} m \text{ if } m <= d_i \\ \psi_i(m - d_i) + d_i \text{ if } m > d_i \end{cases} \tag{11}$$

where d_i and ψ_i are the annual deductible and coinsurance rates, respectively, associated with plan i. Relative to the private plan, Medicaid offers both a lower premium and more generous cost-sharing. The uninsured pay no premiums ($p_U = 0$) and have to pay full price for any medical consumption ($\pi_U = 1$).[43] I assume that the private insurance plan and uninsurance arrive stochastically at the beginning of each period according to a first-order Markov process.[44] I denote an insurance draw with $j_t \sim dF_j(j_t | j_{t-1}, I)$, such that the transition probabilities from insurance to uninsurance differ by income status. Moreover, I impose that private insurance is highly persistent (which would be the case if individuals receive employer sponsored insurance)

[43] While the uninsured do technically have to pay full price for medical care, there are a number of reasons why this is not true in practice. For example, Hadley and Holahan (2003) find that roughly 35 percent of all care received by the uninsured in 2001 was not paid for. In estimation, I account for this by assuming $\psi_U < 1$, but much higher than the other plans.

[44] This is a simplification that I make to capture the fact that private insurance is often tied to employment, rather than actively purchased. However, I plan to change in future versions of the paper. In particular, I plan to allow agents to choose to enroll in private insurance, rather than them exogenously receive it.

and that uninsurance is also persistent, but somewhat less so.[45]

If eligible, an individual may choose to give up their existing plan and replace it with Medicaid ($M_t = 1$), where eligibility is determined based on labor earnings and whether Medicaid has been expanded at time t. However, if the individual does enroll, they must pay the utility cost Ψ, introduced previously. This term is meant to capture both the pecuniary and non-pecuniary burdens associated with enrolling in the program.[46] Once an individual reaches retirement age ($t = t_R$) they switch onto Medicare, which is summarized by premium p_R and cost sharing $\pi_R(\cdot)$. Thus, in retirement individuals no longer make any choice related to insurance.

Income, Assets, and Taxes.—In the pre-retirement phase ($t < t_R$), agents supply labor inelastically and earn a per-period income that is determined by a combination of permanent income (I) and an unobserved stochastic productivity draw (ω_t):

$$w_t = I \times (1 + \omega_t) \text{ for } t = 0, \ldots, t_R - 1$$

where productivity evolves according to a five state first-order Markov process:[47]

$$\omega_t \sim \mathbb{P}(\omega_t | \omega_{t-1}).$$

Once retired ($t \geq t_R$), agents receive a (fixed) transfer from the Social Security Administration (SSA) according to a schedule $S(\cdot)$ which provides some level of replacement based on permanent income during one's working years:

$$w_t = S(I) \text{ for } t = t_R, \ldots, T,$$

where the replacement rate is decreasing in permanent income.

In addition to earning labor income and receiving Social Security payments, agents can also accumulate assets, a_t, which earn a risk-free return, r. For a given insurance plan i, assets evolve according to

$$a_{t+1} = a_t + w_p(ra_t + w_t, \tau) + b_t - \pi_i(m_t) - p_i - c_t, \tag{12}$$

where $w_p(ra_t + w_t, \tau)$ denotes post-tax income, τ describes the tax structure, and b_t denotes government transfers (in addition to SSA payments while retired).[48] Further, I assume that agents cannot borrow (i.e., $a_t \geq 0$). As is standard in the literature, I assume that non-SSA government transfers provide a consumption floor which will vary across working years and retirement:

$$b_t = \max\left\{0, \underline{c} + \pi_i(m_t) + p_i - (a_t + w_p(ra_t + w_t, \tau))\right\}.$$

[45] In practice, I estimate transitions from private insurance to uninsurance directly from the data.
[46] For example, see Fox, Feng, and Reynolds (2023).
[47] I summarize the distribution of these shocks with a persistence component, ρ, and a transitory component, σ^2_ω.
[48] Note that the above equation defines the per-period budget constraint.

Figure 5. Model timing while working

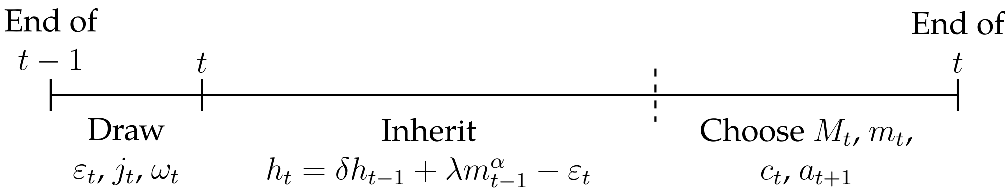

Notes.—Figure provides an outline of the timing in the model.

The prior equation states that government transfers ensure that an individual's per-period wealth less medical expenditures never falls below the consumption floor, $\underline{c}$. Following De Nardi, French, and Jones (2010), I assume that these transfers are only positive if next period's assets are zero.

Timing.—The timing of the model is outlined in Figure 5. Before the start of each period, the individual inherits their augmented health capital $\bar{h}_t$ and receives a health shock ε_t, which jointly determine health in the period h_t and if the agent will survive to $t+1$.[49] The period begins and the agent receives an initial insurance plan (either private insurance or no insurance) and a productivity draw ω_t. If labor earnings are below the Medicaid eligibility threshold and the individual is located in an expansion state, the agent can choose whether to enroll in Medicaid. Following that choice, the agent decides how to divide their cash on hand between health augmenting medical consumption m_t, consumption c_t, and savings a_t. Instantaneous utility is then realized given h_t and c_t and the period concludes. The timing during the Medicare phase is similar, except there are no insurance or productivity draws and no choice over Medicaid coverage.

6.2 Value Functions and Solution Method

For any set of parameter values, I solve the model numerically through backward induction beginning at the maximum attainable age, T. During the Medicare eligible phase, individuals inherit their insurance coverage, so the only control variables are consumption (c_t), medical expenditures (m_t), and assets (a_t). Given the environment outlined above, the individual's problem throughout the retirement phase can be written recursively with the following value function:

$$V_t(h_t, a_t, I) = \max_{c_t, m_t, a_{t+1}} u\big(c_t, h_t, j_t = R\big) + \beta \int V_{t+1}\big(h_{t+1}, a_{t+1}, I\big) dF_{\varepsilon}(\varepsilon_{t+1}|t) \tag{13}$$

[49]Notably, this formulation allows me to exclude the health shock from the list of state variables because h_t fully describes an agent's health status at a given point in time.

subject to

$$\begin{aligned} c_t + a_{t+1} + \pi_R(m_t) + p_R \leq a_t + \tau(S(I) + ra_t) \\ h_{t+1} = f(h_t, m_t, \varepsilon_{t+1}) \\ V_{T+1} = 0, \end{aligned} \tag{14}$$

where β is the discount factor and the final constraint normalizes the value of death to zero. For the final working period, the individual's value function is given by

$$V_t(h_t, a_t, j_t, \omega_t, I) = \max_{\substack{c_t,\, m_t,\, a_{t+1}, \\ M_t \,\in\, \{0,1\}}} u\big(c_t, h_t, M_t\big) + \beta \int V_{t_R}\big(h_{t_R}, a_{t_R}, I\big) dF_\varepsilon(\varepsilon_{t_R}|t_R) \tag{15}$$

subject to

$$\begin{aligned} c_t + a_{t+1} + \pi_i(m_t) + p_i \leq a_t + \tau(I(1+\omega_t) + ra_t) \\ h_{t_R} = f(h_t, m_t, \varepsilon_{t_R}). \end{aligned} \tag{16}$$

For every other working period, an individual's value function is given by

$$\begin{aligned} V_t(h_t, a_t, j_t, \omega_t, I) = \max_{\substack{c_t,\, m_t,\, a_{t+1}, \\ M_t \,\in\, \{0,1\}}} & u\big(c_t, h_t, M_t\big) \\ & + \beta \int\int\int V_{t+1}\big(h_{t+1}, a_{t+1}, j_{t+1}, \omega_{t+1}, I\big) dF_\varepsilon(\varepsilon_{t+1}|t) dF_j(j_{t+1}|j_t) dF_\omega(\omega_{t+1}|\omega_t) \end{aligned} \tag{17}$$

s.t.

$$\begin{aligned} c_t + a_{t+1} + \pi_i(m_t) + p_i \leq a_t + \tau(W(1+\omega_t) + ra_t) \\ h_{t+1} = f(h_t, m_t, \varepsilon_{t+1}) \\ h_0 = h \text{ and } a_0 = a \text{ given.} \end{aligned} \tag{18}$$

7 Model Estimation

In this section I discuss the identification of the model parameters and estimation. Estimation proceeds in two steps. First, I estimate or calibrate a subset of model parameters that can be determined directly from the data or taken from the literature, including the discount factor β and the coefficient of risk aversion σ. In the second step, I estimate the following Θ vector of

Table 7. Externally calibrated parameters

Description	Parameter	Value	Source
Discount factor	β	0.96	Estimates in the literature.
Coefficient of risk-aversion	σ	2.5	Estimates in the literature.
Risk-free rate of return	r	0.025	Ozkan (2017).
Consumption floor pre-65	$b_{t<65}$	\$2000	
Consumption floor 65+	$b_{t\geq 65}$	\$3500	
Income process			
Income persistence	ρ	0.90	Russo (2023).
Variance in transitory income shock	σ^2_ω	0.09	Russo (2023).
Insurance transitions			
Probability of remaining uninsured (low-income)		0.84	
Probability of remaining privately insured (low-income)		0.82	Author's calculations.
Probability of remaining uninsured (high-income)		0.70	
Probability of remaining privately insured (high-income)		0.90	
Cost-sharing			
Uninsured			
Coinsurance	π_U	0.7	
Deductible	d_U	\$0	
Premiums	p_U	\$0	
Private			
Coinsurance	π_P	0.32	
Deductible	d_P	\$1500	
Premiums	p_P	\$3000	
Medicaid			Author's calculations.
Coinsurance	π_M	0.07	
Deductible	d_M	\$20	
Premiums	p_M	\$200	
Medicare			
Coinsurance	π_R	0.25	
Deductible	d_R	\$300	
Premiums	p_R	\$1500	

Notes.—Table presents the model parameters that are set externally. The insurance transition rates are determined from coefficient estimates from the regression model presented in Equation (19) estimated using data from the MEPS. "High-income" and "low-income" refer to individuals in the second and first permanent income quintiles, respectively. Coinsurance rates are calculated based on differences in out-of-pocket spending and total expenditures by insurance coverage status in the MEPS.

parameters

$$\Theta = \{\underbrace{\theta, \gamma, A, \Psi,}_{\text{Preferences}} \underbrace{\delta, \lambda, \alpha,}_{\text{Health production}} \underbrace{\mu_t, \sigma_t,}_{\text{Health shocks}} \underbrace{\phi_{1,t}, \phi_{2,t}}_{\text{Mortality risk}} \}$$

via indirect inference.[50]

[50]For both the health shock and mortality risk parameters, I make a simplifying assumption to reduce the parameter space. Namely, I estimate values of the parameters at every 5 years, and approximate the fully age-varying formulation with a linearly spaced vector between each of those values.

7.1 First-Step Calibration

Insurance Transitions and Cost Sharing.—In order to determine the exogenous transitions across private insurance and uninsurance during working years, I estimate the following regression separately by permanent income groups for individuals aged 55 to 64 in the MEPS:

$$Insurance_{i,t} = \beta_0 + \beta_{Insurance} Insurance_{i,t-1} + u_{i,t}, \tag{19}$$

where $Insurance$ can be an indicator for either private or uninsured, depending on the state. I then use the estimated values of $\hat{\beta}_{Insurance}$ to construct the transition matrix. The probabilities for low- and high-income groups of remaining uninsured or privately insured are presented in Table 7.[51]

I also use the MEPS to determine the coinsurance rates associated with each type of insurance plan. In order to do this, I divide total out-of-pocket expenditures by total expenditures for each individual separately by the insurance status, which yields an estimated set of coinsurance rates. I also externally set the values of premiums and deductibles for each plan based on information I gathered from various sources. For example, the associated premiums for Medicare come from the total annual premiums you would need to pay in order to enroll in Medicare Part B.

Other Fixed Parameters.—I set the discount factor to $\beta = 0.96$ and the coefficient of risk aversion to $\sigma = 2.5$, which are values common in the literature.[52] I also externally set the parameters governing the wage process, the tax and SSI structure, the risk-free rate of interest, and the consumption floor. Though I could estimate the wage process directly in the HRS, I opt instead to take values from the literature. I use values for the persistent and transitory wage components from Russo (2023). The values for the tax schedule are set based on the 2012 tax brackets and federal marginal tax rates. I simplify the social security reimbursement rate schedule to roughly match the estimates from actuaries at the Social Security Administration (Clingman, Burkhalter, and Chaplain 2022). I set the interest rate to 2.5 percent. Finally, I set the consumption floor while old to \$3,500, which falls between estimates in the literature.[53] I assume that the consumption floor while young is less generous, and I set it to \$2,000.

[51] Note that "low-income" and "high-income" refer to the first and second quintiles of the permanent income distribution, respectively.

[52] For example, in a model with heterogeneous preferences, De Nardi, Pashchenko, and Porapakkarm (2017) estimate a discount factors of 0.88 and 0.99, Achou (2023) estimates a value of 0.95, and Lockwood (2018) calibrates the discount factor to 0.975. For the risk-aversion parameter, many studies assume a value of $\sigma = 2$ (e.g., Hall and Jones 2007) and work specifically focusing on retirees often finds values of σ larger than 3 (e.g., De Nardi, French, and Jones 2010). Thus, I believe the assumption of $\sigma = 2.5$ is a reasonable and somewhat conservative value.

[53] For example, De Nardi, French, and Jones (2010) estimate that the consumption floor for retirees is \$2,700, but De Nardi, French, and Jones (2016) estimate a value of \$4,600.

7.2 Estimation of Structural Parameters

Sample Initialization.—Given that I do not observe the entire life-cycle of each individual, I determine the initial distribution of state variables at age 55 in a way that is designed to accomplish two goals: (1) ensure that the sample generating the simulated analogs to the results presented in Section 5 draw initial states consistent with what I see in the data for the estimating sample and (2) that the groups that may not have been relevant in my empirical analysis, but *are relevant* in the counterfactual analyses are simulated based on accurate initial conditions in the data. I satisfy (1) by drawing initial observations (a collection of initial state variables) directly from the data. In particular, I jointly draw (randomly and with replacement) 5000 values for assets, health, permanent income, and insurance status from a subset of individuals aged 53-57 who are in the bottom two income quintiles. Based on the policy functions determined through the backward induction, I then simulate each of these 5000 observations twice, once in a control group and once in a treatment group, where the treated group receives (unexpectedly) the Medicaid expansions at age 62. I include the bottom two quintiles instead of just the bottom quintile to satisfy (2). Though many of the individuals will never end up eligible for Medicaid, I include them in the model to provide a reasonable group with which to simulate policy counterfactuals. For example, if I want to determine the effects of increasing the Medicaid eligibility threshold, I need to make sure that my estimated parameters are appropriate for individuals in a higher income group.

Indirect Inference Procedure—. I estimate the remaining structural parameters, Θ, via indirect inference.[54] Indirect inference relies on matching moments from an *auxiliary model* which can be estimated on real data as well as simulated data generated within the structural model.[55] Formally, the indirect inference estimator is

$$\hat{\Theta} = \underset{\Theta}{\text{argmin}}[\mathbf{m}(\Theta_0) - \mathbf{m}(\Theta)]' W [\mathbf{m}(\Theta_0) - \mathbf{m}(\Theta)]$$

where $\mathbf{m}(\Theta_0)$ and $\mathbf{m}(\Theta)$ are vectors of moments generated by the auxiliary model on the true data and the simulated data at a set of model parameter values Θ, respectively. W is a weighting matrix which I set in a way that appropriately scales the moments and puts more emphasis on moments that contribute to identification of the key structural parameters.[56]

[54] See Gourieroux, Monfort, and Renault (1993).

[55] This auxiliary model need not describe the correct data generating process as long as the moments it generates are related to the structural parameters of interest. The vector of parameters, Θ, can then be estimated by minimizing the distance between moments of the auxiliary model estimated on the observed data and those generated from the simulated data. If the structural model is correctly specified, any bias generated by the estimated auxiliary model will be reproduced when the same model is estimated on the simulated data.

[56] Specifically, I put a weight of 2 on each of the medical expenditure life-cycle moments, a weight of one on each of the health and mortality life-cycle moments, and a weight of ten on each of the regression based moments. The intuition for this is that there are only 9 causal regression based moments (compared to 120 life-cycle profile moments), but they are critical for identification of the structural parameters. I put additional weight on the ex-

What remains, then, is to choose which auxiliary moments to match. In the model individuals make choices over Medicaid enrollment, healthcare expenditures, consumption, and savings. Therein, I choose moments that closely reflect these choices and that provide a link between observed choices and structural parameters of interest. In particular, I use: (i) the evolution of total medical expenditures over the life-cycle; (ii) the evolution of assets over the life-cycle; (iii) the evolution of health by age; (iv) survival probabilities by age; and (v) a series of regressions of the effect of the Medicaid expansion on enrollment into Medicaid, initial medical expenditures, future medical expenditures, and future health as presented in Sections 4 and 5.

Moments: Expenditure, Health, and Survival Profiles.—I first construct the medical expenditure-age profile using the MEPS. Given that I currently only model low-income individuals who could potentially become eligible for Medicaid, I limit construction of the asset-age profile to individuals in the bottom two quintiles of the income distribution. I then calculate the mean of total expenditures at each age. Following De Nardi, Paschenko, and Porapkkarm (2017), I then scale mean medical expenditures up by a factor of 1.7 in order to make medical expenditures in the model consistent with aggregate spending measures in the National Health Expenditure Accounts. One limitation of MEPS data is that there are very few individuals in the sample over the age of 85. As a result, the expenditure profile for the most elderly in the sample is very noisy. I attempt to overcome this issue by forecasting expenditures out from age 85 to 95 as a second-order polynomial of age.

I construct life-cycle profiles of health and mortality using the HRS. As with the medical expenditure profile, I limit construction of the asset and health by age profiles to individuals in the bottom two quintiles of the permanent income distribution.[57] With the sample limited to low-income individuals, I construct the health-age profile using the mean of the health index described in Section 3 at each age. I also rescale the health index to lie between 0 and 10 for numerical tractability. I do not limit the sample to low-income people when constructing the mortality rate by age because there are too few observations to generate a reliable description of mortality. However, using the entire HRS sample has been shown to generate mortality-age profiles that are consistent with national data sets (see Weir, 2016).

These age profiles provide me with 40 moments each, yielding 120 moments, which are summarized in Table 8.

Moments: Treatment Effects of the Medicaid Expansions.—The next set of moments that I use in estimation are the reduced form treatment effect estimates presented in Section 5. Specifically, I target the ATE of the expansions on Medicaid enrollment, the ITT and LATE of the expansions

penditure moments in order to derive accurate cost predictions from the model. This method deviates from the most common practice of using the inverse of the variance-covariance of the data moments, but is similar to the method used in Vreugdenhil (2023). In future iterations of the paper I will explore how sensitive my results are to alternative weighting schemes.

[57]I calculate permanent income in the HRS as the average income within individual for all years that they are observed in the data.

Table 8. Estimation moments

Moment description	Targeted parameters
Life-cycle profiles	
Mortality rate by age in the HRS (40 moments)	$A, \phi_{1,t}, \phi_{2,t}, \sigma_t$
Health profile of low-income individuals in the HRS (40 moments)	$\delta, \phi_{2,t}, \mu_t$
Medical expenditure profile for low-income individuals in the MEPS (40 moments)	$\lambda, \alpha, \theta, \gamma$
Treatment effects	
ATE of the Medicaid expansion on Medicaid enrollment	Ψ
ITT and LATE of the Medicaid expansion on initial expenditures (2 moments)	$\theta, \gamma, \lambda, \alpha$
ITT and LATE of the Medicaid expansion on expenditures during first Medicare wave (2 moments)	θ, δ
ITT and LATE of the Medicaid expansion on health during first Medicare wave (2 moments)	λ, α
ITT of the Medicaid expansion on expenditures during second Medicare wave (1 moment)	θ, δ
ITT of the Medicaid expansion on health during second Medicare wave (1 moment)	λ, α, δ

Notes.—This table presents the 129 moments that are used in estimation as well as the structural parameters that each moment is closely connected to in identification.

on initial medical expenditures, and the ITT and LATE of the expansions on expenditures and health once on Medicare.[58] In the model, a period is a year, but in the data a period is two years. In order to replicate this fact with the simulated data, I collapse the simulated data down into waves (groups of two years). With the data in the correct panel structure, I estimate regressions analogous to those presented in Section 4 for the simulated pre-Medicare sample, as well as the simulated first and second wave of Medicare samples. Since, in the simulated data each initial observation will appear in both the treatment and control groups, I do not implement the TD and TDIV estimators on the simulated data. Instead, I estimate the DD and DDIV analogs. Ultimately, each estimator aims to recover the same empirical objects so the auxiliary moments are the same.

The estimates from the regression model provide me with another 9 moments which are listed in Table 8. In sum, there are 129 moments to identify 40 parameters.

7.3 Identification

Since the parameters are jointly identified, no single moment identifies a single structural parameter in the joint estimation. However, each parameter is closely linked to particular features of the data. In this section I outline which key moments provide identification of the structural parameters.

Preferences.—Two features of the data help pin down the elasticity of substitution between health and consumption, θ. First, medical spending patterns over the life-cycle provide information about how much consumption people are willing to trade-off for small improvements in health as health deteriorates. For example, if health and consumption were highly substitutable,

[58]In addition to targeting the ITT and LATE during the first wave of Medicare, I also target the ITTs for health and spending in the second wave. I choose not to target the estimated LATEs because the first-stage F-statistics for those estimates are quite small, so I am less confident in the validity of those moments.

agents would have little incentive to make increasingly costly health investments towards the end of life (at the expense of additional consumption) as health declines. Alternatively, in the extreme where health and consumption are perfectly complementary, large health investments become worthwhile even if it requires forgoing significant consumption. Second, θ is also identified by the reduced form estimates of the effect of Medicaid coverage on medical expenditures, since gaining Medicaid changes the relative prices of healthcare and other consumption. Thus, the extent to which individuals substitute away from consumption and to medical expenditures when the out-of-pocket price falls, provides additional information on the complementarity between health and consumption.[59]

As with θ, the term governing the share of utility coming from health and consumption, γ, is also identified from the response of medical expenditures to Medicaid and Medicare coverage. The overall level of medical expenditures also helps pin down γ, since total expenditures summarize the share of total resources devoted to health versus consumption. The value of A is identified from the average mortality rates over time. Larger values of A increase the flow utility of being alive at every age, and, therefore, increase the life expectancy that agents aim to achieve. The final preference term, Ψ, which denotes the disutility associated with enrolling in Medicaid, is primarily identified from the reduced form estimate of the effect of the Medicaid expansion on Medicaid enrollment. Therefore, Ψ allows the model to rationalize why many individuals who would otherwise benefit from Medicaid coverage (if eligible) choose not to enroll in a manner that is directly linked to the observed patterns outline in Section 3.

Health Production.—The health depreciation term, δ, is pinned down by the evolution of health over the life-cycle. If δ is too low, then individuals have less incentive to maintain high health because investments in health are highly transitory making maintaining even moderate levels of health very costly. The scale, λ, and curvature, α, parameters governing the degree to which medical expenditures produce health present a key identification challenge. This is due to the fact that the individuals with the highest demand for care are typically sick, which makes determining the relationship between health and healthcare consumption notoriously difficult. I am able to overcome this challenge by the inclusion of my reduced form estimates which provide me with a measure of the causal relationship between healthcare consumption and future health outcomes (see Table 5). Specifically, these two parameters are identified by the reduced form estimates for the effects of Medicaid coverage on expenditures and future health, where the combination of the initial expenditure response and the effect that has on future health outcomes helps to pin down these parameters.

Health Shocks and Mortality Risk.—Given that agents are risk-averse, they prefer to smooth their utility across time. This yields a result that most agents aim to maintain a relatively consistent level of health as they age. Therein, the degree to which average health declines in the

[59]The change in average expenditures due to Medicare coverage provides similar variation to help identify this term.

TABLE 9. SELECT ESTIMATED PARAMETERS

Preferences		Health production	
Parameter	Estimate	Parameter	Estimate
γ	0.593	δ	0.967
θ	-49.03	λ	0.121
A	19.03	α	0.077
Ψ	0.101		

Notes.—This table reports a select set of model parameters estimated via indirect inference.

population is directly related to the distribution of health shocks. Specifically, as the mean of the distribution decreases, average health declines and as the variance increases, more individuals are likely to experience particularly bad shocks which increase the pool of individuals who are a high mortality risk. In this sense, the means of the shocks are most closely related to the mean health while the variances are more related to the survival profiles. Finally, the mortality risk parameters are identified jointly by the health and mortality gradients. For example, a relatively steep health profile and a fairly flat mortality curve would be consistent with health becoming less important for mortality over time (i.e., $\phi_{2,t} > \phi_{2,t+1}$).

7.4 Parameter Estimates and Model Fit

In this section I first present the results from my estimation procedure and describe how well the model performs fitting the targeted moments. I also provide a discussion of external validity based on the estimated model's ability to match additional untargeted moments of economic significance.

Estimated Parameters.—Table 9 presents a select set of estimated parameters.[60] The estimate of $\gamma = 0.539$ is suggestive that consumption and health share somewhat equal weight in terms of resource expended on each, with a slightly higher weight on consumption. This is plausible given that I am focused on more elderly individuals who are much more likely to expend significant resources on health than younger, healthier individuals. The estimated value of $\theta = -49.03$ suggests that health and consumption are highly complementary in utility. Though not directly comparable, this is qualitatively consistent with the findings of Finkelstein, Luttmer, and Notowidigdo (2013) who find that health and consumption are likely complements for individuals in the HRS. The parameter defining the flow utility of being alive (A) is sufficiently large to ensure that utility is positive over relevant values of health and consumption. Finally, the disutility of enrolling in Medicaid, Ψ, is a non-trivial share of the flow-value of being alive (roughly 0.5 percent). This suggests that the hassle costs and administrative burdens associated with

[60]The remaining estimated parameters are presented in Table 25 of Appendix G. In future iterations of this paper I will include the bootstrapped standard errors of the parameter estimates.

Medicaid enrollment could play an important role in discouraging eligible individuals from enrolling.

The parameter governing health depreciation, $\delta = 0.967$, is close to one, indicating that health is highly persistent over time. This is consistent with other work that often takes health as exogenously determined, but finds that transition probabilities across health states are typically quite low (see De Nardi, Pashchenko, and Porapakkarm, 2017 and Aizawa and Fang, 2020). The values of the scale and curvature parameters, $\lambda = 0.121$ and $\alpha = 0.077$, respectively, are also reasonable and suggest that the marginal return to health expenditures is steeply decreasing. This is consistent with the argument that there is clearly a "flat of the curve" for healthcare, where an additional dollar of healthcare spending produces significantly less than a dollar in health benefits.

Model Fit.—Figure 6 displays the model fit for the life-cycle profiles where the data are plotted with solid lines and the simulations are plotted with dashed lines. Panel 6a shows the evolution of average medical expenditures over the life-cycle for low-income individuals in the MEPS.[61] Expenditures gradually increase as people age due to the deterioration of health. The mean of the health index at each age in the HRS is presented in panel 6b and demonstrates that health steadily declines as individuals age, which complements the trends in panel 6a. Finally, panel 6c plots the average mortality rate by age in the HRS.

Generally, the model is able to closely match the patterns across all three profiles. Importantly, the model also is able to reproduce the pattern that is observed in the data where individuals have a noticeable increase in expenditures around age 65, when they become eligible for Medicare. For health, the model simulations are slightly below the data, but the overall trend is accurate. The model struggles to generate the empirically observed mortality rates towards the end of life (i.e., people in the model live somewhat longer), however the model fits the data well over the relevant range of ages for my empirical setting. I suspect that the difficulty in fitting mortality at the end of life may be due to error in my expenditure prediction over the same ages. Specifically, my predictions likely underestimate the rate at which expenditures increase near the end of life.[62] Thus, the model struggles to rationalize the (relatively) low level of expenditures and higher mortality rates simultaneously.

Figure 7 displays the model fit for the regression based moments. Panel 7a presents the treatment effects related to Medicaid enrollment and spending. All of the model generated moments fall within the 95 percent confidence intervals, and many of them are also very close in magnitude to the targeted point estimates. On the other hand, the model has a difficult time matching the health related regression moments. Though the model based moments are near (and in some cases within) the 95 percent confidence interval of the data moments, they are smaller in

[61] As previously discussed, the MEPS expenditures are scaled by a factor of 1.7 and the data from ages 85-95 are forecasts.

[62] Saha (2021) shows that expenditures in the MCBS increase markedly after age 85.

Figure 6. Model fit of targeted life-cycle profiles

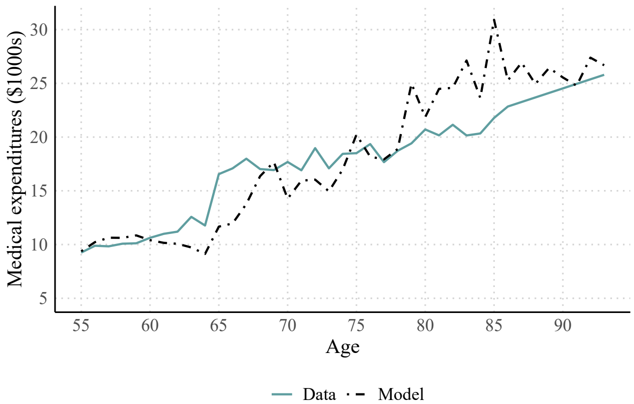

(a) Mean medical expenditure profile

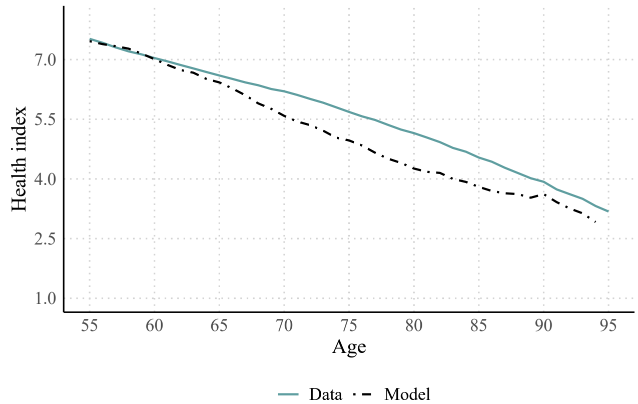

(b) Mean health profile

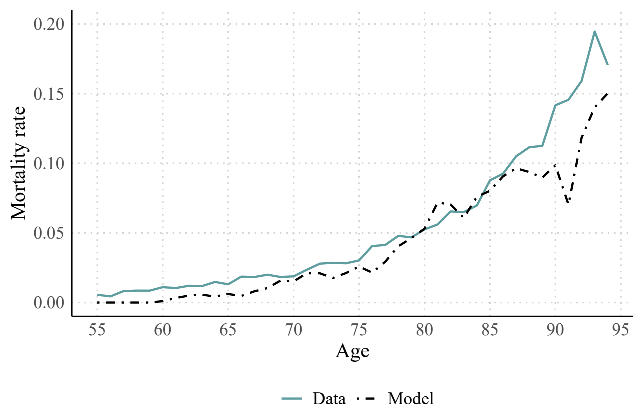

(c) Mortality rate

Notes.—Each figure plots the targeted data moments (solid lines) and the simulated analogs (dashed lines). Expenditure data is from the MEPS and inflated by a factor of 1.7 to better match aggregate measures in the National Health Expenditure Accounts. Expenditures after age 85 are forecasts. Age specific means of the health index and mortality rates are calculated in the HRS.

magnitude.[63] That said, the model generated moments are all positive and, themselves, statistically different from zero, which suggests that the model is still capturing the main underlying mechanisms identified in the empirical results.

External Validation.—To validate the model I also explore how the model replicates four key patterns observed in the data: (1) the effect of the Medicaid expansions on mortality, (2) the extent of adverse selection into Medicaid, (3) the heterogeneity in expenditures by health status, and (4) the heterogeneity in expenditures by past insurance coverage. As shown in Table

[63] I believe this could be due to discrepancies in the way I am measuring health in the data (i.e., the health index), and what units of health represent in the model. In future iterations of the paper, I plan to re-estimate the model targeting the elasticity between past insurance coverage and future health rather than the level change. Based on initial attempts at estimation using these new moments, I suspect that this will not qualitatively change the model's main results.

Figure 7. Model fit of targeted treatment effect moments

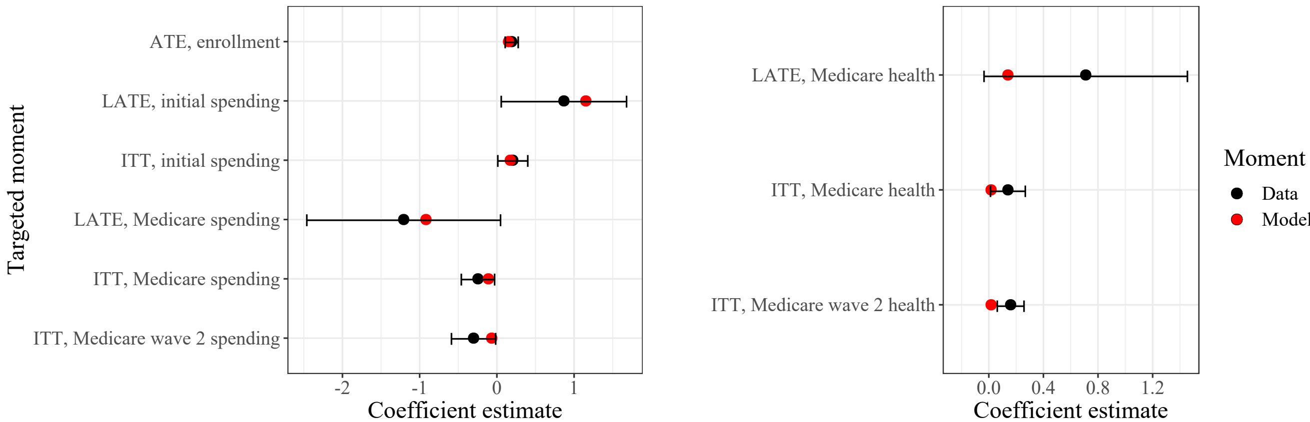

(a) Enrollment and expenditure moments (b) Health moments

Notes.—Figures plot the targeted regression based data moments and their 95 percent confidence intervals in black. The analogous simulation generated moment is plotted in red.

10, the model is able to replicate all of these patterns well. Existing work demonstrates that the Medicaid expansions significantly reduced mortality for people aged 55-64 (by 9.4 percent over the sample mean). The model reproduces a reduction in mortality of roughly 8.2 percent relative to the control group mean, indicating that the model is able to generate realistic effects of Medicaid on survival.[64] Additionally, in the data less healthy individuals are more likely to enroll in Medicaid following the expansions. They are also more likely to consume more care once enrolled. The model is able to match both of these patterns quite well, suggesting that the reaction to the expansions across health levels is consistent across the data and the model. As shown in Appendix Table 19, the previously uninsured also have a larger expenditure response to the Medicaid expansions, which the model also reflects. I discuss external fit further in Appendix G.

8 Assessment of the Medicaid Expansions

With the parameters of the model estimated, I now use the model to answer two key questions: (1) how much did expanding Medicaid to the near-elderly change welfare and (2) how much is this policy going to cost the government? To answer these questions, I conduct a cost-benefit analysis of expanding Medicaid to the near-elderly population. Specifically, I simulate 25,000 agents in both the treatment and control groups, where each set of agents are given the same initial conditions (drawn from the data) and are exposed to the same series of shocks. From the simulated data, I determine the welfare effects of the Medicaid expansions and the net govern-

[64] This is important for later exercises that aim to determine the value of the expansions or counterfactual policies, since decreased mortality is an important channel through which health insurance can increase lifetime utility.

TABLE 10. MODEL VALIDATION

Untargeted moment	Data	Model	Data source
Effect of Medicaid expansions on mortality	-0.094	-0.082	Miller, Johnson, and Wherry (2021)
Adverse selection into Medicaid	-0.065	-0.069	Author's calculations
Effect of interaction between Medicaid expansions and health on expenditures	-0.200	-0.081	Author's calculations
Effect of interaction between Medicaid expansions and uninsurance on expenditures	0.19	0.10	Table 19

Notes.—This table presents a series of untargeted empirical moments and their simulated analog. The data estimate for the effect of Medicaid expansion on mortality (row 1) denotes the relative change in mortality rates in expansion and non-expansion states for individuals aged 55-64. Simulated moment is calculated estimating the relative difference in mortality rates across the treated and non-treated simulated agents aged 62-65. The data moment in row 2 is the coefficient estimate from a regression model similar to Equation (1) that includes an interaction term between the TD term and the health index. The data moment in row 3 is estimated from a regression model similar to Equation (1) where predicted log expenditures is the dependent variable and that includes an interaction term between the TD term and the health index. The data moment in row 4 is estimated from a regression model similar to Equation (1) where predicted log expenditures is the dependent variable and that includes an interaction term between the TD term and an indicator for if the individual was previously uninsured. For more details on these moments, see Appendix G.

ment costs associated with expanding.

8.1 The Value of Medicaid

Welfare Measurement.—To quantify the value of the Medicaid expansions, I measure welfare in terms of the compensating variation (CV) that makes agents in the control group indifferent between remaining in the control group and switching into the treatment group at age 55. This notion of welfare provides a measure of the lump-sum transfer at age 55 the government would have to provide each individual in order to make them as well off as expanding Medicaid.[65] Formally, I calculate the asset differential Δ_a satisfying the indifference condition

$$V_{55}^{C}(h, a + \Delta_a, j, \omega, I) = V_{55}^{T}(h, a, j, \omega, I). \tag{20}$$

where V_{55}^{i} denotes the value of being in the treatment (T) and control (C) groups at age 55 with a given set of initial conditions for health (h), assets (a), insurance coverage (j), productivity (ω), and permanent income (I). The per-capita CV generated by the Medicaid expansions is reported in the first row of panel A in Table 11.[66] This average, however, masks substantial

[65]Note that because I assume the expansions were unanticipated, I restrict the treatment group's policy functions to be the same as the control group up to age 62. However, I calculate their value function based on the realized continuation values. Therein, the treatment group's age 55 values reflect the future value of Medicaid even though they do not internalize that value when making decisions prior to expansions.

[66]In order to arrive at a per-capita estimates of the costs and benefits, I divide the totals by the number of simulated agents and then multiply by 0.4 to account for the fact that I only model individuals in the bottom two income quintiles. This calculation makes the implicit and potentially strong assumption that for people outside of the model (i.e., individuals above the second permanent income quintile) do not benefit from the expansions, nor do they cost

TABLE 11. THE PER-CAPITA BENEFITS AND COSTS OF THE MEDICAID EXPANSIONS

A. Benefits		
CV of Medicaid expansions		\$3,142
NPV of uncompensated care transfer to providers		\$102
	Total benefit:	\$3,245
B. Costs		
NPV of Medicaid expenditures		\$2,853
NPV of dynamic fiscal externality		\$250
	Total net cost:	\$3,102
C. Change in welfare per net cost		
$\frac{CV + \text{NPV of transfers to providers}}{\text{Net costs}} = \mathbf{1.047}$		
D. Welfare decomposition		
CV of Medicaid expansions without moral-hazard		\$2,245
Additional CV due to moral-hazard		\$897

Notes.—This table reports the per-capita welfare benefits and net government costs of the Medicaid expansions for the near-elderly population. "CV" refers to welfare as measured by the notion of compensating variation presented in Equation (20). The result in panel D are determined by an exercise in which I fix agent's medical expenditure policies in a given period, but I allow them to enroll in Medicaid to lower their out-of-pocket prices. To calculate per-capita estimates I divide the simulated totals by the number of simulated agents and then multiply by 0.4 since I only include the bottom two income quintiles in the simulations.

heterogeneity with the 10th percentile of modeled agents placing no value on the Medicaid expansions and the 90th percentile valuing the expansions at \$15,190.[67]

Drivers of Welfare.—These results provide two new insights into *why* public health insurance is valuable. First, a large portion of the value of the expansions is due to Medicaid lowering the out-of-pocket price for costly investments in future health. To illustrate this, I fix the health policy functions for the simulated agents, but allow them to enroll in Medicaid (if eligible) and benefit from lower out-of-pocket prices. This means that the insurance does not lead to additional healthcare investments due to the lower prices.[68] Panel D of Table 11 presents the results

the government anything. The assumption that these people do not cost the government, seems reasonable since they are very unlikely to ever become eligible. However, the assumption that they do not benefit most likely leads me to understate the benefits of the program, since they are still risk-averse and benefit from the existence of the insurance even if their likelihood of using it is very low.

[67]The minimum and maximum valuations are \$0 and \$64,670, respectively. One important note is that the model does not increase taxes on higher earners in order to fund the program; so, by construction, the minimum possible valuation is \$0. If, however, taxes were to increase in order to implement the expansions, some individuals would certainly be made worse of. I plan to explore the importance of taxes in future iterations of the paper.

[68]Technically, this exercise allows people to consume more healthcare, but only due to intertemporal income effects. For example, Medicaid allow an agent to consume the same amount of healthcare at a lower cost, the savings from which can be immediately consumed or carried over as assets into the next period. If this savings response is large enough, then even with the original policy function, the optimal level of health investment has increased.

from this decomposition. As shown in the table, roughly 30 percent of the welfare generated by the expansions is due the fact that additional healthcare spending, i.e., moral hazard spending, increases future health.

Second, due to the level of risk-aversion in the population, the Medicaid expansions generate considerable value for many individuals who are unlikely to ever become eligible for coverage. This is consistent with findings in other work which finds that even high income households receive considerable insurance value from means-tested social insurance programs (Braun, Kopecky, and Koreshkova 2016). This also helps reconcile differences between my results and those in Finkelstein, Hendren, and Luttmer (2019), who find a fairly low willingness to pay for Medicaid for Oregon Health Insurance lottery winners.

8.2 Program Costs

To determine program costs I compare the amount of medical consumption that is paid for by the government across the simulated expansion and non-expansion groups. The share that the government pays depends on the type of insurance that an agent has (i.e., Medicare or Medicaid) and the amount of their total expenditures that are paid for by the insurer. To get a sense for how dynamics factor into program costs, I separately look at costs for the initial expansion period and then the Medicare period.

Initial Costs.—As expected and shown in Section 5, initial program costs are large. For the simulated treatment group, annual per-capita expenditures go up by \$300, of which \$279 is paid for by the government.[69] To help ground the estimate of initial costs note that simulated median annual government expenditures per-Medicaid enrollee are about \$4,000, which is consistent with the median level of medical expenditures for Medicaid enrollees in this age group in the MEPS.[70] Given that a large share of initial Medicaid expenditures come from otherwise uninsured individuals, I also account for the fact that a non-trivial share of the government's expenditures are actually direct transfers to providers for what would have been uncompensated care. To determine how much, I calculate the simulated levels of uncompensated care across the treatment and control groups. As shown in panel A of Table 11, I offset these payments by including the transfer that providers receive as a benefit (where the payments the government makes are included in the \$2,853 number in panel B).

However, this is a wealth effect rather than a response to the prices directly and, therefore, non-distortionary.

[69] By construction, government expenditures are zero for the untreated group prior to turning 65. Also, not all additional spending induced by the expansions is covered by the government. For example, an agent may only be covered by Medicaid for one periods, but their expenditures could increase in subsequent periods as well if they are able to save more than they would have in the control group. However, that the government is covering most of the expenditures suggest that this happens infrequently.

[70] Median expenditures for 62-64 year olds in the MEPS ranges between \$6,000 and \$8,000 when scaled to match aggregate spending measures in the National Health Expenditure Accounts (see. De Nardi, Paschenko, and Porapkkarm 2017). The fact that my estimate is lower is consistent with the fact that expansion individuals are less costly than traditionally covered Medicaid enrollees (see Guth et al. 2021).

Dynamic Fiscal Externalities.—The empirical results presented in Section 5 demonstrate that the Medicaid expansions imposed a positive fiscal externality on the Medicare program, as newly eligible Medicare beneficiaries had relatively lower expenditures throughout their early Medicare tenure. As shown in Figure 7a, the model also replicates this pattern well. However, the positive value presented in the second row of panel B in Table 11 demonstrates that my reduced form result is somewhat misleading. Specifically, when accounting for endogenous survival, the Medicaid expansions *increase* net costs for the Medicare program. This result is driven by two patterns in the model. First, the Medicaid expansions decrease mortality. While this is an important benefit of the expansions, it also produces an additional cost in the form of more, and longer living Medicare enrollees whom the government must insure during retirement. Second, the agents who survive due to the Medicaid expansions (but die in the un-treated group) are also in worse health and have higher healthcare demand. These two results put upward pressure on both total and per-capita Medicare expenditures.

To help illustrate, consider a uninsured individual with chronic conditions who, at age 62, gains Medicaid coverage due to the expansions. The initial provision of coverage induces more consumption and (for the example) better disease management. By the time she turns 65, she is healthier than she otherwise would have been in the absence of the Medicaid expansions, resulting in a relative decrease in healthcare needs. However, her chronic conditions have not been cured and her better health increases her longevity. In this example, it is possible that even though initial Medicare expenditures are lower, the increase in longevity for costly beneficiaries can increase a given individual's cumulative expenditures. In fact, the results from the model are well aligned with this example. Specifically, the model predicts that per-capita Medicare expenditures remain lower in the control group for the first 12 years of Medicare coverage, after which they switch.

Additional Considerations.—There are several other ways in which the Medicaid expansions could affect government expenditures that I have not accounted for, some of which are included in the model and others that are not. For example, the fact that individuals live longer means that they receive SSA benefits and pay taxes for a longer period of time. Though my model can speak to the change in the effect this will have on net costs, I exclude it from my main calculations due to data limitations. That said, when I make the additional assumptions needed to take these into account the results do not change much. The provision of Medicaid may also change reliance on other forms of social insurance, which, in the model, would show up as changes in receipt of the government provided consumption floor. In fact, though not presented, I find that Medicaid reduces reliance on other forms of social insurance (i.e., the consumption floor), which, if anything would lead me to overstate the program's net costs.

One limitation of the model is the exclusion of an endogenous labor supply margin (either intensive or extensive). For example, if the Medicaid expansions increase early retirements this would lead to a decline in tax revenues due to the program. While I acknowledge this limitation, there is little empirical evidence that it is likely a relevant margin in this context.[71] However, this could be more relevant for my counterfactual evaluations in Section 9.

8.3 Final Evaluation

Panel C of Table 11 displays the ratio of total benefits to total net costs of the Medicaid expansions. The value of roughly 1.05 indicates that every dollar that the government spends on the Medicaid expansions for the near-elderly generates roughly $1.05 of benefits. This result is consistent with ranges provided in other studies.[72] Overall, this suggests that expanding Medicaid to the low-income near-elderly is a relatively productive use of public funds even though the healthcare consumption response to Medicaid coverage of this group is considerable.

9 Medicare Reform

In this section, I use the model to assess how allowing low-income individuals to enroll in Medicare starting at age 62 would affect welfare and total government expenditures. To do this, I simulate agents in the control group and a new counterfactual treatment group where instead of Medicaid expanding at age 62, individuals have the option to enroll in Medicare at age 62. The key difference between this counterfactual exercise and the evaluation of the Medicaid expansions is that eligibility for Medicare is not means-tested. However, unlike the automatic enrollment that occurs at age 65, in this counterfactual agents need to actively enroll in Medicare between 62 and 64.[73] The final difference between this counterfactual and the Medicaid expansion is that I assume agents anticipate the changes to Medicare. This is so that I can interpret the results as the longer-run effect of a new Medicare regime, rather than specifically how individuals aged 62 would immediately respond to the program.

Table 12 presents results analogous to those presented for the Medicaid expansions in Table 11. The total CV associated with the policy change is substantially lower than the Medicaid expansions, even though every one in the simulation would be eligible to enroll in coverage through the Medicare reform. This is driven by two differences in the policies. First, the cost-sharing of Medicare is less generous than Medicaid, so the implicit transfer value to consumers is lower. Moreover, anything beyond very marginal improvements in health are extremely costly, so the relatively lean benefits of Medicare (as compared to Medicaid) do not reduce out-of-pocket prices enough to lead to longer-run health improvements for the treated groups. That

[71] For more information, refer to the discussion in Section 4.

[72] See De Nardi, French, and Jones (2016); Finkelstein, Hendren, and Luttmer (2019); and Hendren and Sprung-Keyser (2020).

[73] I assume that there is a utility penalty of enrolling that is equal to one half of the Medicaid enrollment penalty. This is to capture the fact that individuals may still need to fill out extensive paperwork for enrollment, but there is less likely to be any social stigma attached to enrolling. In future work I plan to examine how the ratio of benefits to costs vary by the size of the enrollment cost parameter.

TABLE 12. THE PER-CAPITA BENEFITS AND COSTS OF MEDICARE REFORM

A. Benefits		
CV of Medicare reform		\$2,707
NPV of uncompensated care transfer to providers		\$136
	Total benefit:	\$2,843
B. Costs		
NPV pre-65 Medicare expenditures		\$1,824
NPV of net post-65 Medicare expenditures		\$653
	Total net cost:	\$2,477
C. Change in welfare per net cost		

$$\frac{\sum_i CV_i + \text{NPV of transfers to providers}}{\text{Net costs}} = \mathbf{1.148}$$

Notes.—This table reports the per-capita welfare benefits and net government costs of a counterfactual Medicare reform for the near-elderly population. "CV" refers to welfare as measured by the notion of compensating variation presented in Equation (20).To calculate per-capita estimates I divide the simulated totals by the number of simulated agents and then multiply by 0.4 since I only include the bottom two income quintiles in the simulations.

said, there is still a positive mortality effect of the Medicare reform (with mortality rates decreasing by roughly 4 percent of baseline mortality), which drives the welfare improvements. The second difference is the lack of means-testing. In particular, the Medicare reform does provide a substantial consumption smoothing benefit. However, Medicaid is more valuable in this dimension because it provides much more generous coverage in the states of the world where the value is highest (i.e., when people are very low income, uninsured, and in bad health). This targeting of the Medicaid program generates a much larger insurance value than the Medicare reform even though far fewer individuals enroll in Medicaid.

Turning to program costs, the results in panel B of Table 12 indicate that, though the program covers more people, the less generous coinsurance discourages additional spending such that the pre-65 costs to the government are much lower than those generated by the Medicaid expansions. In addition to the less generous cost-sharing, the fact that the program is anticipated also allows individuals to alter consumption and savings patterns prior to age 62 leading to a somewhat muted response to the insurance provision. The surprising result that net (post-65) Medicare expenditures increased following the Medicaid expansions is also present in panel B of the table, but even more marked. As with the Medicaid expansions, this is due, in part, because more people are surviving to age 65 and because relatively sick and costly individuals are living longer than they otherwise would have, driving up program costs. This second channel is even stronger in the case of the Medicare reform because the improvements in health that people realized before turning 65 are smaller than those from the Medicaid expansions. This is, again, due to less generous cost-sharing and lower corresponding health investments prior to age 65.

Finally, the size of the transfers to private providers for otherwise uncompensated care is higher than for the Medicaid expansion. This is due to the fact that more people are able to move from uninsurance to Medicare coverage because of the lack of an income test.[74]

The change in welfare per dollar of net cost to the government is presented in panel C of Table 12. It indicates that every dollar the government spends on the Medicare reform produces roughly $1.15 of social benefit. This suggest that an opt-in style expansion of the Medicare program to younger low-income individuals could generate considerable welfare gains relative to program costs.

10 Conclusion

In this paper I explore the extent to which the Medicaid expansion under the ACA led to increased healthcare expenditures (i.e., moral hazard), and how the introduction of insurance earlier in life affected the spending habits for those on Medicare. Using data from the Health and Retirement Study, I document a large effect of the Medicaid expansions on insurance coverage and healthcare consumption for low-income near-elderly individuals. Next I provide novel evidence that Medicaid coverage in the periods before retirement led to substantial declines in expenditures and increased health for new Medicare beneficiaries. These findings highlight how earlier provision of public health insurance may impose a positive dynamic fiscal externality on the Medicare system. Finally, I present additional evidence that these effects are persistent, suggesting that there may be longer-term health benefits to providing insurance earlier in life.

Motivated by this evidence, I develop a late-in-life-cycle model that incorporates endogenous health investment, survival, and public insurance enrollment. I am able to identify key model parameters governing the relationship between medical expenditures and health using the reduced form estimates of the causal effects of Medicaid coverage on healthcare consumption and future health outcomes. Model results suggest that expanding Medicaid to the near-elderly population generated substantial welfare gains, where each dollar spent by the government on the program produced roughly $1.05 of social surplus. Contrary to the reduced form results, I find that accounting for endogenous survival causes net Medicare costs to increase following the expansions due to increased longevity. The model also provides insight into why individuals value social insurance programs. Specifically, a decomposition in the value of the Medicaid expansions suggest that roughly 30 percent of the welfare generated by public insurance is due to Medicaid lowering the cost of investments in future health, rather than simply reducing out-

[74]Given that the ACA provides premium subsidies for individuals between Medicaid eligibility and 400 percent of the FPL, it is possible that uncompensated care for middle-income previously uninsured would fall throughout the sample period regardless. However, I abstract from this change in the model given that private insurance coverage is not a choice.

of-pocket risk.

Finally, I use the model to evaluate a reform of the Medicare program that would allow individuals up to 250 percent of the FPL aged 62 to 64 to opt-in to Medicare. Though this reform generates less welfare than the Medicaid expansions, it does so at a considerably lower net cost to the government, yielding nearly $1.15 in social surplus for every additional dollar spent. This suggests that reforming the Medicare program to extend coverage to younger individuals who face high health, earnings, and uninsurance risk could generate substantial value relative to the costs of running the program. The asymmetry in cost-effectiveness across the Medicaid expansion and counterfactual Medicare reform highlights a trade-off associated with the design of public health insurance expansions. Namely, more generous cost-sharing extended to a narrower population can provide larger and more targeted benefits but at much higher costs, whereas broader provision with leaner benefits delivers more benefits per net cost.

References

[1] Abadie, A. Semiparametric Difference-in-Differences Estimators. *The Review of Economic Studies 72*, 1 (01 2005), 1–19.

[2] Abaluck, J., Caceres Bravo, M. M., Hull, P., and Starc, A. Mortality effects and choice across private health insurance plans. Working Paper 27578, National Bureau of Economic Research, July 2020.

[3] Achou, B. Housing in medicaid: Should it really change? *American Economic Journal: Economic Policy 15*, 1 (February 2023), 1–36.

[4] Aizawa, N., and Fang, H. Equilibrium labor market search and health insurance reform. *Journal of Political Economy 128*, 11 (2020), 4258–4336.

[5] Andrews, I., and Stock, J. H. Weak instruments and what to do about them. *NBER Summer Institute, Methods Lectures* (2018).

[6] Aron-Dine, A., Einav, L., and Finkelstein, A. The rand health insurance experiment, three decades later. *Journal of Economic Perspectives 27*, 1 (February 2013), 197–222.

[7] Aron-Dine, A., Einav, L., Finkelstein, A., and Cullen, M. Moral hazard in health insurance: Do dynamic incentives matter? *The Review of Economics and Statistics 97*, 4 (2015), 725–741.

[8] Arrow, K. J. Uncertainty and the welfare economics of medical care. *The American Economic Review 53*, 5 (1963), 941–973.

[9] Baicker, K., Taubman, S. L., Allen, H. L., Bernstein, M., Gruber, J. H., Newhouse, J. P., Schneider, E. C., Wright, B. J., Zaslavsky, A. M., and Finkelstein, A. N. The oregon experiment — effects of medicaid on clinical outcomes. *New England Journal of Medicine 368*, 18 (2013), 1713–1722. PMID: 23635051.

[10] Black, B., Espín-Sánchez, J.-A., French, E., and Litvak, K. The long-term effect of health insurance on near-elderly health and mortality. *American Journal of Health Economics 3*, 3 (2017), 281–311.

[11] Blau, D. M., and Gilleskie, D. B. Health insurance and retirement of married couples. *Journal of Applied Econometrics 21*, 7 (2006), 935–953.

[12] Boudreaux, M. H., Golberstein, E., and McAlpine, D. D. The long-term impacts of medicaid exposure in early childhood: Evidence from the program's origin. *Journal of Health Economics 45* (2016), 161 – 175.

[13] Braun, R. A., Kopecky, K. A., and Koreshkova, T. Old, Sick, Alone, and Poor: A Welfare Analysis of Old-Age Social Insurance Programmes. *The Review of Economic Studies 84*, 2 (03 2016), 580–612.

[14] Buchmueller, T. C., Levy, H. G., and Valletta, R. G. Medicaid expansion and the unemployed. Working Paper 26553, National Bureau of Economic Research, December 2019.

[15] Callison, K., and Sicilian, P. Economic freedom and the affordable care act: Medicaid expansions and labor mobility by race and ethnicity. *Public Finance Review 46*, 2 (2018), 301–324.

[16] Cameron, A. C., and Miller, D. L. A practitioner's guide to cluster-robust inference. *Journal of Human Resources 50*, 2 (2015), 317–372.

[17] Card, D., Dobkin, C., and Maestas, N. The impact of nearly universal insurance coverage on health care utilization: Evidence from medicare. *American Economic Review 98*, 5 (December 2008), 2242–58.

[18] Carey, C. M., Miller, S., and Wherry, L. R. The impact of insurance expansions on the already insured: The affordable care act and medicare. *American Economic Journal: Applied Economics 12*, 4 (October 2020), 288–318.

[19] Centers for Medicare and Medicaid Services. NHE Fact Sheet. https://www.cms.gov/data-research/statistics-trends-and-reports/national-health-expenditure-data/nhe-fact-sheet. Accessed: 2023-09-30.

[20] Centers for Medicare and Medicaid Services. May 2020 medicaid and chip enrollment data highlights, 2020. https://www.medicaid.gov/medicaid/program-information/medicaid-and-chip-enrollment-data/report-highlights/index.html (Accessed April 2020).

[21] Chandra, A., Flack, E., and Obermeyer, Z. The health costs of cost-sharing. Working Paper 28439, National Bureau of Economic Research, February 2021.

[22] Chetty, R., Stepner, M., Abraham, S., Lin, S., Scuderi, B., Turner, N., Bergeron, A., and Cutler, D. The Association Between Income and Life Expectancy in the United States, 2001-2014. *JAMA 315*, 16 (04 2016), 1750–1766.

[23] Clingman, M., Burkhalter, K., and Chaplain, C. Replacement rates for hypothetical retired workers. *Social Security Administration, Actuarial Note*.

[24] Congressional Budget Office. Updated estimates of the insurance coverage provisions of the affordable care act, april 2014. http://www.cbo.gov/sites/default/files/cbofiles/attachments/45231-ACA_Estimates.pdf (Accessed April 2020).

[25] Courtemanche, C., Marton, J., Ukert, B., Yelowitz, A., and Zapata, D. Early impacts of the affordable care act on health insurance coverage in medicaid expansion and non-expansion states. *Journal of Policy Analysis and Management 36*, 1 (2017), 178–210.

[26] Cronin, C. J. Insurance-Induced Moral Hazard: A Dynamic Model Of Within-Year Medical Care Decision Making Under Uncertainty. *International Economic Review 60*, 1 (February 2019), 187–218.

[27] Crossley, T. F., Levell, P., and Poupakis, S. Regression with an imputed dependent variable. *Journal of Applied Econometrics 37*, 7 (2022), 1277–1294.

[28] Cutler, D. M., and Gruber, J. Does Public Insurance Crowd out Private Insurance? *The Quarterly Journal of Economics 111*, 2 (05 1996), 391–430.

[29] De Nardi, M., French, E., and Jones, J. B. Medicaid insurance in old age. *American Economic Review 106*, 11 (November 2016), 3480–3520.

[30] De Nardi, M., Pashchenko, S., and Porapakkarm, P. The lifetime costs of bad health. Working Paper 23963, National Bureau of Economic Research, October 2017.

[31] De Nardi, M., French, E., and Jones, J. Why do the elderly save? the role of medical expenses. *Journal of Political Economy 118*, 1 (2010), 39–75.

[32] Diaz-Campo, C. S. Dynamic moral hazard in nonlinear health insurance contracts. *Unpublished Manuscript* (2022).

[33] Dranove, D., Gartwaite, C., and Ody, C. The impact of the aca's medicaid expansion on hospitals' uncompensated care burden and the potential effects of repeal. *Issue brief (Commonwealth Fund) 12* (05 2017), 1–9.

[34] Einav, L., and Finkelstein, A. Moral hazard in health insurance: What we know and how we know it. *Journal of the European Economic Association 16* (05 2018).

[35] Einav, L., Finkelstein, A., Ryan, S. P., Schrimpf, P., and Cullen, M. R. Selection on moral hazard in health insurance. *American Economic Review 103*, 1 (February 2013), 178–219.

[36] Einav, L., Finkelstein, A., and Schrimpf, P. The Response of Drug Expenditure to Nonlinear Contract Design: Evidence from Medicare Part D . *The Quarterly Journal of Economics 130*, 2 (02 2015), 841–899.

[37] Fang, H., and Gavazza, A. Dynamic inefficiencies in an employment-based health insurance system: Theory and evidence. *American Economic Review 101*, 7 (December 2011), 3047–77.

[38] Fang, H., Keane, M., and Silverman, D. Sources of advantageous selection: Evidence from the medigap insurance market. *Journal of Political Economy 116*, 2 (2008), 303–350.

[39] Finkelstein, A., Hendren, N., and Luttmer, E. F. P. The value of medicaid: Interpreting results from the oregon health insurance experiment. *Journal of Political Economy 127*, 6 (2019), 2836–2874.

[40] Finkelstein, A., Luttmer, E. F. P., and Notowidigdo, M. J. What good is wealth without health? the effect of health on the marginal utility of consumtpion. *Journal of the European Economic Association 11*, s1 (2013), 221–258.

[41] Finkelstein, A., Taubman, S., Wright, B., Bernstein, M., Gruber, J., Newhouse, J. P., Allen, H., Baicker, K., and Group, O. H. S. The oregon health insurance experiment: Evidence from the first year. *The Quarterly Journal of Economics 127*, 3 (2012), 1057–1106.

[42] Fox, A., Feng, W., and Reynolds, M. The effect of administrative burden on state safety-net participation: Evidence from food assistance, cash assistance, and medicaid. *Public Administration Review 83*, 2 (2023), 367–384.

[43] Frean, M., Gruber, J., and Sommers, B. D. Premium subsidies, the mandate, and medicaid expansion: Coverage effects of the affordable care act. *Journal of Health Economics 53* (2017), 72 – 86.

[44] French, E., and Jones, J. B. The Effects of Health Insurance and Self-Insurance on Retirement Behavior. *Econometrica 79*, 3 (May 2011), 693–732.

[45] French, E., and Jones, J. B. Health, health insurance, and retirement: A survey. *Annual Review of Economics 9* (2017), pp. 383–409.

[46] Geruso, M., Layton, T. J., and Wallace, J. What difference does a health plan make? evidence from random plan assignment in medicaid. *American Economic Journal: Applied Economics 15*, 3 (July 2023), 341–79.

[47] Gilleskie, D. B. A dynamic stochastic model of medical care use and work absence. *Econometrica 66*, 1 (1998), 1–45.

[48] Goldin, J., Lurie, I. Z., and McCubbin, J. Health Insurance and Mortality: Experimental Evidence from Taxpayer Outreach. *The Quarterly Journal of Economics 136*, 1 (09 2020), 1–49.

[49] Goodman, L. The effect of the affordable care act medicaid expansion on migration. *Journal of Policy Analysis and Management 36*, 1 (2017), 211–238.

[50] Goodman-Bacon, A. Difference-in-differences with variation in treatment timing. *Journal of Econometrics 225*, 2 (2021), 254–277. Themed Issue: Treatment Effect 1.

[51] Goodman-Bacon, A. The long-run effects of childhood insurance coverage: Medicaid implementation, adult health, and labor market outcomes. *American Economic Review 111*, 8 (August 2021), 2550–93.

[52] Gooptu, A., Moriya, A. S., Simon, K. I., and Sommers, B. D. Medicaid expansion did not result in significant employment changes or job reductions in 2014. *Health Affairs 35*, 1 (2016), 111–118. PMID: 26733708.

[53] Gourieroux, C., Monfort, A., and Renault, E. Indirect inference. *Journal of Applied Econometrics 8*, S1 (1993), S85–S118.

[54] Grossman, M. On the concept of health capital and the demand for health. *Journal of Political Economy 80*, 2 (1972), 223–55.

[55] Gruber, J. *Public Finance and Public Policy*, seventh ed. Worth Publishers, 2022.

[56] Guth, M., Corallo, B., Rudowitz, R., and Garfield, R. Medicaid expansion enrollment and spending leading up to the covid-19 pandemic. *Kaiser Family Foundation*. `https://www.kff.org/medicaid/issue-brief/medicaid-expansion-enrollment-and-spending-leading-up-to-the-covid-19-pandemic/` (Accessed: 2022-09-30).

[57] Hadley, J., and Holahan, J. How much medical care do the uninsured use, and who pays for it? *Health Affairs 22*, Suppl1 (2003), W3–66–W3–81.

[58] Hall, R. E., and Jones, C. I. The Value of Life and the Rise in Health Spending. *The Quarterly Journal of Economics 122*, 1 (02 2007), 39–72.

[59] Hendren, N., and Sprung-Keyser, B. A unified welfare analysis of government policies. *Quarterly Journal of Economics 135*, 3 (2020), 1209–1318. View the estimates online at www.policyinsights.orgWatch the Econimate Video.

[60] Hu, L., Kaestner, R., Mazumder, B., Miller, S., and Wong, A. The effect of the affordable care act medicaid expansions on financial wellbeing. *Journal of Public Economics 163* (2018), 99 – 112.

[61] Hudson, Sally; Hull, P., and Liebersohn, J. Interpreting instrumented difference-in-differences. `https://www.mit.edu/~liebers/DDIV.pdf` (Accessed: 2021-08-30).

[62] Imbens, G. W., and Angrist, J. D. Identification and estimation of local average treatment effects. *Econometrica 62*, 2 (1994), 467–475.

[63] Kaestner, R., Garrett, B., Gangopadhyaya, A., and Fleming, C. Effects of aca medicaid expansions on health insurance coverage and labor supply. Working Paper 21836, National Bureau of Economic Research, December 2015.

[64] Kaiser Family Foundation. Status of state medicaid expansion decisions: Interactive map. https://www.kff.org/medicaid/issue-brief/status-of-state-medicaid-expansion-decisions-interactive-map/ (Accessed: 2023-07-31).

[65] Khatana SAM, B. A., and et al., N. A. Association of medicaid expansion with cardiovascular mortality. *JAMA Cardiology 4* (2019), 671–679.

[66] Khwaja, A. Estimating willingness to pay for medicare using a dynamic life-cycle model of demand for health insurance. *Journal of Econometrics 156*, 1 (2010), 130–147. Structural Models of Optimization Behavior in Labor, Aging, and Health.

[67] Laitner, J., Silverman, D., and Stolyarov, D. The role of annuitized wealth in post-retirement behavior. *American Economic Journal: Macroeconomics 10*, 3 (July 2018), 71–117.

[68] Leung, P., and Mas, A. Employment effects of the aca medicaid expansions. Working Paper 22540, National Bureau of Economic Research, August 2016.

[69] Li, A. Commitment, competition, and preventive care provision. *Unpublished Manuscript* (2023).

[70] Lin, H., and Sacks, D. W. Intertemporal substitution in health care demand: Evidence from the rand health insurance experiment. *Journal of Public Economics 175* (2019), 29–43.

[71] Lockwood, L. M. Incidental bequests and the choice to self-insure late-life risks. *American Economic Review 108*, 9 (September 2018), 2513–50.

[72] Low, H., and Pistaferri, L. Disability insurance and the dynamics of the incentive insurance trade-off. *American Economic Review 105*, 10 (October 2015), 2986–3029.

[73] Luhby, T. Democrats want ot expand medicare benefits through spending bill. *CNN*. https://www.cnn.com/2021/07/14/politics/medicare-dental-vision-benefits-spending-bill-democrats/index.html (Accessed 2022-10-20).

[74] Mazurenko, O., Balio, C. P., Agarwal, R., Carroll, A. E., and Menachemi, N. The effects of medicaid expansion under the aca: A systematic review. *Health Affairs 37*, 6 (2018), 944–950. PMID: 29863941.

[75] McInerney, M., Winecoff, R., Ayyagari, P., Simon, K., and Bundorf, M. K. Aca medicaid expansion associated with increased medicaid participation and improved health among near-elderly: Evidence from the health and retirement study. *INQUIRY: The Journal of Health Care Organization, Provision, and Financing 57* (2020), 0046958020935229. PMID: 32720837.

[76] McWilliams, J. M., Meara, E., Zaslavsky, A. M., and Ayanian, J. Z. Use of health services by previously uninsured medicare beneficiaries. *New England Journal of Medicine 357*, 2 (2007), 143–153. PMID: 17625126.

[77] Medicaid and CHIP Payment and Access Commission. Medicaid enrollment changes following the aca.

[78] Miller, S., Johnson, N., and Wherry, L. R. Medicaid and Mortality: New Evidence From Linked Survey and Administrative Data. *The Quarterly Journal of Economics 136*, 3 (01 2021), 1783–1829.

[79] Miller, S., and Wherry, L. R. The long-term effects of early life medicaid coverage. *Journal of Human Resources 54*, 3 (2019), 785–824.

[80] Murphy, K., and Topel, R. The value of health and longevity. *Journal of Political Economy 114*, 5 (2006), 871–904.

[81] Newhouse, J. P., and Group, T. I. E. Free for all? lessons from the rand health insurance experiment. *Harvard Univeristy Press* (1993).

[82] Norris, H. C., Richardson, H. M., Benoit, M.-A. C., Shrosbree, B., Smith, J. E., and Fendrick, A. M. Utilization impact of cost-sharing elimination for preventive care services: A rapid review. *Medical Care Research and Review 79*, 2 (2022), 175–197. PMID: 34157906.

[83] Nyman, J. A., and Maude-Griffin, R. The welfare economics of moral hazard. *International Journal of Health Care Finance and Economics 1* (2001), 23–42.

[84] Olden, A., and Møen, J. The triple difference estimator. *The Econometrics Journal 25*, 3 (03 2022), 531–553.

[85] Ozkan, S. Preventive vs. curative medicine: A macroeconomic analysis of health care over the life cycle. *Unpublished Manuscript*, University of Toronto (2017).

[86] Palumbo, M. G. Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle. *Review of Economic Studies 66*, 2 (1999), 395–421.

[87] Pauly, M. V. The economics of moral hazard: Comment. *The American Economic Review 58*, 3 (1968), 531–537.

[88] Peng, L., Guo, X., and Meyerhoefer, C. D. The effects of medicaid expansion on labor market outcomes: Evidence from border counties. *Health Economics 29*, 3 (2020), 245–260.

[89] Poterba, J., Venti, S., and Wise, D. A. Health, education, and the postretirement evolution of household assets. *Journal of Human Capital 7*, 4 (2013), 297–339.

[90] Poterba, J. M., Venti, S. F., and Wise, D. A. The asset cost of poor health. *The Journal of the Economics of Ageing 9* (2017), 172–184.

[91] Ridder, G., and Moffitt, R. The econometrics of data combination. In *Handbook of Econometrics*, J. Heckman and E. Leamer, Eds., 1 ed., vol. 6B. Elsevier, 2007, ch. 75.

[92] Russo, N. Health-dependent preferences, consumption, and insurance. *Manuscript, Goethe University Frankfurt* (2023).

[93] Saha, N. A unified framework for valuing human health and mortality. *Unpublished Manuscript* (2021).

[94] Shupe, C. Public health insurance and medical spending: The incidence of the aca medicaid expansion. *Journal of Policy Analysis and Management 42*, 1 (2023), 137–165.

[95] Simon, K., Soni, A., and Cawley, J. 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 (2017), 390–417.

[96] Tipirneni, R., Levy, H. G., Langa, K. M., McCammon, R. J., Zivin, K., Luster, J., Karmakar, M., and Ayanian, J. Z. Changes in Health Care Access and Utilization for Low-SES Adults Age 51-64 after Medicaid Expansion. *The Journals of Gerontology: Series B* (08 2020).

[97] Vreugdenhil, N. Booms, busts, and mismatch in capital markets: Evidence from the offshore oil and gas industry. *Journal of Political Economy, Forthcoming*.

[98] Weir, D. R. Validating mortality ascertainment in the health and retirement study. Tech. rep., Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, Michigan, November 2016.

[99] Yang, Z., Gilleskie, D. B., and Norton, E. C. Health insurance, medical care, and health outcomes: A model of elderly health dynamics. *The Journal of Human Resources 44*, 1 (2009), 47–114.

A Additional Data Details

Further Details on the HRS.—Beginning in 1992, wave 1 of the HRS surveyed a nationally representative sample of people born between 1931 and 1941 and their spouses. This initial panel also included oversamples of Blacks, Hispanics, and Florida residents. The HRS has reinterviewed the wave 1 cohort every other year since 1992. In 1998, the HRS expanded the scope of their study to provide insight into the entire U.S. population over the age of 50. In pursuit of this goal, the HRS merged in a separate study of individuals born before 1924, dubbed the AHEAD cohort. Respondents in AHEAD were first interviewed in 1993 and then in 1995, 1998, and subsequently every two years. The HRS also began surveying two additional cohorts: the "Children of the Depression" (CODA) and the "War Baby" (WB) cohorts. CODA contains individuals born between 1924 and 1930, and the WB cohort includes those born between 1942 and 1947. Each of these cohorts were first surveyed in 1998 and every two years after. Since 1998, three additional cohorts have been added to the HRS as Baby Boomers have approached and entered retirement. HRS added the "Early Baby Boomer" (EBB) cohort in 2004, the "Mid Baby Boomer" (MBB) cohort in 2010, and the "Late Baby Boomer" (LBB) in 2016. Individuals included in the EBB were born between 1948 and 1953, the MBB cohort respondents were born between 1954 and 1959, and the LBB were born 1960 to 1965. Since first included, the EBB, MBB, and LBB cohorts have been reinterviewed every two years.

Further Details on the MEPS.—MEPS is made up of a set of large annual rotating panel surveys of individuals and families, medical providers, and employers. These surveys provide nationally representative estimates of healthcare expenditures, utilization, and insurance coverage for the noninstitutionalized civilian population in the United States. My main purpose of using MEPS is to impute total medical expenditures in the HRS sample, and there are a number of reasons why the Household Component (HC) of MEPS is well suited to this end. In terms of substance, the HC surveys contain a substantial amount of information on demographics, health conditions, health status, healthcare utilization, insurance coverage, income, and employment that mirror analogous measures in the HRS.[75] The HC is also a good fit for imputation due to the survey design. Specifically, the HC surveys households in two consecutive years which is critical since many HRS survey variables cover the entire survey period (or two years). Since I can observe the same household in the HC for the two consecutive years that correspond to the HRS waves, I can generate measures that better match those in the HRS. Thus, to make the HRS and MEPS directly comparable, I collapse the MEPS data at the individual level and keep individuals that were surveyed over a two year period corresponding to a wave in the HRS.[76]

[75] As with the HRS data, I deflate all monetary values in MEPS using the GDP implicit price deflator, with 2012 as the base year.

[76] For example, an individual surveyed in 2015 and 2016 (which overlaps with HRS wave 13) is kept, but an individual surveyed in 2014 and 2015 is dropped.

Based on MEPS survey methodology, the representative nature of the data do not change when dropping staggered observations in this way.

Descriptive Statistics in the HRS and MEPS.—Table 13 presents some descriptive statistics for the elderly, which I define as above age 64 and on Medicare, and non-elderly samples. Overall, the samples are similar, with the exception that the HRS sample is somewhat older and generally less healthy based on the ADL and IADL indicators. The HRS sample also spends more out-of-pocket on healthcare than those in MEPS. Somewhat surprisingly, the MEPS has much higher share of Medicaid enrollees compared to the HRS even though incomes look quite similar.[77] Within the HRS the elderly are much less likely to be working, are poorer, and are generally less healthy. Even with those differences, their out-of-pocket spending is only slightly higher than the non-elderly group. Overall, Table 1 demonstrates that the samples in the HRS and MEPS are generally similar on a range of observables, suggesting that the MEPS is a reasonable option for imputing mean expenditures in the HRS.

A.1 Description of Main Analysis Samples

The Pre-Medicare Sample.—Table 14 presents descriptive statistics for the pre-Medicare sample separately for expansions and non-expansion states.[78] Individuals across expansion and non-expansion states are broadly similar in terms of demographics, with the exceptions that individuals in expansion states are whiter, more likely to have graduated high school, have higher incomes and are wealthier. As a consequence of individuals in expansion states earning higher incomes, there is also a slightly higher share of individuals in non-expansion states that fall below the 138 percent threshold of the FPL. Unsurprisingly, Medicaid enrollment is slightly higher in expansion states which is, in part, due to higher enrollment following the expansions. Private coverage also differs across states, with non-expansion states having a much lower rate of private coverage. Correspondingly, non-expansion states also have a much higher share of uninsured individuals. However, rates of SSDI do not seem to systematically vary across expansion and non-expansion states.

A number of measures of health also vary across expansion and non-expansion states. For example, based on self-reported health, individuals in expansion states appear to be somewhat healthier than their non-expansion peers. However, across many diagnostic measures (e.g., diabetes, stroke, heart disease) there is little difference across the states. These differences suggest that controlling for health status is likely important for my analysis. In addition to including subjective and diagnostic measures of health, I also construct a general health index in the spirit of Poterba, Venti, and Wise (2013) and Poterba, Venti, and Wise (2017). The index is given by the first principal component a large list of diagnostic and subjective measures of health and

[77] It is possible that this is due to under-reporting of Medicaid status in the HRS.

[78] Note that an expansion state is one that expanded at any point throughout the sample period.

Table 13. Descriptive statistics of selected variables in the MEPS and HRS

	MEPS			HRS		
	All	Non-elderly	Elderly	All	Non-elderly	Elderly
Medicaid	0.26	0.24	0.27	0.16	0.14	0.17
	(0.44)	(0.42)	(0.44)	(0.37)	(0.34)	(0.38)
Total household income	25,138	29,043	22,385	25,889	31,117	23,705
	(17,188)	(19,217)	(14,997)	(18,454)	(22,837)	(15,756)
Female	0.61	0.60	0.62	0.64	0.64	0.64
	(0.49)	(0.49)	(0.49)	(0.48)	(0.48)	(0.48)
Age	66.04	56.37	72.83	69.76	57.79	74.77
	(10.80)	(4.31)	(8.59)	(10.90)	(3.94)	(8.80)
Black	0.22	0.24	0.21	0.22	0.28	0.20
	(0.42)	(0.43)	(0.41)	(0.42)	(0.45)	(0.40)
Hispanic	0.21	0.28	0.17	0.16	0.25	0.12
	(0.41)	(0.45)	(0.37)	(0.37)	(0.43)	(0.33)
Married	0.46	0.51	0.41	0.50	0.60	0.46
	(0.50)	(0.50)	(0.49)	(0.50)	(0.49)	(0.50)
Widowed	0.23	0.08	0.34	0.29	0.10	0.37
	(0.42)	(0.27)	(0.47)	(0.45)	(0.30)	(0.48)
Divorced	0.18	0.23	0.15	0.13	0.18	0.10
	(0.39)	(0.42)	(0.36)	(0.33)	(0.38)	(0.31)
Less than high school	0.34	0.29	0.38	0.42	0.36	0.44
	(0.47)	(0.46)	(0.48)	(0.49)	(0.48)	(0.50)
High school	0.28	0.29	0.28	0.35	0.35	0.35
	(0.45)	(0.45)	(0.45)	(0.48)	(0.48)	(0.48)
Some college	0.14	0.15	0.13	0.16	0.20	0.14
	(0.34)	(0.36)	(0.33)	(0.37)	(0.40)	(0.35)
College	0.09	0.10	0.08	0.07	0.08	0.06
	(0.28)	(0.30)	(0.27)	(0.25)	(0.27)	(0.24)
Working	0.28	0.56	0.08	0.19	0.52	0.05
	(0.45)	(0.50)	(0.28)	(0.39)	(0.50)	(0.22)
ADL help or difficulty	0.06	0.02	0.08	0.11	0.06	0.14
	(0.23)	(0.15)	(0.27)	(0.32)	(0.24)	(0.34)
IADL help or difficulty	0.11	0.05	0.16	0.25	0.21	0.27
	(0.31)	(0.21)	(0.36)	(0.44)	(0.40)	(0.45)
Smoker	0.17	0.22	0.13	0.17	0.28	0.13
	(0.37)	(0.42)	(0.34)	(0.38)	(0.45)	(0.34)
Total out-of-pocket spending	1,236	1,021	1,387	2,525	2,123	2,694
	(2,167)	(1,976)	(2,281)	(5,178)	(4,162)	(5,540)
Total medical expenditures	8,523	5,782	10,450	.	.	.
	(11,840)	(95,18)	(12,880)	.	.	.

Notes.—Standard deviations reported in parentheses. Number of observations varies by variable and sample. All monetary values are deflated with 2012 as the base year.

Table 14. Summary statistics for the pre-Medicare sample, by state expansion status

	Non-expansion states		Expansion states		Difference
	Mean	Std. dev.	Mean	Std. dev.	Non-exp. − Exp.
Demographics					
Age	62.24	(1.476)	62.21	(1.480)	0.026
Female	0.58	(0.494)	0.57	(0.496)	0.010
Black	0.17	(0.377)	0.12	(0.326)	0.051***
Hispanic	0.10	(0.296)	0.08	(0.276)	0.014**
Married	0.75	(0.431)	0.76	(0.425)	-0.010
Less than high school	0.20	(0.402)	0.15	(0.361)	0.049***
High school	0.33	(0.470)	0.34	(0.475)	-0.015*
College	0.25	(0.433)	0.26	(0.440)	-0.013*
Household income	78360	(92136)	93655	(156572)	-15295***
Household assets	337546	(856592)	425455	(1483757)	-87910***
$\leq 138\% FPL$	0.13	(0.334)	0.11	(0.316)	0.016**
Insurance					
Medicaid	0.02	(0.139)	0.03	(0.171)	-0.010***
Private coverage	0.78	(0.417)	0.84	(0.363)	-0.068***
Uninsured	0.16	(0.366)	0.10	(0.303)	-0.058***
SSDI	0.00	(0.067)	0.00	(0.069)	-0.000
Health					
Health index	0.70	(0.047)	0.70	(0.046)	-0.001
Self-reported health	3.45	(0.997)	3.50	(0.983)	-0.053***
BMI	28.23	(5.455)	28.04	(5.244)	0.186*
Smoker	0.17	(0.375)	0.15	(0.356)	0.020***
Ever smoke	0.57	(0.495)	0.57	(0.494)	-0.006
High blood-pressure	0.47	(0.499)	0.44	(0.497)	0.030***
Diabetes	0.14	(0.352)	0.14	(0.342)	0.010
Cancer	0.08	(0.268)	0.09	(0.283)	-0.010*
Heart ailment	0.13	(0.332)	0.12	(0.328)	0.003
Stroke	0.03	(0.170)	0.03	(0.164)	0.002
Lung ailment	0.04	(0.206)	0.05	(0.217)	-0.005
Healthcare utilization					
Any hospital stay	0.16	(0.363)	0.16	(0.362)	0.001
Hospital stays	0.21	(0.577)	0.21	(0.564)	0.003
Nights in hospital	0.81	(3.395)	0.85	(3.888)	-0.037
Any doctor visit	0.91	(0.291)	0.93	(0.262)	-0.019***
Num. doctor visits	6.92	(8.581)	7.55	(10.170)	-0.635***
Any outpatient care	0.17	(0.379)	0.17	(0.377)	0.002
Take Rxs regularly	0.73	(0.445)	0.72	(0.450)	0.010
Any dental visit	0.66	(0.475)	0.73	(0.445)	-0.071***
Any specialty facility	0.05	(0.224)	0.07	(0.252)	-0.015***
Imputed log expenditures, (ins. incld.)	8.10	(1.120)	8.15	(1.052)	-0.052**
Imputed log expenditures, (ins. excld.)	8.09	(1.091)	8.12	(1.045)	-0.032*
Total out-of-pocket spending	3023.45	(5403.035)	2527.63	(4465.841)	495.818***
Observations	7780		12270		

Note: * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Table 15. Summary statistics for the first Medicare wave sample

	Non-expansion states		Expansion states		Difference
	Mean	Std. dev.	Mean	Std. dev.	Non-exp. − Exp.
Demographics					
Age	65.66	(0.731)	65.65	(0.725)	0.010
Female	0.57	(0.494)	0.57	(0.495)	0.005
Black	0.17	(0.380)	0.12	(0.324)	0.055***
Hispanic	0.10	(0.298)	0.09	(0.279)	0.013*
Married	0.72	(0.450)	0.73	(0.443)	-0.015
Less than high school	0.22	(0.413)	0.16	(0.371)	0.053***
High school	0.32	(0.468)	0.35	(0.477)	-0.028**
College	0.25	(0.431)	0.25	(0.435)	-0.007
Household income	73808	(140389)	81965	(155607)	-8157*
Household assets	392839	(1371778)	459842	(1231275)	-67003*
Health					
Health index	0.66	(0.051)	0.66	(0.049)	-0.002
Self-reported health	3.35	(1.012)	3.39	(1.006)	0.049*
BMI	28.29	(5.490)	28.19	(5.455)	0.095
High blood-pressure	0.56	(0.497)	0.52	(0.500)	0.043***
Diabetes	0.20	(0.397)	0.18	(0.388)	0.011
Cancer	0.11	(0.318)	0.12	(0.326)	-0.007
Heart ailment	0.18	(0.384)	0.17	(0.372)	0.014
Stroke	0.05	(0.212)	0.04	(0.206)	0.003
Lung ailment	0.07	(0.253)	0.07	(0.255)	-0.001
Healthcare utilization					
Any hospital stay	0.20	(0.402)	0.19	(0.395)	0.009
Hospital stays	0.31	(0.755)	0.30	(0.774)	0.012
Nights in hospital	1.45	(5.820)	1.36	(6.126)	0.089
Any doctor visit	0.92	(0.277)	0.94	(0.245)	-0.020***
Num. doctor visits	8.13	(10.899)	8.54	(11.929)	-0.406
Any outpatient care	0.19	(0.394)	0.20	(0.403)	-0.011
Take Rxs regularly	0.81	(0.396)	0.78	(0.412)	0.021*
Any dental visits	0.63	(0.484)	0.70	(0.460)	-0.070***
Any specialty facility	0.08	(0.273)	0.09	(0.287)	-0.009
Imputed log expenditures, (ins. incld.)	8.22	(1.026)	8.20	(1.001)	0.022
Imputed log expenditures, (ins. excld.)	8.20	(1.016)	8.17	(0.996)	0.026
Total out of pocket spending	3070.24	(5441.578)	2694.45	(4742.681)	375.788***
Observations	3534		5409		8943

Note: * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

well-being. Though there is a slight difference in the health index across expansion and non-expansion states, it is not significant. The lack of a difference in the health index but notable difference in self-reported health suggests that incorporating diagnostic measures are likely important in understanding health. In the bulk of my analysis I limit attention to either self-reported health or the health index to understand the evolution of health over the life-cycle and in response to insurance coverage. Though no measure of health is perfect, I am confident with these measures because they are strongly associated with future mortality as shown in Table 21 of Appendix E.

Some measures of healthcare utilization also differ significantly across states. For example, individuals in expansion states are more likely to go to the doctor (along both extensive and intensive margins), the dentist, and specialty facilities; but across other measures the states are quite similar (e.g., hospital visits, outpatient care, and prescription drugs). Finally, both measures of imputed expenditures are higher in expansion states, but the relative difference is higher for the measure that includes insurance in the imputation. At the same time, out-of-pocket spending is actually lower in expansion states. For the imputation that includes insurance coverage, this could be, in part, driven by lower uninsurance rates in expansion states. However, I interpret the fact that this asymmetry exists for the imputation measure that excludes insurance coverage as indicating that including other measures of utilization (which are, for the most part, weakly higher in expansion states) is critical in determining true expenditures.

The Medicare Sample.—A similar set of summary statistics for the wave one Medicare sample are presented in Table 15. As expected, since these are the same people in both samples, the trends in demographics, health measures, and healthcare utilization across expansion and non-expansion states are similar to those in the pre-Medicare sample. However, the difference in imputed log total medical expenditures now disappears even though the difference in out-of-pocket spending remains (but shrinks somewhat). Overall, the summary statistics in Table 15 demonstrate that, across many observables, individuals in expansion and non-expansion states on their first wave of Medicare are fairly similar.

B Imputation Details

B.1 Imputation of Medical Expenditures in the HRS

In order to conduct the empirical analysis in this paper, I would ideally have a data set that augments the HRS sample with information on each individual's total medical expenditures for each sample period. Though the HRS can be linked to administrative records from the Center for Medicare and Medicaid Services, those linked records only provide expenditures paid for by either Medicare or Medicaid. Since I am also interested in the expenditures of individuals not covered by either Medicare or Medicaid, those data will not serve my purposes.[79] Alternatively, MEPS provides detailed information on individual expenditures for all individuals regardless of their insurance coverage. Unfortunately, MEPS does not have a long enough panel component to allow me to examine how medical expenditures in one period affect outcomes long into the future. As a result, I combine the HRS and the HC from MEPS for my empirical analysis, where my imputation strategy is informed by Ridder and Moffitt (2007) and Fang, Keane, and Silverman (2008).

Imputation Strategy.—Denote the MEPS data as

$$\{E_i, \mathbf{I}_i, \mathbf{H}_i, \mathbf{F}_i, \mathbf{D}_i\}_{i \in \mathcal{I}_{MEPS}}$$

and the data in HRS as

$$\{\mathbf{I}_j, \mathbf{H}_j, \mathbf{F}_j, \mathbf{D}_j\}_{j \in \mathcal{I}_{HRS}}$$

where $\mathcal{I}_{MEPS}$ and $\mathcal{I}_{HRS}$ denote the MEPS and HRS samples, respectively. In this case, all of the key variables from the HRS–insurance coverage ($\mathbf{I}$), health measures and healthcare utilization $\mathbf{H}$, financial variables ($\mathbf{F}$), and demographics ($\mathbf{D}$)–are found in MEPS, but total medical expenditures (E) are only in MEPS.[80] My imputation strategy is to first use the variables that are consistent across both HRS and MEPS to estimate a prediction equation for total expenditures, which I then use to impute expected expenditures, $\hat{E}_j$, for each person in the HRS sample. Specifically, I obtain a prediction equation of the following form from MEPS:

$$\hat{E}_i = \hat{\theta}_0 + \hat{\theta}_1 \mathbf{I}_i + \hat{\theta}_2 \mathbf{H}_i + \hat{\theta}_3 \mathbf{F}_i + \hat{\theta}_4 \mathbf{D}_i. \tag{21}$$

Now I can use this prediction equation to impute expected expenditures for the remainder of the HRS sample, where for each $j \in \mathcal{I}_{HRS}$:

$$\hat{E}_j = \hat{\theta}_0 + \hat{\theta}_1 \mathbf{I}_j + \hat{\theta}_2 \mathbf{H}_j + \hat{\theta}_3 \mathbf{F}_j + \hat{\theta}_4 \mathbf{D}_j. \tag{22}$$

In practice, I estimate each of these equations separately for elderly and non-elderly individuals. The motivation for doing so is that, ex-ante, one might expect that spending behavior changes once individuals go on Medicare and around the time that they stop working (e.g., Card, Dobkin, and Maestas 2008). Also, as individuals get older, the data generating process for expenditures may change substantially as well. Therefore, I find it appropriate to split the

[79] Additionally, HRS will only allow researcher to link the claims data with the geographic location data, which is critical for my empirical strategy, under very unique circumstances. This hurdle makes constructing the ideal data set nearly impossible.

[80] The vector of health related variables, $\mathbf{H}$, contains a number of measures for health as well as specific utilization variables like number of nights spent in the hospital, which I will not include as controls in my later analysis. Having these additional measures of utilization is critical for identification.

sample for imputation. Additionally, I include a subset of the HRS sample who do have observed medical expenditures in Equation (21).[81] I include this random subset to hopefully provide additional information on spending patterns that could be important in the HRS data that is, for some reason, not effectively captured by the MEPS.

Table 16 presents the regression results from the imputation process. The table demonstrates that the regression model does a pretty good job of explaining the variation in total medical expenses across both samples (based on the adjusted-R^2 values of roughly 0.5 and 0.55 across the different models). Moreover, the results across the models suggest that estimating these separately for the Pre-Medicare and Medicare groups is likely important since the relative contribution of the covariates seems to differ somewhat substantially across the samples. For example, the coefficients on educational attainment and insurance coverage are much larger for the Pre-Medicare sample. Overall, the results in the table suggest that utilization, health, and insurance coverage have the strongest association with expenditures. This is expected and encouraging, since I argue that these utilization measures drive the variation in my imputed expenditure measure which is central to the identification of my parameters of interest.

Following imputation, I trim the data by dropping all observations with imputed expenditures above the 99th percentile of imputations in the sample.

Imputation Results.—Table 17 presents a summary of the imputations by different groups, and provides expenditures in MEPS and total out-of-pocket spending and the working shares in both sample for comparison. Across MEPS and HRS, the HRS imputations are generally inflated relative to the average expenditures observed in MEPS. This discrepancy exists in the out-of-pocket spending in both samples as well, suggesting that the imputations, while higher, most likely reflect actual differences in spending across the two samples. Table 13 also shows that individuals in the HRS are more likely to report having difficulties with both ADLs and IADLs which suggests the individuals in the HRS are generally less healthy than those in MEPS. Thus, these differences across the samples in terms of health and utilization are reflected in the imputations.

Focusing on the non-elderly, we see that individuals on Medicaid spend much more on healthcare than those not in the program. This is most likely driven by a number of factors including the moral hazard response that I am interested in. However, this is not evidence of moral hazard directly since the sample of individuals that select into Medicaid are most likely the sickest low-income individuals and, thus, the most in need of healthcare. As shown in the table, the working share for those on Medicaid is slightly over a third of the non-elderly non-Medicaid sample in the HRS. This could reflect that large negative health shocks lead to worse employment outcomes and higher rates of Medicaid eligibility as well as increased expenditures. Motivated by the potential for selection, I use two versions of my imputed measure of healthcare

[81] I use the excluded observations in the HRS as a "hold-out" sample with which I can assess the accuracy and validity of my predictions.

expenditures: the first includes insurance information and the second excludes any insurance information. The rational for exploring multiple imputation measures is that if Medicaid is traditionally adversely selected (i.e., enrollees are less healthy and consume more care), then the imputed measures that account for Medicaid coverage could *overstate* true expenditures for individuals enrolling in Medicaid due to the expansions. Alternatively, not accounting for the effect that Medicaid coverage has on out-of-pocket prices in the imputation process may lead to an *underestimate* of expenditures if, for example, more generous insurance coverage causes individuals to substitute from lower to higher cost services within a utilization category.[82] Ultimately, I employ both measures of imputed expenditures with the understanding that true expenditures are likely to lie somewhere in between both of my measures.

Assessment of Imputation.—In order to assess the quality of my imputation I exclude a random subset of HRS observations for which I have the survey measure of expenditures, which I denote as the "HRS hold-out" sample. First, Figures 8 - 11 plot binned scatters of the relationship between actual log expenditures (as measured in the HRS) and the predicted level of log expenditures for the two imputation measures across various subgroups of the HRS hold-out sample.[83] Figure 8 plots these values for the entire HRS hold-out sample, whereas Figures 9, 10, and 11 plot them for the main analysis sample, individuals in expansion states, and individuals in non-expansion states, respectively. Overall, these figures suggest that my predicted log expenditures measure is strongly correlated with actual expenditures for the hold-out group.

Additionally, Figures 12 and 13 plot the distribution of the prediction errors for several subgroups of the hold-out sample for each of the imputation measures. Broadly, these figures demonstrate that the error in the predicted measure does not seem to be systematically different across expansion and non-expansion states in the hold-out group. There does, however, seem to be a slight difference in the distribution of errors across Medicaid status, where predicted expenditures for Medicaid enrollees are somewhat more likely to be underestimates than for individuals not enrolled in Medicaid. To the extent that this affects my estimates, I suspect that this may lead to my estimates being biased towards zero. In this case, my estimates could be perceived as a lower-bound of the effect of Medicaid coverage on current expenditures.

[82]Two examples where there would be no change in the utilization measures, but would lead to differences in total expenditures would be switching from generic to name-brand drugs or switching from a low cost clinic to a higher cost doctor's office.

[83]The figures use the insurance excluded measure of predicted log expenditures, but the results are similar for either measure. Conveniently, all of the observations in the HRS hold-out sample have strictly positive medical expenditures which allows me to log transform the variable without biasing the sample included in the figure.

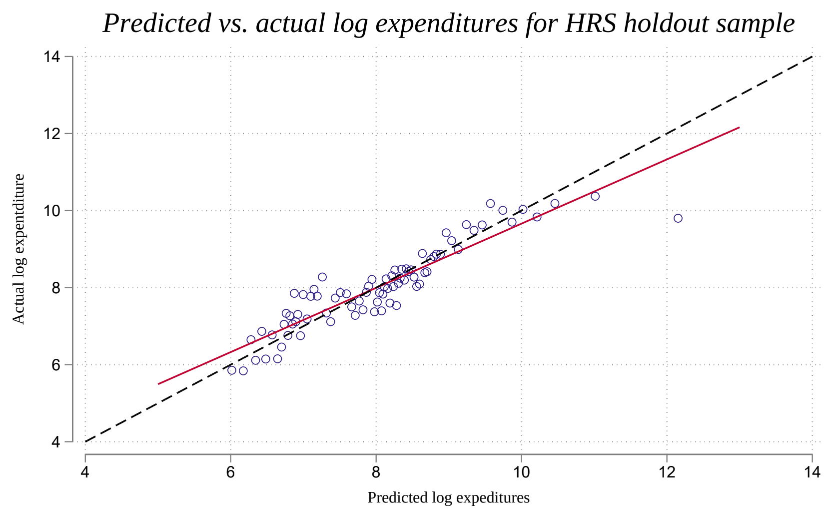

Note: Includes all HRS holdout observations between ages 60 and 70.

FIGURE 8

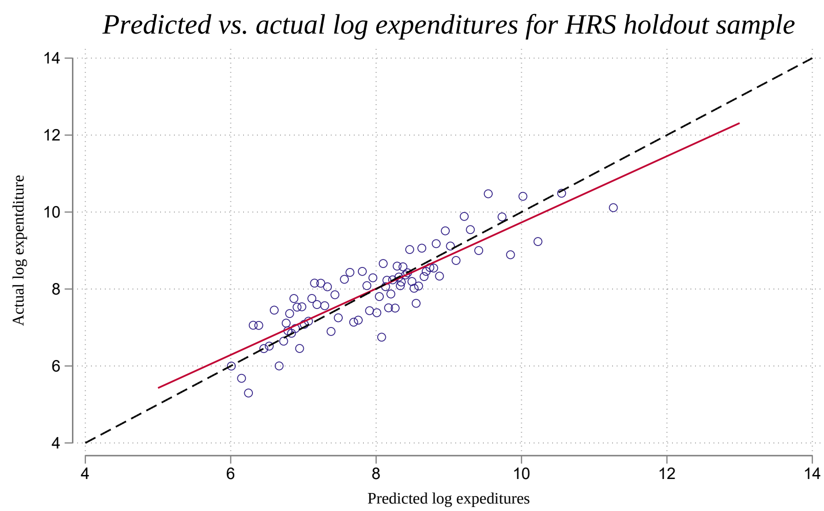

Note: Includes only observations in main analysis sample.

FIGURE 9

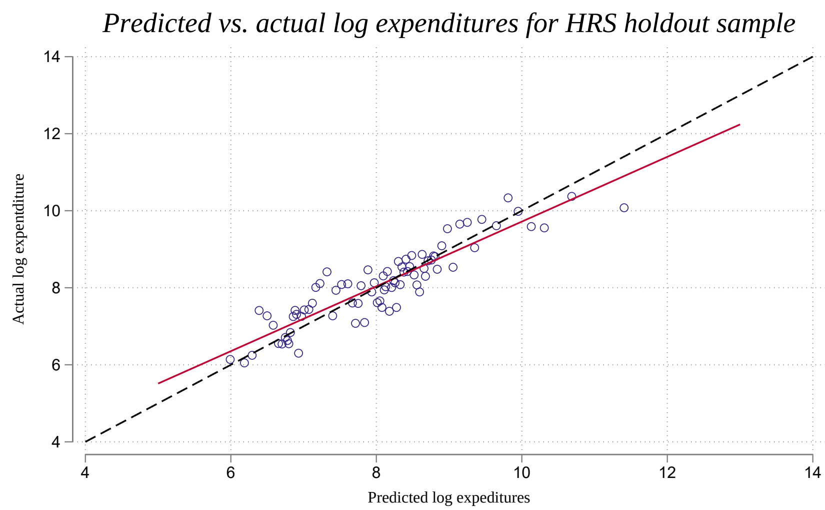

Note: Only includes observations in expansion states.

Figure 10

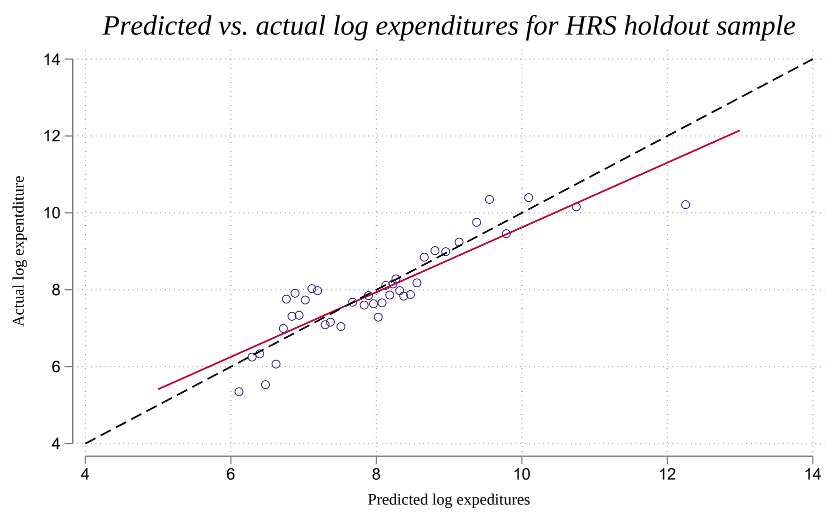

Note: Only includes observations in non-expansion states.

Figure 11

Table 16. Regression results for imputation process

	(1)	(2)	(3)	(4)
Age	-0.827	0.748	-0.916	0.565
	(2.448)	(1.169)	(2.486)	(1.174)
Age^2	0.016	-0.009	0.017	-0.007
	(0.043)	(0.016)	(0.044)	(0.016)
Age^3	-0.000	0.000	-0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Female	-0.008	0.017	-0.002	0.020
	(0.033)	(0.029)	(0.034)	(0.029)
Race:				
Black	-0.087**	-0.121***	-0.079*	-0.120***
	(0.041)	(0.035)	(0.041)	(0.035)
Hispanic	-0.104**	-0.161***	-0.172***	-0.159***
	(0.044)	(0.040)	(0.043)	(0.039)
Other	-0.116*	-0.089	-0.127*	-0.071
	(0.064)	(0.057)	(0.065)	(0.056)
Marital status:				
Married	-0.078	0.176***	-0.094*	0.161***
	(0.053)	(0.061)	(0.054)	(0.061)
Widowed	-0.072	0.131**	-0.085	0.112*
	(0.072)	(0.062)	(0.073)	(0.062)
Divorced	0.005	0.203***	-0.010	0.200***
	(0.056)	(0.065)	(0.057)	(0.065)
Separated	0.007	0.103	0.001	0.084
	(0.087)	(0.101)	(0.088)	(0.101)
Educational attainment:				
Some high-school	0.059	0.018	0.005	-0.002
	(0.057)	(0.047)	(0.050)	(0.042)
High-school	0.105**	0.077*	0.067	0.049
	(0.052)	(0.042)	(0.042)	(0.034)
Some college	0.202***	0.075	0.147***	0.048
	(0.061)	(0.051)	(0.052)	(0.044)
College	0.293***	0.119*	0.235***	0.086
	(0.080)	(0.063)	(0.074)	(0.057)
Graduate degree	0.190*	0.116	0.134	0.087
	(0.106)	(0.084)	(0.102)	(0.080)
Income:				
$\$5,000-\$10,000$	0.010	0.093	0.092	0.112
	(0.120)	(0.121)	(0.121)	(0.121)
$\$10,000-\$15,000$	0.092	0.124	0.154	0.125
	(0.135)	(0.136)	(0.137)	(0.137)
$\$15,000-\$20,000$	-0.023	0.065	0.062	0.064
	(0.149)	(0.152)	(0.151)	(0.153)
$\$20,000-\$15,000$	0.013	0.061	0.119	0.056
	(0.160)	(0.165)	(0.162)	(0.166)
$\$25,000-\$30,000$	0.071	0.106	0.207	0.100
	(0.169)	(0.176)	(0.171)	(0.176)
$\$30,000-\$35,000$	0.085	0.158	0.236	0.161
	(0.177)	(0.185)	(0.179)	(0.185)
$\$35,000-\$40,000$	0.078	0.070	0.217	0.059
	(0.186)	(0.193)	(0.187)	(0.193)
$\$40,000-\$45,000$	-0.007	0.129	0.151	0.130
	(0.191)	(0.201)	(0.193)	(0.202)
$\$45,000-\$50,000$	-0.005	0.101	0.168	0.100
	(0.199)	(0.212)	(0.201)	(0.212)
$\$50,000+$	0.107	0.038	0.277	0.027
	(0.207)	(0.217)	(0.209)	(0.218)
Working	-0.154***	-0.186***	-0.128***	-0.180***
	(0.036)	(0.046)	(0.035)	(0.046)
Insurance coverage:				
Private	0.494***	0.174***		
	(0.039)	(0.031)		
Medicaid	0.536***	0.172***		
	(0.049)	(0.037)		
Other public	0.405***	0.088*		
	(0.076)	(0.052)		
Healthcare utilization:				
Total out-of-pocket spending	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
Any doctor visit	0.019***	0.014***	0.020***	0.015***
	(0.001)	(0.001)	(0.001)	(0.001)
Any hospital visits	0.486***	0.426***	0.494***	0.429***
	(0.027)	(0.019)	(0.027)	(0.019)
Nights in the hospital	0.011***	0.005***	0.012***	0.005***
	(0.003)	(0.002)	(0.003)	(0.002)
Any at home healthcare	0.592***	0.619***	0.647***	0.629***
	(0.086)	(0.041)	(0.087)	(0.041)
Any dentist visits	0.506***	0.246***	0.554***	0.255***
	(0.032)	(0.028)	(0.032)	(0.028)
Any R_xs	1.026***	1.149***	1.079***	1.155***
	(0.044)	(0.062)	(0.045)	(0.062)
Health:				
Cancer	0.385***	0.240***	0.418***	0.237***
	(0.060)	(0.034)	(0.061)	(0.034)
Heart problems	0.312***	0.258***	0.332***	0.264***
	(0.044)	(0.029)	(0.044)	(0.029)
Lung problems	0.292***	0.142***	0.289***	0.141***
	(0.059)	(0.041)	(0.060)	(0.041)
Stroke	0.201***	0.093**	0.233***	0.096**
	(0.069)	(0.038)	(0.070)	(0.038)
High blood-pressure	0.128***	0.055*	0.123***	0.052
	(0.034)	(0.032)	(0.035)	(0.032)
Diabetes	0.409***	0.364***	0.419***	0.367***
	(0.041)	(0.030)	(0.042)	(0.030)
Arthritis	0.252***	0.188***	0.270***	0.191***
	(0.033)	(0.028)	(0.034)	(0.028)
Smoker	0.028	0.050	0.018	0.048
	(0.039)	(0.044)	(0.040)	(0.044)
Struggles with any ADLs	0.036	0.219***	0.005	0.231***
	(0.102)	(0.060)	(0.104)	(0.061)
Struggles with any IADLs	0.095	0.115**	0.093	0.123**
	(0.059)	(0.047)	(0.059)	(0.048)
Intercept	20.061	-13.703	22.186	-8.950
	(46.249)	(29.049)	(46.967)	(29.169)
Sample	Pre-Medicare	Medicare	Pre-Medicare	Medicare
Insurance included	✓	✓		
Adj-R^2	.49	.56	.48	.55
Obs.	7130	5889	7130	5889

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 17. Summary of imputed expenditures and out-of-pocket spending, by group

	MEPS				HRS			
	Non-Elderly			Elderly	Non-Elderly			Elderly
	All	Medicaid	No Medicaid		All	Medicaid	No Medicaid	
Total expenditures	5,700	8,7987	4,965	10,132	6,575	13,341	5,406	11,573
	(9,597)	(11,793)	(8,609)	(12,603)	(21,223)	(32,613)	(18,263)	(27,478)
Total out-of-pocket spending	1,034	564	1,179	1,398	2,085	634	2,298	2,576
	(1,991)	(1,200)	(2,157)	(2,268)	(4,628)	(2,104)	(4,851)	(5,226)
Working share	0.55	0.25	0.64	0.08	0.55	0.25	0.60	0.06
	(0.50)	(0.43)	(0.48)	(0.28)	(0.50)	(0.43)	(0.49)	(0.23)

Note.—Total expenditures in the HRS are imputed. Number of observations varies by variable and sample.

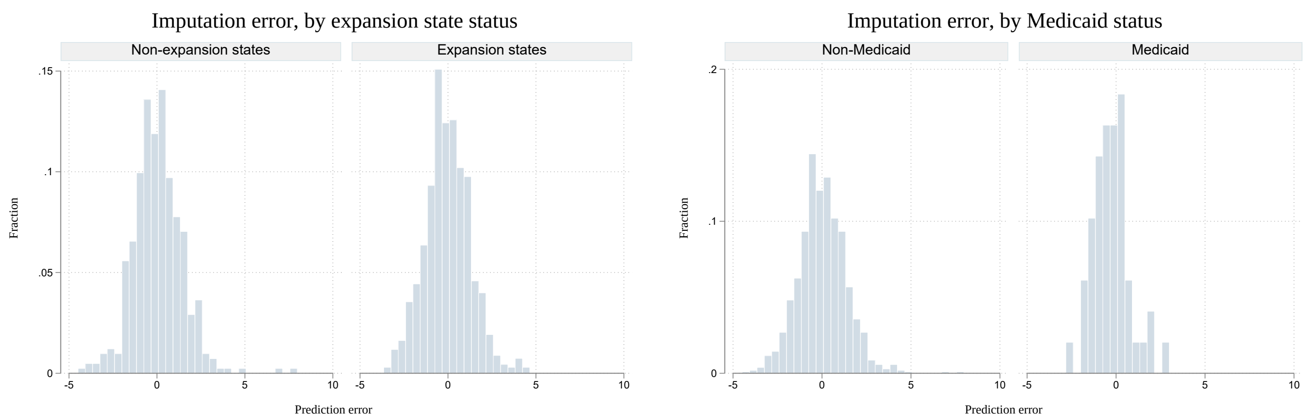

Figure 12

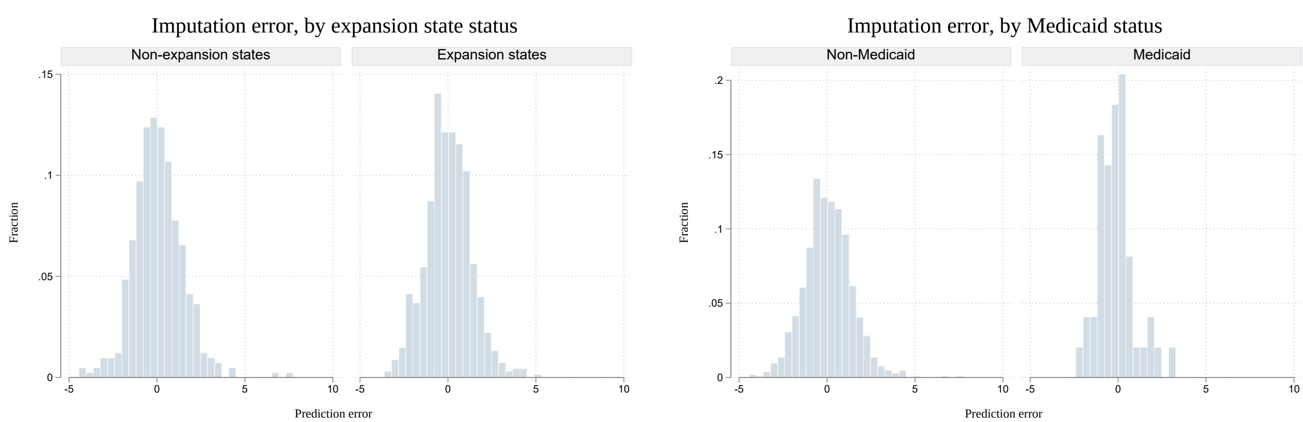

Figure 13

C Additional Discussion of Reduced Form Identification

Identification of the TD Estimator.—The validity of my TD strategy also requires that the composition of the treatment and control groups is stable across the repeated cross-sections. In practice this requires that individuals do not shift from the non-eligible to the eligible population or move from non-expansion to expansion states in response to the policy. A possible violation would arise if individuals are strategic in their labor supply choices such that they lower their income level below the eligibility threshold in response to the expansions. While this is possible, existing work has found little evidence that the Medicaid expansions affected employment outcomes or earnings. For example, Leung and Mas (2016); Buchmueller, Levy, and Valleta (2019); Callison and Sicilian (2016); Kaestner et al. (2015); and Gooptu et al. (2016) all find either no effects of the expansions on employment or earnings, or small positive gains in employment for some groups. Alternatively, Peng, Guo, and Meyerhoefer (2019) find a somewhat small and temporary decline in employment following the expansions but no effects on wages. Overall, these prior findings are suggestive that the low-income group is likely stable across treatment and control periods for expanding states. However, to be confident in this assumption, I also examine in the HRS whether eligibility, employment, income, or retirement were significantly affected by the Medicaid expansions. Though not reported, I find that these outcomes were unaffected by the expansions for individuals in my sample. The other key threat to this assumption arises if individuals were more likely to move to expansion states in the pre-Medicare period in order to obtain insurance coverage. However, previous evidence suggests that migration decisions are uncorrelated with state Medicaid expansions (Goodman 2017).[84] Moreover, considering the age and income of my sample, I find this type of threat even less likely.

Additional Threats to Identification.—In addition to the cases outlined previously, there are several other threats to identification. First, the data collection in the HRS somewhat confounds the interpretation of either π_1 or α_1 as the LATE of Medicaid. Respondents are reported as on Medicaid if they are enrolled at the time of the survey, but respondents provide information on healthcare utilization retrospectively for the two years prior to the current wave. As a result, these specifications may confound the role that Medicaid is playing in determining total expenditures. For instance, if individuals spend down their income in the two years leading up to the survey due to medical costs and then go on Medicaid, my estimates would overestimate the role of Medicaid on spending. Alternatively, if individuals switch off of Medicaid throughout the survey period, but spend substantially while on Medicaid, my estimates would be understating the effect of coverage. At the moment, I do not see a clear way to avoid this concern with the current data limitations.

Second, Finkelstein et al. (2012) also point out that exogenous expansion in Medicaid eli-

[84] Though not currently presented, I also do not find evidence of migration in response to the Medicaid expansions in the HRS.

gibility may affect enrollment in other government programs such as Supplemental Nutrition Assistance Program (SNAP) or Temporary Assistance for Needy Families (TANF). Enrollment in these programs could result in higher levels of nutrition and health which may decrease total expenditures, which would introduce a downward bias in my estimates. On the other hand, cash-equivalent transfer programs alter the household budget constraints for low-income individuals. This may result in an even greater spending response than if the individual received Medicaid alone. To explore this possibility I estimate equation (1), but I replace the dependent variable with the total income from veterans' benefits, food stamps, and welfare. The results, which are imprecisely estimated, suggest that government payments actually decrease around the time of the expansions, but only by around $100 over a two-year period. This highlights a potentially interesting relationship between various public assistance programs, but does not indicate a clear threat to my identifying assumptions.

Finally, prior work documents that the Medicaid expansions significantly decreased mortality for near-elderly individuals (e.g., Miller, Wherry, Johnson, and Wherry 2021). This may lead my estimates for the affect of Medicaid coverage on initial expenditures to be upward biased or attenuate my estimates for how prior Medicaid coverage affects future expenditures. Specifically, if the reduction in mortality in expansions states leads to a decline in health for the population of low-income people in those states, this may drive my estimates of the expenditure response up. At the same time, reduced mortality in expansion states may also lead to sicker people aging onto Medicare, which could artificially drive up healthcare consumption in the treatment group. To try and address these concerns, I explore the extent to which the Medicaid expansions affected mortality in individuals aged 60-64 in the HRS. Though currently unreported, I do find suggestive evidence that the expansions may have decreased the likelihood of death for low-income 60-64 year old respondents in the HRS.[85]

Understanding exactly how mortality may affect my estimates is complex. For example, if individuals with more future expenditure risk consume more healthcare due to receiving Medicaid through the expansions and that prolongs their lives, this has the potential to inflate expenditures in the treatment group before and after aging into Medicare. As previously state, this is likely to attenuate my estimates during the Medicare period, but may upward bias the estimates in the pre-Medicare period as well. However, for the pre-Medicare period, it is not obvious that incorporating this effect into the estimate is problematic, given that this is a relevant channel through which insurance, if it has dynamic affects, is likely to alter expenditures across time. In fact, from a policy maker's point of view, accounting for the additional spending accruing to sick individuals who would have died in the absence of the expansions is extremely important in understanding the overall costs of the policy. Ultimately, I do not directly address these concerns in estimation, but they are important to keep in mind when interpreting my results.

[85]The estimated coefficients for linear probability and logistic TD regression models are negative and economically significant, but only the point estimate of the linear model is marginally statistically significant.

My hope is that by conditioning on observable health measures, I am able to control for any underlying compositional change in the health of the treatment and control group populations that might be introduced through this mortality response to the expansions.

D Identification of the TDIV Estimator

To demonstrate identification of the TDIV estimator employed in this paper, I adopt the notation and simplifying assumptions of Hudson, Hull, and Libersohn (2017) (henceforth, HHL). Namely, I will assume that there is a time varying outcome Y_{it}, a binary treatment M_{it}, and a binary instrument Z_{it} for individuals i in time periods $t \in \{0,1\}$.[86] Let $T_t = 1\{t=1\}$ denote the period in which the instrument is introduced, such that $Z_{it=0} = 0 \forall i$. Additionally, assume that the instrument can be expressed as the product of two binary indicators S_{it} (e.g. geographic location) and L_{it} (e.g., low-income status) which jointly determine period-1 exposure to the instrument: $Z_{i1} = S_{i1}L_{i1}$. In a manner analogous to HHL, the TDIV coefficient β_{TDIV} comes from the following IV system:

$$Y_{it} = \alpha_i + \beta_{TDIV} M_{it} + \gamma_1 L_{it} + \gamma_2 T_t + \gamma_3 L_{it} T_t + \gamma_4 S_{it} T_t + \gamma_5 S_{it} L_{it} + \varepsilon_{it} \tag{23}$$

$$M_{it} = \phi_i + \pi Z_{it} T_t + \psi_1 L_{it} + \psi_2 T_t + \psi_3 L_{it} T_t + \psi_4 S_{ti} T_t + \psi_5 S_{it} L_{it} + \eta_{it}. \tag{24}$$

Given that the IV model is just-identified, β_{TDIV} can be re-expressed as the ratio of the reduced form and first stage coefficients, where the reduced form coefficient is ρ in the following regression equation:

$$Y_{it} = \mu_i + \rho Z_{it} T_t + \theta_1 L_{it} + \theta_2 T_t + \theta_3 L_{it} T_t + \theta_4 S_{it} T_t + \theta_5 S_{it} L_{it} + \nu_{it}, \tag{25}$$

and the first-stage coefficient is π from Equation (24). As shown in Olden and Møen (2022), each of these estimands can be expressed as:

$$\begin{aligned} \rho =& \big(E[Y_{i1} - Y_{i0}|S_i = 1, L_i = 1] - E[Y_{i1} - Y_{i0}|S_i = 1, L_i = 0]\big) \\ &- \big(E[Y_{i1} - Y_{i0}|S_i = 0, L_i = 1] - E[Y_{i1} - Y_{i0}|S_i = 0, L_i = 0]\big) \end{aligned} \tag{26}$$

$$\begin{aligned} \pi =& \big(E[M_{i1} - M_{i0}|S_i = 1, L_i = 1] - E[M_{i1} - M_{i0}|S_i = 1, L_i = 0]\big) \\ &- \big(E[M_{i1} - M_{i0}|S_i = 0, L_i = 1] - E[M_{i1} - M_{i0}|S_i = 0, L_i = 0]\big) \end{aligned} \tag{27}$$

where

$$\beta_{TDIV} = \frac{\rho}{\pi}.$$

[86] The assumptions on the value of the treatment and the assumption that the instrument is binary are met in my setting. As with standard difference-in-differences, it is possible to extend this analysis to the case with repeated cross-sections and multiple time periods as in Abadie (2005).

In order to arrive at Equations (26) and (27), note that it was required to drop the t subscript from the L and S terms. Implicitly, this change in notation assumes that the triple-difference assumption of group-stability (in both S and L) is fixed in the pre- and post-treatment periods. Specifically, by writing L_i and S_i as time-invariant, I am imposing that the composition of the treatment and control groups is invariant to the introduction of the instrument (i.e., policy change at $t = 1$).

Now, let Y_{it}^{m} denote the potential outcome of individual i in time t if exposed to treatment level m, and let M_{it}^{z} denote an individual's potential treatment level if exposed to instrument value z in t. As discussed in HHL, this notation implies the standard IV exclusion restriction that the instrument only affects outcomes through treatment as well as an additional restriction that treatment and outcomes in period $t = 0$ are unaffected by the value of the instrument in period $t = 1$.

In addition to the restrictions outlined above, I make two additional assumptions. The first is the TD parallel trend assumption for both treatment and outcomes:[87]

A1:

$$E[M_{i1}^{0} - M_{i0}^{0}|S_i = 1, L_i = 1] - E[M_{i1}^{0} - M_{i0}^{0}|S_i = 1, L_i = 0]$$
$$=$$
$$E[M_{i1}^{0} - M_{i0}^{0}|S_i = 0, L_i = 1] - E[M_{i1}^{0} - M_{i0}^{0}|S_i = 0, L_i = 0]$$

and

$$E[Y_{i1}^{M_{i1}^{0}} - Y_{i0}^{M_{i0}^{0}}|S_i = 1, L_i = 1] - E[Y_{i1}^{M_{i1}^{0}} - Y_{i0}^{M_{i0}^{0}}|S_i = 1, L_i = 0]$$
$$=$$
$$E[Y_{i1}^{M_{i1}^{0}} - Y_{i0}^{M_{i0}^{0}}|S_i = 0, L_i = 1] - E[Y_{i1}^{M_{i1}^{0}} - Y_{i0}^{M_{i0}^{0}}|S_i = 0, L_i = 0].$$

The second assumption is the standard IV assumption of monotonicity:

A2 (*Monotonicity*): $P(M_{i1}^{1} \geq M_{i1}^{0}) = 1$.

[87] I provide intuition behind this assumption in Section 4.

Under **A1** and the assumption that Z_i is exogenous, the first stage coefficient identifies:

$$\begin{aligned}
\pi &= \big(E[M_{i1} - M_{i0}|S_i = 1, L_i = 1] - E[M_{i1} - M_{i0}|S_i = 1, L_i = 0]\big) \\
&\quad - \big(E[M_{i1} - M_{i0}|S_i = 0, L_i = 1] - E[M_{i1} - M_{i0}|S_i = 0, L_i = 0]\big) \\
&= \big(E[M^1_{i1} - M^0_{i0}|S_i = 1, L_i = 1] - E[M^0_{i1} - M^0_{i0}|S_i = 1, L_i = 0]\big) \\
&\quad - \big(E[M^0_{i1} - M^0_{i0}|S_i = 0, L_i = 1] - E[M^0_{i1} - M^0_{i0}|S_i = 0, L_i = 0]\big) \\
&= E[M^1_{i1} - M^0_{i1}|S_i = 1, L_i = 1] \\
&= E[M^1_{i1} - M^0_{i1}|Z_i = 1] \\
&= E[M^1_{i1} - M^0_{i1}]
\end{aligned}$$

and the reduced form identifies:

$$\begin{aligned}
\rho &= \big(E[Y_{i1} - Y_{i0}|S_i = 1, L_i = 1] - E[Y_{i1} - Y_{i0}|S_i = 1, L_i = 0]\big) \\
&\quad - \big(E[Y_{i1} - Y_{i0}|S_i = 0, L_i = 1] - E[Y_{i1} - Y_{i0}|S_i = 0, L_i = 0]\big) \\
&= \big(E[Y^{M^1_{i1}}_{i1} - Y^{M^0_{i0}}_{i0}|S_i = 1, L_i = 1] - E[Y^{M^0_{i1}}_{i1} - Y^{M^0_{i0}}_{i0}|S_i = 1, L_i = 0]\big) \\
&\quad - \big(E[Y^{M^0_{i1}}_{i1} - Y^{M^0_{i0}}_{i0}|S_i = 0, L_i = 1] - E[Y^{M^0_{i1}}_{i1} - Y^{M^0_{i0}}_{i0}|S_i = 0, L_i = 0]\big) \\
&= E[Y^{M^1_{i1}}_{i1} - Y^{M^0_{i1}}_{i1}|S_i = 1, L_i = 1] \\
&= E[Y^{M^1_{i1}}_{i1} - Y^{M^0_{i1}}_{i1}|Z_i = 1] \\
&= E[Y^{M^1_{i1}}_{i1} - Y^{M^0_{i1}}_{i1}],
\end{aligned}$$

such that:

$$\beta_{TDIV} = \frac{E[Y^{M^1_{i1}}_{i1} - Y^{M^0_{i1}}_{i1}]}{E[M^1_{i1} - M^0_{i1}]}. \tag{28}$$

If the denominator is nonzero (i.e., the instrument is relevant in the first stage), then by **A2** and Theorem 2 of Imbens and Angrist (1994), TDIV identifies the LATE:

$$\beta_{TDIV} = E[Y^1_{i1} - Y^0_{i1}|M^1_{i1} \neq M^0_{i1}].$$

E Additional Empirical Results

E.1 Effects of the Medicaid Expansions on Initial Utilization and Health

Effects on Other Utilization.—Overall, the results in Tables 2 indicate that there was a large and positive expenditure response to the increased Medicaid coverage generated by the expansions. However, the results presented thus far rely on imputed expenditures. To explore what is driving the increase in expenditures, I also estimate versions of Equations (3) and (4) where I replace the expenditure measure with out-of-pocket spending as well as other utilization measures that are observed directly rather than imputed. I present these results in panel A of Table 18, where each row of columns 1 and 2 present the coefficient estimate on the triple-difference term and (instrumented) Medicaid coverage, respectively, for each dependent variable.

Generally, the results in Table 18 suggest that other measures of utilization also increased or stayed the same following the Medicaid expansions. In particular, the estimates suggest that the Medicaid expansions increased the likelihood of going to the doctor, with an ITT of 9 percentage points and a LATE of 46 percentage points. The coefficient estimates for the models of intensive margin doctor visits are also positive and large in magnitude, but not statistically significant. Interestingly, there is a small but imprecise negative effect on extensive hospitalization, but a positive (but imprecise) effect on intensive margin hospitalization. Though imprecisely estimated, these results could be suggestive of a substitution of uncompensated emergency department care for doctor office visits. At the same time, for individuals who truly require hospital services (beyond what a doctor's office or clinic can provide), their consumption of services increases (albeit statistically insignificantly). There was also a positive and (marginally) statistically significant effect on the use of any at home healthcare services (ITT of 3.7 percentage points and LATE of 19.2 percentage points). Also of note is that out-of-pocket spending seems to decrease (though the estimate is very imprecise) even though overall expenditures increased. This makes sense given the dramatic reduction in cost-sharing that Medicaid offers in comparison to either private insurance or uninsurance.[88] Overall, I interpret this evidence as reaffirming the results in Table 2 and, in a way, further validating the healthcare expenditure imputations as a reliable measurement of actual healthcare spending.

Effects on Health.—In addition to determining the effect of the Medicaid expansions on expenditures and utilization, I assess the extent to which these policies contributed to contemporaneous health. To do so I re-estimate Equations (3) and (4) with two different measures of health as the dependent variable. The first health measure is self-reported health, which ranges from 1 ("Poor" health) to 5 ("Excellent" health).[89] The second measure of health is the health index described in Section 3. The results from this analysis are presented in panel B of Table 18. All of the coefficient estimates are positive, but imprecise. The positive sign indicates that increased insurance coverage though the expansions increases self-reported health (e.g. a higher probability of moving from "Fair" to "Good" health) as well as the health index. The lack of precision suggest that healthcare coverage does not necessarily increase health outcomes immediately. In subsequent analysis, I explore how this changes overtime.

Heterogeneity by Past Insurance Coverage.—Before moving on to explore how the Medicaid ex-

[88]I suspect that out-of-pocket expenses did fall, but the design of the survey questions and the low-frequency of the data makes this decrease difficult to detect. Using a different empirical design and slightly different measure of out-of-pocket spending, McInerney et al., 2020 find that out-of-pocket costs fell by between \$500 and \$650 for low-income near elderly individuals in the HRS.

[89]Since the self-reported health outcome is categorical, the linear specification in (3) is not completely appropriate; however, I think it is a decent approximation since I am only interested in a broad relationship between health and insurance coverage.

TABLE 18. EFFECT OF EXPANSIONS ON OTHER OUTCOMES

	TD (1)	TDIV (2)
A. ***Other utilization***		
Any doctor visits	0.090*	0.462*
	(0.050)	(0.265)
Doctor visits (counts)	1.595	8.191
	(1.249)	(6.916)
Any hospitalizations	-0.034	-0.176
	(0.058)	(0.296)
Hospitalizations (counts)	0.013	0.069
	(0.102)	(0.526)
Nights in hospital	0.121	0.622
	(0.476)	(2.468)
Any Rxs	-0.002	-0.011
	(0.040)	(0.204)
Any at home care	0.037*	0.192†
	(0.023)	(0.121)
Total out-of-pocket spending	-150	-770
	(579)	(2957)
B. ***Health outcomes***		
Self-reported health	0.087	0.442
	(0.099)	(0.515)
Health index	0.002	0.013
	(0.004)	(0.026)
Sample	2014 expansions only	2014 expansions only
Controls included	✓	✓
Individual FEs	✓	✓
First-stage F-stat		20.82
Underid. P-val		0.001
Obs.	17081	17081

Note.—Standard errors clustered at the state level in parentheses.
† $p < 0.15$, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

pansions affect future outcomes once on Medicare, I also examine heterogeneity in the response to the expansions across prior insurance coverage. As shown in Table 1, a large proportion of new Medicaid enrollees appear to have been previously uninsured. Motivated by this evidence, I explore how my main findings for expenditures and other outcomes vary across prior uninsured status. To do this I estimate versions of the following TD equation:

$$Y_{i,t} = \phi + \delta Expand_{s,t} \times LI_{i,t} \times Uninsured_{i,t-1} + X_{i,t}'\Omega + \Psi + \varepsilon_{i,t} \tag{29}$$

where $Uninsured_{i,t-1}$ is an indicator if individual i was uninsured in the previous wave and Ψ is a vector of indicators which includes interactions between $Expand$ and $Uninsured$, LI and $Uninsured$, as well as all of the standard indicators and interactions in the TD model and individual-level fixed effects. Under the identifying assumptions described previously, I can interpret δ_1 as the differential effect of the expansions on low-income individuals who were previously uninsured in expansion states. Table 19 presents results from estimating Equation (29) with all of the previously presented dependent variables.[90]

[90]Note that the sample size decreases slightly from the previous tables. This is due to the fact that I cannot observe previous insurance status for some individuals in the pre-Medicare sample.

Table 19. Heterogeneous response to the expansions by previous uninsured status

	$\hat{\delta}$
Dependent variable:	
$\widehat{\ln(E_{i,t})}$, (ins. included)	0.290**
	(0.130)
$\widehat{\ln(E_{i,t})}$, (ins. excluded)	0.187†
	(0.123)
Any doctor visits	0.196**
	(0.089)
Doctor visits (count)	1.961
	(1.687)
Any hospitalizations	-0.057
	(0.083)
Hospitalizations (counts)	-0.122
	(0.125)
Nights in hospital	-0.346
	(0.471)
Any Rxs	0.135*
	(0.081)
Any at home care	0.014
	(0.039)
Self-reported health	-0.155
	(0.167)
Health index	-0.008
	(0.010)
Sample	2014 expansions only
Controls included	✓
Individual FEs	✓
Obs.	16260

Note.—Standard errors clustered on state in parentheses.
† $p < .15$, * $p < .1$, ** $p < .05$, *** $p < .01$

The results in the table show that a large share of the increase in healthcare utilization is driven by the previously uninsured. This makes sense given that they experience the greatest change in out-of-pocket prices when receiving Medicaid coverage. However, the coefficient estimate for the model with “Any Rxs” as the dependent variable is positive and statistically significant (which was not the case previously). This finding suggests that not only are previously uninsured more likely to consume healthcare, but that they are specifically more likely to go to the doctor and consume prescription drugs. These consumption patterns are encouraging because they suggest that the expansions are increasing access to certain types of preventative care and chronic disease management services that are likely to have dynamic benefits. Somewhat puzzling, however, is that contemporaneous health measures also remain unaffected for this group as well.

E.2 Effects of the Medicaid Expansions on Future Health and Diagnoses

Figure 14 plots results from an event study specification where the health index is the dependent variable. The figure demonstrates that prior to treatment, there were (noisy but) common trends between the treatment and control groups, but following the expansions, health during

Figure 14. Variation in aggregate health index once on Medicare

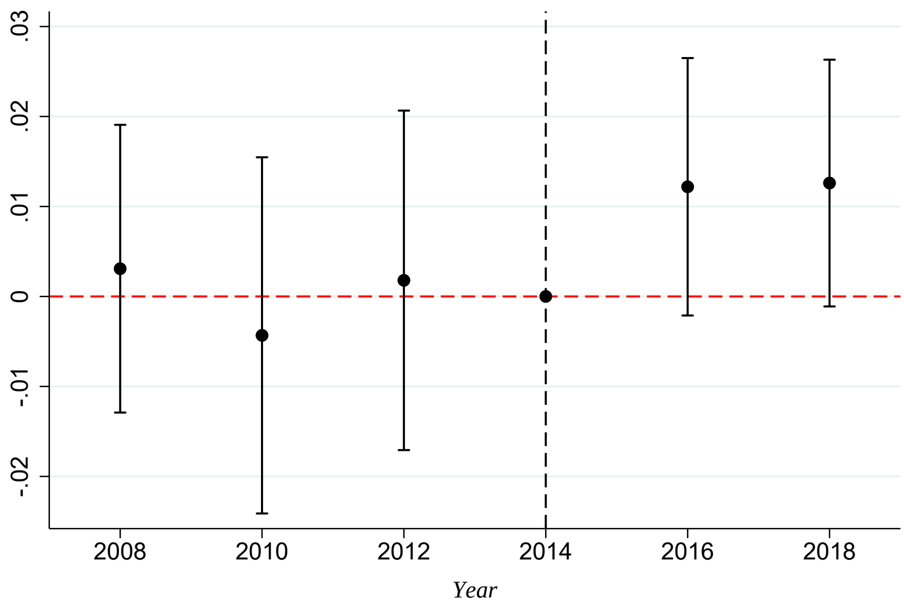

Notes.—Figure displays coefficient estimates and 95% state-level clustered confidence intervals from an event study regression of the heath index during the first wave of Medicare coverage onto the indicators for the survey wave relative to the Medicaid expansion interacted with previous low-income status. The estimating sample is limited individuals in states that expanded in 2014 or did not expand during the sample period. The regression model includes a rich set of observable controls for demographics and lagged health status.

the first two years of Medicare coverage increased for individuals previously exposed to the expansions. Moreover, the consistency in the point estimates for 2016 and 2018 suggest that this effect persisted for both waves of individuals affected by the Medicaid expansions prior to age 65.

Table 20 presents results for the effect of previous exposure to the Medicaid expansions on diagnoses during the first wave of Medicare. Interestingly, even though overall health appears to increase for these individuals, rates of diabetes diagnoses actually increase. The point estimate for the Diabetes model in column 2 suggests that past exposure to the expansions increases rates of diabetes diagnoses by 5 percentage points. Given the results that I find for other health measures, I think it is unlikely that the expansions are *causing* diabetes to become more common in the population. Rather, having access to insurance is most likely increasing detection of cases that would have otherwise gone undetected. Alternatively, the negative (and marginally significant) point estimates on the Stroke model suggest that the expansions may have led to lower rates of strokes (which are more acute phenomena that may be avoided through proper management of chronic disease).

TABLE 20. EFFECT OF LAGGED MEDICAID EXPANSION EXPOSURE ON CURRENT HEALTH DIAGNOSES

Dependent variables:	(1)	(2)	(3)
High blood pressure	-0.009	-0.016	-0.006
	(0.028)	(0.029)	(0.030)
Diabetes	0.039†	0.050*	0.052**
	(0.026)	(0.026)	(0.024)
Stroke	-0.007	-0.016†	-0.009
	(0.011)	(0.010)	(0.010)
Heart disease	0.013	0.024	0.024
	(0.019)	(0.020)	(0.022)
Lung disease	0.006	0.009	0.013
	(0.025)	(0.027)	(0.027)
Sample	2014 expansions only	2014 expansions only	2014 expansions only
Health$_{t-2}$ controls		✓	
Health$_{t-3}$ controls			✓
Obs.			

Note.—Standard errors clustered at the state-level in parentheses.
† $p < .15$, * $p < .1$, ** $p < .05$, *** $p < .01$

E.3 Relationship Between Health and Mortality

Throughout the paper I use self-reported health and a general health index as key measures of health. Table 21 present results that describe the association between these measures of health and future mortality. Specifically. I estimate regressions of the following form:

$$Mortality_{i,t+1} = \beta_0 + \beta_1 Health_{i,t} + X_{i,t}\Gamma + \varepsilon_{i,t} \tag{30}$$

where $Mortality$ is an indicator equal to one if the individual died in between survey waves t and $t+1$, $Health$ is one measure of health status in t, and X is a vector of additional controls. Columns 1 and 2 show that self-reported health is strongly associated with mortality. For the main sample in my analysis (column 1) and for the entire HRS sample aged 55-95 (column 2) there is substantially higher rates of mortality among individuals in "Fair" or "Poor" health. Relative to individuals in "Excellent" health, individuals in "Poor" health in my main analysis sample are 4.8 percentage points more likely to die over the next two years. For the entire sample aged 55-95, the difference is 8.6 percentage points. Columns 3 and 4 report similar results but when I replace self-reported health with the general health index. The index lies within 0 and 1, with a mean of 0.7 and standard deviation of 0.047 for the pre-Medicare sample. Thus the coefficient estimate in column 3 suggests that a standard deviation decrease in the health index is associated with an increase in the likelihood of mortality by roughly 1.4 percentage points. The last two columns of the table report similar results using another measure of health called a frailty index that is commonly used in papers that have life-cycle models (see Russo 2023). This index is constructed from counts of various diagnoses and difficulties with ADLs and IADLs. I include this measure to demonstrate that both self-reported health and the general health index are similarly associated with mortality as the frailty index.

Table 21. Relationship between various health measures and future mortality

	(1) Self-report	(2) Self-reported	(3) Health index	(4) Health index	(5) Frailty	(6) Frailty
Self-reported health						
Very good	0.002	0.001				
	(0.002)	(0.002)				
Good	0.007***	0.014***				
	(0.003)	(0.002)				
Fair	0.019***	0.042***				
	(0.003)	(0.002)				
Poor	0.048***	0.086***				
	(0.005)	(0.003)				
Health index			-0.288***	-0.731***		
			(0.016)	(0.006)		
Frailty					0.007***	0.023***
					(0.000)	(0.000)
Sample	Main sample	Everyone aged 55-95	Main sample	Everyone aged 55-95	Main sample	Everyone aged 55-95
Obs.	27771	175457	27443	173319	27777	175566

Note.—Standard errors clustered at the state-level in parentheses.

* $p < .1$, ** $p < .05$, *** $p < .01$

F Robustness of Empirical Findings

F.1 Results for Alternative Imputation Methods

Panel A of Table 22 re-displays the results from Table 2 and panel B presents estimates from the same sets of specifications and estimating samples, but where the dependent variable is the version of imputed log of total expenditures that does not incorporate insurance coverage directly in the imputation process. The coefficient estimates are largely consistent with those in panel A, though consistently smaller in magnitude and the point estimates in columns 3 and 6 are also no longer statistically significant. This asymmetry across the panels is consistent with my expectation outline in Appendix B that my imputation including insurance coverage may *overestimate* expenditures whereas the measure of expenditures excluding insurance in the imputation likely *underestimates* true expenditures. Importantly, though, the point estimates across each of the panels are not statistically different from each other which I take as consistent evidence that the Medicaid expansions lead to meaningful increases in medical expenditures for the low-income near-elderly population.

Table 22. Effect of expansion status on medical spending

	TD (1)	TD (2)	TD (3)	TDIV (4)	TDIV (5)	TDIV (6)
	A. First imputation method: Insurance status included					
Medicaid expansion × Low-income	0.265**	0.292**	0.172*			
	(0.106)	(0.098)	(0.100)			
	[0.296]	[0.333]	[0.182]			
Medicaid coverage				1.329**	1.228***	0.886*
				(0.528)	(0.449)	(0.515)
				[2.286]	[2.087]	[1.124]
	B. Second imputation method: Insurance status excluded					
Medicaid expansion × Low-income	0.189**	0.206**	0.104			
	(0.093)	(0.099)	(0.094)			
Medicaid coverage				0.946*	0.868**	0.534
				(0.483)	(0.414)	(0.495)
Sample	Whole sample	2014 expansions only	2014 expansions only	Whole sample	2014 expansions only	2014 expansions only
Controls included	✓	✓	✓	✓	✓	✓
Individual FEs			✓			✓
First-stage F-stat				18.2	25.51	20.82
Underid. P-val				0.001	0.001	0.001
Obs.	20050	18197	17081	20050	18197	17081

Note: * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

F.2 Alternative Eligibility Definitions

Table 23 presents the main empirical findings of the paper, but using the legal threshold for Medicaid expansion eligibility. Columns 1 and 2 present results from the first-stage regression (Equation (1)). As discussed in Section 3, this measure of eligibility is strongly associated with increased Medicaid enrollment, but somewhat less so than the rougher measure of eligibility used throughout the main text of the paper. Columns 3 and 4 present the TD results for the initial spending response to the Medicaid expansions for the pre-Medicare sample for both imputation measures. Though no longer statistically significant, the point estimates are of the same sign and magnitude as those reported in the main text. Moreover, the point estimates are not statistically different from one another when taking into account their confidence intervals. Columns 5-7 report point estimates from the TD regression on the Medicare sample for specifications with varying degrees of controls for past health status. Overall, these estimates are very consistent with those presented in Table 3 in the main text.

Table 23. Robustness using FPL ratio to define eligibility

	(1) First-stage	(2) First-stage	(3) Initial spending (ins. exclud.)	(4) Initial spending (ins. includ.)	(5) Medicare, no health controls	(6) Medicare, twice lagged health	(7) Medicare, thrice lagged health
Expansion$_t \times$ income$_t \leq$ 138% $\times$ FPL	0.148***	0.151***	0.172	0.195	-0.336**	-0.238*	-0.280**
	(0.046)	(0.039)	(0.129)	(0.139)	(0.134)	(0.121)	(0.135)
Obs.	20066	18875	17352	17352	8948	8032	6970

Note.—Standard errors clustered on state in parentheses.
* $p < .1$, ** $p < .05$, *** $p < .01$

F.3 Results for Additional Alternative Specifications

Table 24. Effect of expansions on various outcomes during the first wave of Medicare, TD

	(1)	(2)	(3)
A. *Expenditures*			
$\widehat{\ln(E_{i,t})}$, ins. included	-0.326**	-0.250**	-0.300**
	(0.145)	(0.113)	(0.126)
$\widehat{\ln(E_{i,t})}$, ins. excluded	-0.338**	-0.264**	-0.317**
	(0.145)	(0.113)	(0.126)
B. *Other utilization*			
Total out-of-pocket spending	-708.545	-815.627†	-851.892†
	(514.742)	(551.265)	(573.123)
Any hospitalizations	-0.092†	-0.077	-0.095*
	(0.059)	(0.059)	(0.056)
Hospitalizations (counts)	-0.177	-0.137	-0.150
	(0.127)	(0.128)	(0.126)
Nights in hospital	-2.088*	-1.716	-2.210*
	(1.253)	(1.322)	(1.194)
Any doctor visits	-0.088*	-0.044	-0.065
	(0.050)	(0.052)	(0.056)
Doctor visits ≥ 5	-0.183**	-0.162**	-0.223***
	(0.073)	(0.069)	(0.074)
Any Rxs	-0.044	-0.015	-0.043
	(0.055)	(0.051)	(0.056)
Any at home care	-0.052**	-0.050**	-0.052**
	(0.026)	(0.025)	(0.024)
C. *Health outcomes*			
Self-reported health	0.271**	0.230**	0.197†
	(0.133)	(0.099)	(0.122)
Health index	0.019**	0.014**	0.013*
	(0.010)	(0.006)	(0.007)
Sample	2014 expansions only	2014 expansions only	2014 expansions only
Health$_{t-2}$ controls		✓	
Health$_{t-3}$ controls			✓
Obs.	8949	8032	6971

Note.—Standard errors clustered on state-wave in parentheses.
† $p < .15$, * $p < .1$, ** $p < .05$, *** $p < .01$

Table 25. All estimated structural parameters

Preferences		Health production	
Parameter	Estimate	Parameter	Estimate
γ	0.593	δ	0.967
θ	-49.03	λ	0.121
A	19.03	α	0.077
Ψ	0.101		
Survival parameters		**Health shocks**	
$\phi_{1,55}, \phi_{2,55}$	-4.54, -4.92	μ_{55}, σ_{55}	-10.68, 3.38
$\phi_{1,60}, \phi_{2,60}$	-1.63, -1.50	μ_{60}, σ_{60}	-9.20, 4.83
$\phi_{1,65}, \phi_{2,65}$	-1.18, -4.53	μ_{65}, σ_{65}	-9.19, 4.97
$\phi_{1,70}, \phi_{2,70}$	-0.74, -5.27	μ_{70}, σ_{70}	-9.12, 5.84
$\phi_{1,75}, \phi_{2,75}$	-0.59, -8.72	μ_{75}, σ_{75}	-9.09, 6.16
$\phi_{1,80}, \phi_{2,80}$	-0.19, -7.40	μ_{80}, σ_{80}	-8.93, 7.05
$\phi_{1,85}, \phi_{2,85}$	-0.16, -5.49	μ_{85}, σ_{85}	-8.61, 7.39
$\phi_{1,90}, \phi_{2,90}$	-0.04, -9.90	μ_{90}, σ_{90}	-8.36, 7.50
$\phi_{1,95}, \phi_{2,95}$	-0.04, -9.90	μ_{95}, σ_{95}	-7.57, 9.08

G Internal Estimation

G.1 Estimated Parameters

Table 25 presents the full set of model parameters. The health shocks evolve as expected, with both the mean shock and the variance of the shock increasing in age. The rate at which the shocks decline is (except for the difference from 55 to 60) roughly increasing in age as well. This suggests that health begins deteriorating more quickly due to health shocks later in life. The growth of the variance also suggests that the dispersion of shocks is also increasing in age, which is intuitive. Therefore, as individuals age, the mean shock gets larger and relatively bad shocks become more likely.

G.2 External Validity

Table 26 re-displays the untargeted model fit from Table 10. The first row is a measure of how adversely selected Medicaid is following the expansions. To generate the data moment, I estimate the regression model in Equation (1), but with additional interaction terms between the triple-difference term and the health index. I then estimate the analog DD model on the simulated data with the additional interaction. The sign of these coefficients suggests that sicker individuals are much more likely to select into Medicaid post-expansion. Given that enrolling in Medicaid is costly, the model's ability to replicate this pattern suggests that the types of individuals who are likely to value Medicaid the most (in terms of health) are consistent across the model and the data. The second row of the table presents an estimate of the relationship between the expenditure response to the Medicaid expansions and health. The data moment

Table 26. Model validation

Untargeted moment	Data	Model	Data source	Description
Effect of Medicaid expansions on mortality	-0.094	-0.082	Miller, Johnson, and Wherry (2021)	Data estimate denotes the relative change in mortality rates in expansion and non-expansion states for individuals aged 55-64. Simulated moment is calculated estimating the relative difference in mortality rates across the treated and non-treated simulated agents aged 62-65.
Adverse selection into Medicaid	-0.065	-0.069	Author's calculations	Data moment estimated from a regression model similar to Equation (1) that includes an interaction term between the TD term and the health index.
Effect of interaction between Medicaid expansions and health on expenditures	-0.200	-0.081	Author's calculations	Data moment estimated from a regression model similar to Equation (1) where predicted log expenditures is the dependent variable and that includes an interaction term between the TD term and the health index.
Effect of interaction between Medicaid expansions and uninsurance on expenditures	0.19	0.10	Table 19	Data moment estimated from a regression model similar to Equation (1) where predicted log expenditures is the dependent variable and that includes an interaction term between the TD term and an indicator for if the individual was previously uninsured.

presented is the coefficient estimate on the triple interaction term from a regression model like that presented in Equation (29), but where the indicator for uninsured status is replaced with the health index. The value of the coefficient estimate suggests that the expenditure response to the expansions is decreasing in health, which is consistent with the idea that sicker individuals have a higher demand for healthcare. The model also reproduces this trend, with a point estimate generated from the simulated data that is of the same sign and similar magnitude.

The next untargeted moment is the coefficient estimate presented in table 19 that describes the relative expansion induced expenditure response for previously uninsured individuals. The positive point estimate suggest that individuals who were uninsured prior to the expansions consume more healthcare than those who had some other form of coverage. The model is able to produce a similar pattern, where the relative expenditure response for individuals who were previously uninsured is larger than simulated agents with prior private insurance. The fact that he model replicates this pattern is encouraging, as The final targeted moment comes from Miller, Johnson, and Wherry (2021), and is the estimated effect that the Medicaid expansions had on mortality rates for individuals aged 55-64. The value of -0.094 reflects that mortality rates decreased in expansions states by 9.4% over the sample mean. I construct an analogous measure of the reduction in mortality relative to mean mortality in the simulated data, which yields a reduction of 7.6% over the mean mortality rate for the control group.

One way in which the model somewhat struggles to match aggregate patterns is in terms of savings. In the data, the asset to income ratio for individuals in the lower income groups is roughly 5.73. However, agents in the simulated data have an asset to income ratio of approximately 7.19. This indicates that individuals over save in the model relative to the data. I believe the primary reason for this is that I do not directly model dual eligibility for Medicare enrollees. While I have a means-tested consumption floor for retirees in the model, I restrict those transfers such that they can only be used for consumption and not medical spending. Further, there is no notion of transfers based on medical need. This leads to agents over saving in order to self-insure against future health risk. Ultimately, the level of self-insurance is, if anything, likely to lead me to *understate* the value of the Medicaid expansions and *overstate* future Medicare expenditures if income effects are large (i.e., individuals with higher savings later in life consume more healthcare, all else equal). Therein, while this is a shortcoming of the model, I do not think it affects the direction of the quantitative findings reported in Section 8.