

Health, families and private and public health insurance*

Francisco Bullano[†]

November 26, 2025

[Click here for most recent version](#)

Abstract

This paper evaluates the welfare consequences of Medicaid expansions targeting low-income children and pregnant women between the late 1980s and 2010 and their effects on private health insurance premiums. By extending public coverage to dependents, these reforms reshaped household insurance decisions and led to significant crowding out of private insurance. Motivated by new evidence on shifts in insurance coverage, family-level dynamics, and medical spending patterns, I develop and estimate a model of household insurance choice that incorporates heterogeneous household composition and health risks, Medicaid eligibility constraints, and endogenous premiums. I find that while these expansions lead to significant crowding out, welfare gains were substantial among low-income households. These gains come from increased consumption and medical treatment among low-income families. I find that the effect on private health insurance prices is small because opposing shifts in the risk profile of enrollees and in the size of insurance plans offset each other in equilibrium.

*I am especially grateful to my advisors Mariacristina De Nardi and Joseph Mullins for their continuous encouragement and support, as well as for very thoughtful discussions. I am also grateful to Alessandra Fogli, Andrew Goodman-Bacon, Marco Bassetto, Anmol Bhandari, Jeremy Lise for very useful discussions and to Johanna Torres Chain for her very useful comments and infinite support. All errors are my own.

[†]bullano083@umn.edu, University of Minnesota, Department of Economics

I Introduction

For more than half a century, U.S. policymakers have pursued broader access to health care and insurance coverage. Signed into law in 1965, Medicaid provides public health insurance to the poor, disabled, and medically needy individuals—groups who lacked access to private coverage. Through successive reforms, Medicaid expanded significantly, becoming one of the largest components of U.S. public spending and a cornerstone of the social safety net. Because the United States relies heavily on private health insurance, understanding how Medicaid affects household behavior, private insurance markets, and welfare is therefore central to public policy.

Between 1986 and 2010, a series of policy changes dramatically expanded Medicaid eligibility for children and pregnant women in low-income families. Because children and pregnant women, if covered, are usually claimed as dependents under a family member’s employer-sponsored plan, these expansions altered household insurance decisions by creating new combinations of public and private coverage within the same household. Also, families could now insure dependents through Medicaid without maintaining private coverage for adults, effectively decoupling dependents’ insurance decisions from those of their parents or spouses. This expansion thus broadened the set of coverage options available to families and introduced new margins along which they could adjust their insurance choices.

The new coverage options substantially changed household behavior. Many individuals who were previously privately insured shifted to Medicaid, leading to substantial crowd-out of private coverage ([Cutler and Gruber, 1996](#)). Because the policy primarily targeted dependents rather than policyholders, it also generated spillover effects: household members ineligible for public insurance—typically parents—were more likely to drop their own coverage and become uninsured ([Bullano, 2023](#)).

While crowd-out and spillover effects are key for understanding the fiscal cost of these expansions and their welfare implications for targeted households, they also could have broader market consequences. If targeted households differ systematically from the privately insured population—particularly in health risk and expected medical spending—the reallocation of coverage can alter the risk composition of the privately insured pool if the expansion is large. Such compositional changes may, in turn, affect equilibrium premiums in private insurance markets.

Evaluating the welfare implications of these expansions therefore requires accounting for both their direct effects on targeted families and their equilibrium consequences for private markets.

This paper examines how the expansion of public health insurance to children and preg-

nant women affects household coverage choices, private insurance markets, and welfare. I first document three facts on Medicaid eligibility, insurance coverage composition, and total healthcare expenditures that motivate a structural model of household insurance decisions and equilibrium pricing in private markets.

First, between 1986 and 2010, Medicaid eligibility expanded substantially for children and pregnant women, while eligibility for adults increased only modestly. Using historical eligibility rules, I construct annual measures of the share of the U.S. population eligible for Medicaid and show that the fraction of eligible children rose sharply during this period, whereas the increase for adults was significantly smaller. As a result, the expansion generated a substantial mismatch in public insurance eligibility within households, making children and pregnant women increasingly eligible for public coverage even when other family members were not.

Second, insurance coverage composition shifted markedly over this period, affecting millions of individuals. Consistent with the expansion of Medicaid eligibility, public coverage rose sharply for children and more modestly for adults, while private insurance declined substantially for both groups. These changes represent a large reallocation of coverage across insurance sources. The magnitude of this shift suggests that the redistribution of enrollees between public and private programs may have altered the risk composition of the privately insured population.

Third, to assess how these changes could affect premiums and welfare among low-income families, I document how total health care expenditures vary across demographic and income dimensions. Along the demographic dimension, the populations targeted by the Medicaid expansions—children and young adults—turn out to have substantially lower total health care expenditures than the average privately insured individual. This suggests that when these relatively low-cost groups exit private coverage, the average cost of those who remain privately insured may rise. Along the income dimension, individuals in low-income families exhibit a higher probability of zero annual medical expenditures than their higher-income counterparts, suggesting that foregoing medical treatment is an important behavioral margin for these households.

Motivated by these facts, I develop a model of household insurance choice and pricing in private markets. I follow [Einav, Finkelstein, Ryan, Schrimpf, and Cullen \(2013a\)](#) and model health insurance decisions in a static model, and incorporate three key extensions. First, I explicitly model that households are heterogeneous in their composition, from single individuals (with and without dependents) to couples, and that each member faces idiosyncratic health risks and has differential access to public insurance. This structure captures the targeted nature of the Medicaid expansions—which primarily extended coverage to children and

pregnant women—and allows household insurance decisions to respond to institutional constraints in the U.S. health insurance system. Second, I model health risk directly rather than medical expenditure risk and allow income to influence treatment behavior. Thus medical spending emerges endogenously from the interaction between health shocks and treatment decisions, a distinction that is essential for capturing income-related behavioral responses among low-income households (Acquatella, 2023; Capatina and Keane, 2025; Ozkan, 2024; Lockwood, 2024). Third, I allow private insurance prices to adjust endogenously, reflecting how individuals sort between public and private coverage. As public insurance expands, selection into private plans changes, altering the risk composition of the privately insured pool and, consequently, equilibrium premiums. This mechanism is critical for capturing feedback effects between public policy and private market dynamics. To the best of my knowledge, no prior study has analyzed the Medicaid expansions for children and pregnant women and their implications for welfare and private insurance pricing.

I estimate the model using data from 1998–2000, a period representing an intermediate stage of the Medicaid expansion—when eligibility for children and pregnant women had already increased substantially but remained well below its 2010 levels. I then use the model to conduct two counterfactual exercises that quantify the effects of changes in eligibility rules.

In the first, I evaluate how the Medicaid expansions up to this intermediate stage affected coverage, premiums, and welfare. Specifically, I simulate an economy in which these groups are subject to the same Medicaid eligibility rules as adults, removing the differential thresholds that favored them. For this exercise, I don’t impose fiscal neutrality. The counterfactual economy exhibits lower public coverage and higher private insurance enrollment, but the composition of these shifts varies across groups. Among children, gains in private coverage do not offset the loss of Medicaid, leading to higher uninsurance. Among adults, the opposite occurs: as families move into private family plans, adult coverage rises and uninsurance falls. Despite large reallocation between public and private insurance, equilibrium adjustments in private premiums remain modest. Premiums in the baseline economy are slightly higher than in the counterfactual without the expansions, but the differences are small. This is because, for family plans, the average cost per insured person is higher in the baseline because fewer children remain privately covered. However, this upward pressure is offset by a decline in the average plan size, as larger families exit private coverage altogether in response to the more generous Medicaid eligibility. The welfare comparison between the two economies reveals substantial heterogeneity across households. The expansion generates large welfare gains for the targeted population, particularly among households below 150% of the federal poverty line (FPL). For these families, gains arise from both higher consumption and greater medical

utilization: Medicaid lowers the effective price of care, reduces out-of-pocket expenses, and eliminates private premiums. Among higher-income households, welfare improvements are smaller and stem mainly from increased disposable income, as these families already received medical care prior to the expansion. For untargeted households, welfare effects are slightly negative, reflecting the modest premium increases associated with the policy.

In the second exercise, I examine to what extent the continued increase in Medicaid eligibility over the 2000s can account for the observed changes in coverage. For this exercise, I simulate a counterfactual economy in which eligibility thresholds are updated to their 2010 levels and recompute the equilibrium. A distinctive feature of this period is that eligibility expanded not only for children and pregnant women but also for adults, bringing adult eligibility thresholds closer to those for children. Here, too, I do not impose fiscal neutrality. The results show that changes in eligibility rules explain about a third of the observed increase in Medicaid enrollment and the decline in private insurance for children, while for adults it capture about two third in the increase in Medicaid, while only 20% of the decline in private coverage. Overall, both adult and children uninsurance rate decline. Because over this period eligibility rules aligned adult thresholds with those for children, about half of the increase in children's Medicaid participation stems from higher adult eligibility—many already-eligible children gained coverage once their parents became eligible. The equilibrium effects on private premiums are quantitatively small. Expanding eligibility for children and pregnant women places modest upward pressure on premiums, but this effect is partly offset by the increased eligibility of adults, who on average crowd out relatively more costly individuals from private coverage. These opposing forces cancel out, leaving only a limited net impact on premiums. Welfare analysis among these two economies reveals, larger welfare gains among targeted households, mostly explained by extending eligibility to adults, which are on average subject to higher health risk. For untargeted households, the welfare implications from higher premiums is negligible. When I impose a revenue-neutral policy and increase the progressivity of the income tax to balance the government budget, aggregate welfare still rises, but the gains become more uneven. Low-income targeted households retain most of their welfare improvements, while higher-income targeted households experience welfare losses as higher taxes offset part of the benefits from Medicaid. Untargeted households also face higher taxes and experience small welfare declines.

The analysis shows that the Medicaid expansions for children and pregnant women generated substantial welfare gains for the low-income families they targeted, while leaving private health insurance premiums largely unchanged. This finding matters because it demonstrates that public insurance can expand coverage and improve welfare without materially distorting selection in private insurance markets. The results indicate that the main trade-off in

expanding public coverage is fiscal: welfare gains for low-income households come at limited cost to non-targeted groups. The model clarifies the mechanism behind this result. By decoupling dependents’ coverage from that of adults, the expansions shifted relatively low-cost members out of private pools, but premiums remained stable because adjustments in plan size and composition offset these selection effects. Taken together, these findings provide a coherent quantitative account of how household-level insurance decisions link public policy to private markets and show that targeted public programs can redistribute effectively at limited cost to the private sector.

This paper contributes to three strands of the literature. First, this paper contributes to the literature on health insurance choices and health risk by emphasizing the role of family-level decision-making and institutional constraints. I extend the standard models of health insurance choice and medical utilization (Cardon and Hendel, 2001; Einav, Finkelstein, Ryan, Schrimpf, and Cullen, 2013a) to incorporate within household heterogeneity for public insurance and allow for multiple insurance sources and income effects, which are key features for low-income families. The model allows for under-treatment among low-income households which is empirically consistent with findings in Capatina and Keane (2025); Acquatella (2023); Lockwood (2024). I integrate these elements into framework that captures price adjustments and behavioral feedback, both of which are essential for evaluating the implications of public insurance expansions.

Second, it adds to the extensive literature analyzing the effects of expanding public health insurance on private health insurance. Cutler and Gruber (1996) first showed that Medicaid expansions in the late 1980s and early 1990s crowded out private insurance. Since then, several studies have examined crowd-out across different periods and datasets.¹ A more recent strand of this literature examines spillover effects on non-eligible household members. Koch (2015) studies the late-1990s and 2000s expansions and finds that increases in children’s Medicaid eligibility reduced parents’ private coverage. Similarly, Bullano (2023) analyzes the late-1980s and early-1990s expansions and finds that expanding Medicaid eligibility for children and pregnant women led some parents to drop their own insurance. This paper provides a unified framework to study both mechanisms and their welfare and equilibrium implications. I model household insurance decisions in the presence of public coverage, explicitly incorporating family linkages and institutional features of the U.S. health insurance system. The results show that while these mechanisms make it more difficult to reduce the uninsured population, they generate sizable welfare gains among low-income households,

¹A non-exhaustive list studying crowd-out includes Dubay and Kenney (1996), Dubay and Kenney (1997), Thorpe and Florence (1998), Blumberg et al. (2000), Yazici and Kaestner (2000), Card and Shore-Sheppard (2004), Ham and Shore-Sheppard (2005), and Shore-Sheppard (2008)

with only modest effects on premiums.

Finally, this paper contributes to the literature examining the large decline in employer-sponsored private insurance coverage over this period. Existing research attributes these trends primarily to rising healthcare costs and insurance premiums (Chernew, Cutler, and Keenan, 2005a,b; Kronick and Gilmer, 1999; Shen and Long, 2006), as well as declining take-up rates among low-income and less educated adults. This paper complements this literature by developing a structural model that links changes in Medicaid eligibility for children and pregnant women—groups typically covered as dependents under private insurance—to household coverage decisions and equilibrium pricing in private markets. The analysis shows that these eligibility expansions account for a sizable share of the decline in private coverage, while having only limited effects on premiums.

The remainder of the paper is structured as follows: Section II presents the motivating facts. Section III describes the model. Section IV outlines the estimation strategy. Section V presents estimation results and model fit. Section VI describes the counterfactual exercise, and Section VII concludes.

II Data, historical changes and facts

My empirical analysis draws on two nationally representative datasets: the National Health Interview Survey (NHIS) and the Household Component of the Medical Expenditure Panel Survey (MEPS-HC), both representative of the U.S. civilian non-institutionalized population. The NHIS provides long-run information on insurance coverage and is used to document trends over the relevant period. The MEPS-HC is a longitudinal sub-sample of the NHIS that follows selected individuals over two years, collecting detailed data on demographic characteristics, medical expenditures, insurance coverage, and health status.

I apply a consistent sample selection across both datasets, focusing on heads of household, their spouses (ages 19–64), and dependent children (ages 0–18). This group excludes individuals 65 and older, who are universally covered by Medicare, and captures nearly 90 percent of the U.S. population under age 65. Because private insurance eligibility for non-working individuals often depends on family ties, this definition aligns with the structure of insurance units. Appendix A provides additional details.

In this section, I document three sets of facts that motivate the model. First, I show that between 1986 and 2010 Medicaid eligibility expanded sharply, with changes concentrated among children and pregnant women. Second, I examine how insurance coverage evolved over the same period, highlighting that the increase in public coverage occurred simultaneously with a pronounced decline in private coverage, implying a substantial recomposition of the

privately insured population. Third, I explore how this shift in coverage composition relates to the underlying cost structure of private insurance by analyzing medical expenditures across demographic and income groups, with particular attention to how treatment decisions vary by family income.

II.1 Expanding public health insurance in the US: 1986 - 2010

Between 1986 and 2010, the U.S. health insurance landscape underwent substantial changes, with a strong increase in Medicaid eligibility for children and pregnant women. Medicaid eligibility is typically determined by comparing family income to individual-specific income thresholds.² Before October 1986, these thresholds were uniform between individuals within a household, that is, children and adults faced the same income thresholds.³ As a result, if one household member was eligible, all members typically were. This changed with the Omnibus Budget Reconciliation Act (OBRA) of 1986 and subsequent legislative reforms, which introduced higher income eligibility thresholds for specific demographic groups, particularly children and pregnant women.⁴ These reforms generated within-household variation in eligibility, different members could face different thresholds based on their demographic characteristics.

Figure 1 illustrates how changes in Medicaid rules affected income eligibility between 1986 and 2010. I compile annual Medicaid eligibility rules by demographic group and calculate the fraction of a simulated population eligible for coverage.⁵ Panel 1a shows the fraction of children (ages 0–18) and the fraction of adults (ages 19–64) eligible. Children’s eligibility increased sharply in the early and late 1990s, rising from under 10% in 1986 to over 40% by 2010. In contrast, adult eligibility expanded more modestly over the same period.

Panel 1b presents Medicaid eligibility by age group across three periods. The solid black line corresponds to 1986–1988, when expansions primarily targeted infants and pregnant women. The dashed black line reflects the 1996–1998 period, just prior to the implementation of the State Children’s Health Insurance Program (SCHIP). The dotted line represents the 2006–2008 period, after SCHIP had been fully implemented and before the Affordable Care Act (ACA). Eligibility for children increased substantially over all this period, as well as for pregnant women. For other adults, eligibility increased during the last decade of the


²There are other eligibility criteria as well, such as asset test criteria, eligibility based on disability, being medically needed, among others.

³Until October 1986, Medicaid eligibility was tied to eligibility for Aid to Families with Dependent Children (AFDC).

⁴Including OBRA 1987–1990, the Medicare Catastrophic Coverage Act of 1989–1990, and the State Children’s Health Insurance Program (SCHIP) of 1997.

⁵See Appendix B for further information.

Figure 1: Medicaid Eligibility: 1986 to 2010




Note: Simulated eligibility using Medicaid income eligibility rules (1986–2010). Sample: Heads of household, spouses/cohabitants (age 18–64), and their dependent children from the 1996–1997 Medical Expenditure Panel Survey. Eligibility is based on income rules for children (0–18), pregnant women, and adults with dependent children. Other Medicaid eligibility categories are excluded. *Left panel:* Share of children (0–18) eligible (solid black line with circles) and share of adults (19–64) eligible (dashed black line). *Right panel:* Share eligible by age group. Solid black line with circles corresponds to 1986–1988; dashed black line with squares corresponds to 1996–1998; dotted black line with triangles corresponds to 2006–2008.

analysis, and this was concentrated among young adults with children who were already eligible. Overall, this period saw a substantial increase in child eligibility and a structural shift in Medicaid eligibility, creating many families in which children and pregnant women qualified for public insurance while parents did not.

While previous graphs focus on changes in eligibility, we now turn to actual Medicaid enrollment and private insurance. Figure 2 presents aggregate trends in individual health insurance coverage over the analysis period, disaggregated by age group (children aged 0–18 and adults aged 19–64) and source of coverage (Medicaid vs. private insurance). Panel 2a shows those in Medicaid, while 2b shows the share in private insurance. Mirroring eligibility, there was a substantial increase in Medicaid coverage among children, from 8% in 1984 to nearly 35% by 2010, coinciding with a sharp decline in private coverage, which fell from nearly 78% to about 60%.

While Medicaid coverage for adults increased over the period, the magnitude of this increase was considerably smaller than for children, consistent with the more limited eligibility expansions targeting adults. During the same time period, the share of adults covered by private insurance declined markedly. It declined from 84% in 1984 to 70% in 2010. The

Figure 2: Health Insurance Coverage: Children (0–18) and Adults (19–64)



Note: Own calculations using the National Health Interview Survey (1984–2010). Selected years reflect insurance source availability. Sample includes heads of household, spouses/cohabitants (ages 19–64), and their dependent children. Left panels: covered by Medicaid. Right panels: covered by private insurance. Vertical dashed lines indicate time periods when legislation changes expanded eligibility for the targeted population.

decline in adult private insurance led to an increase in the uninsurance rate of adults over this period.

Taken together, these trends reveal a large-scale expansion in public insurance eligibility and a corresponding contraction in private coverage. These pronounced changes motivate an analysis of whether shifts in coverage composition affected equilibrium premiums in private insurance markets.

II.2 Medical expenditure, demographic characteristics and income

Understanding how medical expenditures vary across demographic and income groups is central for interpreting both the welfare and market effects of Medicaid expansions. On the one hand, variation in medical spending across demographic groups determines how crowd-out affects the cost structure of private insurance, since shifts in who remains privately insured change the average cost of coverage. On the other hand, differences in medical spending behavior across income groups are key for understanding treatment decisions and welfare gains from expanded coverage.

Medical expenditure and age. To assess how the cost structure varies with demographic characteristics, I examine how individual total health care expenditures change with

age among the privately insured. Total expenditures include both out-of-pocket spending and payments made by private insurers. Figure 3a shows that average annual health care spending rises steeply with age. Children consume on average about \$970 per year, compared with \$2,470 for adults. Among adults, spending increases sharply with age—from roughly \$1,620 for individuals aged 25–34 to about \$4,635 for those aged 55–64. The average across all privately insured individuals is approximately \$2,100 per year. This pattern implies that as crowd-out reduces private coverage among children, the composition of the privately insured pool shifts toward older, higher-cost individuals. Even if individual utilization patterns remain unchanged, such a compositional shift increases the average cost of providing private insurance, potentially exerting upward pressure on premiums.


Medical expenditure and income. Since Medicaid expansions primarily targeted low-income households, it is important to examine how income correlates with medical spending. Figure 3b shows the probability of incurring any medical expenditure among privately insured individuals, splitting the sample into two income groups based on their position within the age-specific income distribution. The probability of having zero expenditures is higher among low-income individuals. These differences are larger among individuals who report being in good health, while the gap narrows considerably for individuals reporting being in poor health.⁶ These patterns indicate that low-income households are more likely to forgo treatment, consistent with higher opportunity costs of care. The smaller gap among individuals in poor health suggest that treatment became more valuable in poor health. This interpretation aligns with [Acquatella \(2023\)](#), who document similar evidence among low-income households. From a modeling perspective, this behavior highlights the importance of allowing for income effects and endogenous treatment choices when assessing the welfare implications of public insurance expansions.

III Model

This section develops an equilibrium model of health insurance motivated by the empirical patterns documented above. The framework builds on the model developed by [Einav, Finkelstein, Ryan, Schrimpf, and Cullen \(2013b\)](#), who study insurance choice and medical utilization under exogenous prices, and extends it in three key ways. First, I explicitly model household composition, allowing for heterogeneous members facing distinct health risks and Medicaid eligibility, consistent with the asymmetric expansions targeting children and pregnant women. Second, I introduce income effects in medical treatment decisions to capture

⁶Appendix D contains these specific disaggregation.

Figure 3: Individual annual total health care expenditures



Note: Own calculations using the Household component of the Medical Expenditure Panel Survey 1998-2000. Sample includes heads of household, spouses/cohabitants (ages 18-64), and their dependent children. *Top left:* Mean medical expenditure by age *Top right:* Probability of zero expenditures by age and household income. 95% confidence interval reported.

the behavior of low-income families, for whom forgoing care is an important adjustment margin. Third, I embed the model in equilibrium, allowing private insurance premiums to adjust endogenously to changes in the composition of enrollees, linking policy-induced coverage shifts to market outcomes and welfare.

III.1 Environment and timing

The economy consists of heterogeneous households that differ in family composition, income, and health characteristics. Each household k is composed of N_k members indexed by i . Members are described by observable demographics $h_i = (a_i, s_i, hs_i)$, where a_i denotes age, s_i sex, and hs_i health status. Household heterogeneity is summarized by

$$X_k = \{\eta_k, y_k, \Theta_k\}, \quad \text{where} \quad \eta_k = (\eta^a, \eta^c), \quad \Theta_k = \{h_i\}_{i=1}^{N_k}.$$

Here, η^a and η^c denote the number of adults and children, respectively, and y_k denotes household income.

Decisions take place in a static two-stage environment. In the first stage, households choose a health insurance contract κ_k^j from the set of contracts $\Omega(X_k, \bar{Y})$ available given their characteristics and the Medicaid eligibility rules $\bar{Y}$. In the second stage, after individual health shocks $\Lambda_k = (\lambda_1, \dots, \lambda_{N_k})$ are realized, households decide how much medical care to

purchase for each member and how much to consume.

III.2 The household problem

Given the environment described above, households choose medical spending and consumption after health shocks are realized and select health-insurance contracts ex ante to maximize expected utility.

Second-stage problem

Given health shocks $\Lambda_k = (\lambda_1, \dots, \lambda_{N_k})$ household second-stage utility is

$$u(c_k, m_k, \Lambda_k) = \frac{1}{1-\gamma} \left(\frac{c_k}{\zeta^{\eta_k}} \right)^{1-\gamma} + \sum_i w(h_i) \left[(m_i - \lambda_i) - \frac{1}{2\omega(h_i)} (m_i - \lambda_i)^2 \right].$$

Following [Cardon and Hendel \(2001\)](#) and [Einav et al. \(2013b\)](#), I model the value of medical treatment as a quadratic function of the distance between medical spending and the health shock. Health shocks reduce utility, and medical treatment offsets this loss. Household chooses from three discrete treatment levels per individual: $m_i \in \{0, \alpha\lambda_i, \lambda_i\}$, representing no treatment, partial treatment, and full treatment. No treatment implies the health shock reduces utility, partial treatment mitigates this loss, and full treatment fully offsets it.

For consumption, I depart from the quasi-linear specification commonly used in the literature and instead adopt a strictly concave utility function that allows for diminishing marginal utility of consumption, governed by parameter γ . Here, ζ^η denotes an equivalence scale that adjusts for household composition.

This formulation allows for the possibility that no treatment (partial treatment) may be optimal, creating a margin of self-insurance via forgone care, particularly relevant for low-income households ([Acquafella, 2023](#)) and uninsured individuals ([Capatina and Keane, 2025](#)).

The household chooses (c_k, m_k) to maximize this utility subject to its budget constraint:

$$\begin{aligned} v(X_k, \Lambda_k, \kappa_k^j) &= \max_{c_k \geq 0, m_k \in M_k(\Lambda_k)} u(c_k, m_k, \Lambda_k) \\ \text{s.t. } c_k + \text{OOP}_k^j(m_k) &\leq Y(y_k - P_k^j; \eta_k) + \text{Tr}(Y(y_k - P_k^j; \eta_k), \text{OOP}_k^j(m_k); \eta_k), \\ M_k(\Lambda_k) &= \{m_k \in R_+^{N_k} : m_i \in \{0, \alpha\lambda_i, \lambda_i\} \forall i\}. \end{aligned} \tag{1}$$

Where $\text{OOP}_k^j(\cdot)$ and P_k^j denote, respectively, the out-of-pocket payment function and the premium under contract j , $Y(\cdot)$ is after-tax income, and $\text{Tr}(\cdot)$ represents transfers that

guarantee a minimum consumption floor.

A health insurance contract directly affects the household's budget constraint in two ways. First, the premium P_k^j reduces disposable income available for consumption and medical spending. Second, the out-of-pocket payment function $\text{OOP}_k^j(\cdot)$ determines the share of medical expenditures paid by the household. By lowering the effective price of treatment, health insurance reduces the marginal cost of medical care and increases the incentive to seek treatment.

First-stage problem

Before observing health shocks, the household selects the insurance contract that maximizes expected utility over second-stage outcomes:

$$V(X_k) = E \left[\max_{j \in \Omega(X_k, \bar{Y})} E_{\Lambda_k} [U(v(X_k, \Lambda_k, \kappa_k^j), \psi_k^j)] + \epsilon_k^j \right], \quad (2)$$

where

$$U(x, \psi) = \frac{((1 - \gamma)x^{1/1-\gamma} + \psi)^{1-\gamma_r}}{1 - \gamma_r}.$$

The function $U(x, \psi)$ is a non-linear aggregator that combines second-stage utility x with contract-specific utility ψ_k^j . It allows the curvature of second-stage utility, governed by γ , to differ from risk aversion over uncertain outcomes, governed by γ_r . The curvature parameter γ governs optimal treatment decisions, while γ_r governs choices over insurance contracts. This assumption decouples the first- and second-stage problems, making the model tractable enough to estimate. The term ψ_k^j captures contract-specific utility that does not depend on medical or non-medical consumption. It includes non-pecuniary costs associated with applying for Medicaid and with combining public and private plans within the same household. Finally, ϵ_k^j is an idiosyncratic taste shock at the contract level.

The household problem implies the policy functions

$$PF_k = \{\pi(\kappa_k^j | X_k), c_k^j, m_k^j, \text{OOP}_k^j, \text{Tr}_k^j\}_{j \in \Omega(X_k, \bar{Y})}, \quad (3)$$

where $\pi(\kappa_k^j | X_k)$ denotes the probability of choosing contract j .

III.3 Health-insurance contracts and feasible choice set

Households obtain health insurance through either private or public sources. Private insurance is offered competitively in two standardized plans: an individual plan that covers

only the policyholder and a family plan that extends coverage to dependents. Both plans are characterized by an out-of-pocket (OOP) cost schedule and a premium.

Public coverage is available through Medicaid, which provides insurance at the individual level to members who meet the income and demographic criteria summarized by the eligibility rules $\bar{Y}$. When a household applies for Medicaid, eligible members receive coverage according to these rules, while the remaining members may remain uninsured or hold private insurance.

Access to private insurance is not universal. Let π_k^{priv} denote the probability that household k has access to private coverage. With probability π_k^{priv} , the household can choose among the available private plans (none, individual, or family coverage); otherwise, no private option is available.

The feasible set of health-insurance contracts is therefore

$$\Omega(X_k, \bar{Y}) = \Delta(X_k) \times MED(X_k, \bar{Y}),$$

where $\Delta(X_k)$ represents the set of private options determined by access probability π_k^{priv} , and $MED(X_k, \bar{Y})$ represents public options (none or Medicaid). Each element $j \in \Omega(X_k, \bar{Y})$ specifies, for every household member, a coverage assignment, an OOP cost function, and a premium P_k^j .

This formulation captures how public and private insurance interact within the household and how variation in access to private coverage shapes overall insurance choices. The detailed mapping of coverage categories, cost aggregation, and contract enumeration is provided in Appendix E.

III.4 Private insurance and government

Private insurers

Private insurers are risk-neutral and operate competitively in segmented markets. Each firm offers only one product—either an individual or a family plan—and risk is pooled among all enrollees within that market. Both plans have exogenous cost-sharing schedules, $OOP^{ind}(\cdot)$ and $OOP^{fam}(\cdot)$. In equilibrium, firms set premiums to cover expected medical costs of their enrollees. The zero-profit conditions are

$$p^q = E[\text{medical spending} - OOP^q(\text{medical spending}) \mid \text{plan } q \text{ chosen}], q \in \{ind, fam\}. \quad (4)$$

The government

The government imposes a progressive income tax and allows private insurance premiums to be deducted from taxable income. Collected revenues finance Medicaid, welfare transfers, and an exogenous level of public expenditure G that does not yield direct utility to households.

Total revenues are the sum across households of expected tax payments net of premium deductions,

$$Rev = \sum_k E[y_k - P_k - Y(y_k - P_k \mid \eta_k)].$$

Government expenditures include expected Medicaid reimbursements and transfer payments,

$$Med = \sum_k E[\text{Medicaid reimbursements}], \quad Wf = \sum_k E[\text{transfer payments}].$$

The government budget constraint requires

$$Rev = G + Med + Wf. \tag{5}$$

III.5 Equilibrium

Let T denote the exogenous environment:

$$T = \{Y(\cdot; \eta), Tr(\cdot; \eta), G, \bar{Y}, OOP, X, \Pi(X)\},$$

where $OOP = \{OOP^{ind}, OOP^{fam}, OOP^{med}, OOP^{unins}\}$ is the set of exogenous out-of-pocket cost functions, $X = \{X_k\}_k$ is the distribution of household characteristics, and $\Pi(X)$ is the probability distribution over this set.

An equilibrium consists of a set of household policy functions and private insurance premiums such that:

1. Given private insurance premiums $p = \{p^{ind}, p^{fam}\}$ and the environment T :
 - The induced set of available contracts $\Omega(X_k, \bar{Y})$ is determined by Medicaid eligibility and private access.
 - The household policy functions defined in (3) solve the household problem in (2).
2. Given the resulting household decisions:
 - Premiums p satisfy the zero-profit conditions for insurers, as in equation (4).

- The government budget constraint holds:
- Market clearing:

$$C + M + G = Y,$$

where C , M , and Y denote aggregate consumption, medical expenditure, and income, respectively.

IV Taking the model to data

This section describes how I map the model to the data. I first present the functional forms and parameters estimated within the model, followed by those fixed and estimated externally. I then describe the estimation procedure and the data used to discipline the model.

Parameters estimated using the model

Household utility

The model includes two sets of preference parameters: those governing preferences over health insurance contracts and those determining preferences over consumption and medical treatment.

In the first stage, households choose among available insurance contracts. I estimate the coefficient of risk aversion, γ_r , separately for single-adult and two-adult households, distinguishing between those with and without children. The contract-specific utility component, Ψ_k^j , captures two non-pecuniary costs: (i) the disutility associated with applying for Medicaid and (ii) the disutility from combining public and private coverage within the same household. I estimate these parameters separately across household types, in the same way as the risk-aversion coefficients.

In the second stage, households choose medical treatment conditional on coverage and health shocks. Treatment decisions depend on the ratio between the marginal utility of consumption and the marginal utility of treatment. To identify the later, I set the curvature of marginal utility of consumption to $\gamma = 2$. Parameters governing the utility of treatment, $w(h_i)$ and $\omega(h_i)$, vary by sex and self-reported health status. Consumption is adjusted for household size using the OECD-modified equivalence scale.⁷

⁷ $\zeta^\eta = 1 + 0.5(\eta^a - 1) + 0.3\eta^c$, where η^a and η^c denote the number of adults and children.

Health shocks

Individuals face zero-inflated log-normal health shocks. For individual i ,

$$\lambda_i = \begin{cases} \exp(z_i) & \text{with probability } \varphi_i, \\ 0 & \text{with probability } 1 - \varphi_i, \end{cases} \quad z_i = \log(\lambda_i) \sim N(\mu_i, \sigma_i^2).$$

The parameters $(\mu_i, \sigma_i, \varphi_i)$ depend on observed characteristics through

$$\begin{aligned} \mu_i &= x_i \beta^\mu, \\ \log(\sigma_i) &= x_i \beta^\sigma, \\ \log\left(\frac{\varphi_i}{1 - \varphi_i}\right) &= x_i \beta^\varphi, \end{aligned}$$

where x_i includes a third-order polynomial in age interacted with sex and health status.

Household face uncertain health shock represented by the vector $\Lambda_k = (\lambda_1, \dots, \lambda_{N_k})$. I assume that individual health shocks are independent, and therefore the joint distribution obtained as the product of the individual distributions.

This approach differs from [Einav et al. \(2013b\)](#) and [Marone and Sabety \(2022\)](#), who model a single household-level shock and therefore abstract from within-household heterogeneity in insurance sources.

Parameters exogenously fixed or estimated outside the model

Access to private health insurance depends on adult's labor supply. For adult i in household k ,

$$\log\left(\frac{\pi_{i,k}^{\text{priv}}}{1 - \pi_{i,k}^{\text{priv}}}\right) = ls_{i,k} \beta^\pi,$$

where $ls_{i,k}$ denotes labor-supply indicators. At the household level,

$$\pi_k^{\text{priv}} = \begin{cases} \pi_{1,k}^{\text{priv}} & \text{if } \eta^a = 1, \\ 1 - (1 - \pi_{1,k}^{\text{priv}})(1 - \pi_{2,k}^{\text{priv}}) & \text{if } \eta^a = 2. \end{cases}$$

Labor supply is treated as exogenous and policy-invariant.

Tax Function, Transfers, Medicaid Rules, and Health Plans

The remaining institutional components are fixed externally. The income tax function follows [Feldstein \(1969\)](#), [Benabou \(2000\)](#), and [Heathcote et al. \(2017\)](#), and is parameterized using the estimates from [Borella et al. \(2023\)](#), which govern both the level and progressivity of effective taxes. Transfers guarantee that disposable income net of medical expenditures does not fall below a minimum consumption floor, set equal to one-half of the federal poverty line, consistent with [Hubbard et al. \(1994\)](#) and [De Nardi et al. \(2010\)](#). Medicaid income eligibility thresholds are taken directly from contemporaneous policy rules, as documented in Appendix B. Out-of-pocket medical expenses are determined by deductible, coinsurance, and maximum payment limits that differ across private insurance, Medicaid, and uninsured households. Following [Capatina and Keane \(2025\)](#), uninsured individuals face a 35 percent probability of being unable to treat their health shocks. Appendix F provides details.

Data and estimation strategy

The model is estimated using the Medical Expenditure Panel Survey (MEPS-HC) for the years 1998–2000. I follow the same sample restrictions described in Section II, but for tractability, I limit the estimation sample to single- and two-adult households with up to three children.⁸

To capture heterogeneity in household composition, specifically, how demographic characteristics and income are distributed within families, I solve the household problem separately for each household in the data. This ensures that the joint distribution of household characteristics and family income in the model matches that observed in the data. Capturing this heterogeneity while essential to represent the interaction between family composition, health risk, and Medicaid eligibility, it also makes the estimation computationally challenging. To address this challenge, I exploit the model’s two-stage structure, which separates insurance choices from medical treatment decisions. The estimation therefore proceeds in two steps.

In the first step, I estimate the parameters governing the distribution of health shocks and preferences over medical treatment. Conditional on insurance coverage, medical expenditures depend only on second-stage treatment decisions. I exploit this separability and estimate $w, \omega, \beta^\varphi, \beta^\mu, \beta^\sigma$ by solving the treatment problem conditional on coverage type and applying the Simulated Method of Moments (SMM). The estimation targets 128 empirical moments: the mean, variance, and probability of zero medical spending by age, sex, and self-reported health for high-income, privately insured individuals, as well as the difference in the extensive

⁸I exclude households that report Medicaid coverage despite having family income above the program’s eligibility thresholds. These cases likely reflect measurement error or specific eligibility criteria not modeled here.

margin of medical expenditures between high- and low-income privately insured groups. Appendix G provides additional details on the estimation procedure and moment selection.

In the second step, I estimate the parameters governing household insurance choices, conditional on the estimated health-shock process and treatment-preference parameters from the first stage. I solve the insurance-choice problem and use maximum likelihood to match the model’s predicted and observed coverage outcomes across contract types. The estimated parameters include household-type-specific risk aversion, the utility cost of Medicaid participation, and the disutility associated with mixing public and private coverage. Appendix H provides further details.

V Model fit and results

The model reproduces several key features of the data. The discussion proceeds in the order of the estimation procedure. First, I compare model-implied medical expenditures with their empirical counterparts. Second, I assess coverage choices and the model’s implications for selection into private versus public insurance and the resulting premiums. Appendix J contains the full set of fit statistics, figures, and diagnostics.

Medical expenditures. The model reproduces the main empirical patterns in medical expenditures. It captures the life-cycle profile of spending and the extensive margin of medical use (Figure 13b). Mean medical expenditures rise with age, and the model generates the widening gap between individuals in good and poor health (Figures 13a and 13c).

I use expenditure moments for Medicaid and uninsured groups, both excluded from estimation, to validate the model’s mechanisms. Mean medical spending for Medicaid recipients increases sharply during adulthood, while the uninsured exhibit substantially lower levels of spending (Figure 13d). The model explains about 70 percent of the difference in adult medical expenditures between privately insured and Medicaid recipients and a similar share of the gap between the privately insured and the uninsured. For Medicaid recipients, higher spending stems from their relatively poorer health status. On the other hand, for the uninsured, it is explained by lower treatment levels, which stems from lower access to treatment and higher effective prices.

The model reproduces the main empirical differences in medical spending by income. On the extensive margin, it successfully generates lower treatment rates among low-income families (Figure 14a). This gap arises mainly from healthy individuals in low-income households who forgo treatment, while treatment rates among those in poor health remain similar across income groups. The model also captures the positive relationship between medical spending

and family income (Figure 14b).

Coverage choices. In the estimation, I require the model to match household-level coverage choices across the six contracts available in the economy. These contracts represent possible combinations of private (single or family) and public (Medicaid) insurance options that households can select (see Section III.3 and Appendix E). The model replicates two salient patterns in the data: the large share of households choosing family coverage rather than individual policies, and the very low incidence of within-household mixing between public and private plans (Figure 9). Most households opt either for full private coverage or exclusively for public/Medicaid insurance.

Individual coverage outcomes stem directly from households' choices over insurance contracts and are typically the outcomes of interest. Thus, I describe the patterns in the data and the model fit. In the data, private insurance is the main source of coverage for both adults and children. However, children are less likely to be privately insured, more likely to be covered by public insurance, and less likely to be uninsured than adults. The model reproduces these same patterns closely, matching both the dominance of private insurance overall and the higher public coverage rates among children (see Figure 10).

I then examine private, Medicaid, and uninsured rates by age groups and by family income. Along the age dimension (see Figure 11), the data reveal distinct patterns. Among children, younger ones are less likely to hold private coverage and more likely to be enrolled in Medicaid. Among adults, younger individuals have lower private coverage and higher uninsurance rates, while older adults show the opposite pattern, greater private coverage and lower uninsurance. The model closely reproduces these age-related differences in insurance coverage. Along the income dimension (see Figure 12), the data reveal clear gradients in coverage. Private insurance rates rise sharply with income, while Medicaid participation and uninsurance both decline. The model reproduces these income-related patterns closely. Three mechanisms drive this relationship. First, labor-supply decisions shape the exogenous access probability π_k^{priv} , reinforcing the inverse relationship between income and private coverage. Second, higher Medicaid eligibility among low-income adults and children reduces private insurance demand, consistent with the negative relationship between eligibility and private coverage. Third, low-income households are less likely to purchase private insurance because they can adjust their out-of-pocket expenditures while uninsured through their treatment choices.

Private health insurance premiums. Sorting into private insurance, both single and family contracts, determines the cost of such plans in the economy. The equilibrium pre-

miums implied by the model for single and family coverage are \$1,729 and \$4,800. Using observed data on total medical expenditures and choices, I calculate premiums using equation 4. I obtain \$1,709 for single coverage and \$4,705 for family coverage.⁹

V.1 Model prediction and estimates from the literature

I also compare selected model implications with empirical estimates from the existing literature.

The model implies an average price elasticity of demand for private health insurance of -0.05 across the full income distribution. These magnitudes are consistent with the empirical evidence that finds that the price elasticity of take-up conditional on offer, is typically small. (Liu and Chollet, 2006) Elasticities are larger among low-income households, reaching -0.19 for families below the federal poverty line and -0.26 for those with incomes between one and two times that threshold. They then decline rapidly with income and approach zero for higher-income households.

Next, I examine the model’s implications for the responsiveness of medical spending to income. The model implies an average income elasticity of 0.05 , below empirical estimates such as Acemoglu et al. (2013), who report an elasticity of about 0.7 . The elasticity varies along the income distribution: for households between one and two times the poverty line, it reaches 0.22 and declines toward zero at higher income levels. This pattern reflects that higher income families already treat nearly all realized health shocks, while lower-income households exhibit a higher marginal propensity to allocate additional resources to medical care. The lower elasticity relative to empirical estimates likely arises because the model focuses on medical expenditures as a response to health shocks, rather than capturing other channels such as health investment (Ozkan, 2024).

Finally, I examine the model’s implied price elasticity of total medical expenditures. The model generates an average elasticity of -0.15 , consistent with empirical estimates of the price responsiveness of medical spending (Aron-Dine et al., 2013). Elasticities vary substantially across the income distribution: for families below the poverty line, the implied elasticity is -0.54 and it increase toward zero for higher-income households. For higher income families, the effect of price is null as they would treat their shocks irrespectively of the cost-sharing structure.

⁹MEPS-IC collects actual data on premiums for employer provided health insurance. During this time period the average premium for single coverage was \$2,183 and for family coverage average \$5,660.

VI Counterfactual exercises

I estimate the model for the 1998–2000 period, which I refer to as the baseline period. This period represents an intermediate stage of the Medicaid expansion: by then, eligibility for children and pregnant women had already increased substantially since 1986 but remained well below 2010 levels. As shown in Figure 1, eligibility continued to broaden after the baseline period—not only for children and pregnant women but also for adults—bringing adult thresholds closer to those for dependents. I use the model to conduct two counterfactual exercises that quantify the effects of changes in Medicaid eligibility rules.

The first counterfactual evaluates how the expansions up to this baseline period affected coverage, premiums, and welfare. I simulate an economy in which children and pregnant women face the same eligibility rules as adults, removing the differential thresholds introduced by earlier reforms.

The second counterfactual examines how the continued expansion of eligibility between the baseline period and 2010 contributed to the observed changes in coverage. I simulate an economy with eligibility thresholds set to their 2010 levels and recompute the equilibrium, isolating the effect of policy-driven eligibility changes from other contemporaneous factors influencing enrollment and premiums.

In both counterfactuals, I focus on how eligibility changes affect equilibrium outcomes without enforcing government revenue neutrality. I keep taxes fixed to isolate the direct effects of the expansions on coverage, prices, and welfare. For the second counterfactual, in a separate subsection, I extend the analysis to a revenue-neutral setting, where I adjust taxes to balance government spending.

VI.1 Going back to 1986: Removing Medicaid eligibility for children and pregnant women

In this exercise, I set the Medicaid eligibility thresholds for children and pregnant women equal to those for adults, thereby removing the dependent-targeted expansions introduced after 1986. Table 4 in Appendix K reports the resulting changes in coverage.

This policy change sharply reduces Medicaid eligibility. Overall, 18.4 percent of children and 1.5 percent of adults lose access to the program (Panel A). The losses concentrate among the groups that previously benefited the most from generous rules. Younger children face the largest declines because the original thresholds were highest for early ages: eligibility drops by about 24 percentage points for children aged 0–5 and by 13 points for those aged 12–18. Among adults, losses occur mainly among pregnant women of childbearing age (19–34) (Panel B).

Low-income families experience nearly all of the eligibility reductions. Households below 200 percent of the federal poverty line (FPL) account for almost the entire decline. In the baseline economy, about 98 percent of children below poverty qualify for Medicaid, but under this counterfactual, only 62 percent remain eligible. Among children in families between 100 and 200 percent of the FPL, eligibility falls by more than 60 percentage points because most of these children had qualified through the earlier expansions. Adults lose eligibility primarily among low-income pregnant women (Panels C and D).

The drop in eligibility produces large changes in coverage. Medicaid enrollment among children falls from 14.5 to 8.2 percent, a 6.3 percentage-point (pp) decline. Private coverage among children rises by 4.7 pp, while their uninsurance rate increases by 1.6 pp. Among adults, private insurance coverage increases by 1.5 pp, driven mainly by adults with children (+3.4 pp) rather than those without (+0.3 pp). For adults with children, the rise in private coverage reduces their uninsurance rate by 1.1 pp, producing an overall 0.4 pp decline in adult uninsurance (Panel A).

The asymmetric change in uninsurance between adults and children aligns with [Bullano \(2023\)](#), who finds that expanding eligibility for children and pregnant women leads non-eligible household members to drop private coverage and become uninsured.

Private health insurance premiums As more children and adults enroll in private insurance, the average cost of private coverage falls slightly. Premiums decline by 0.7% for family coverage (−\$33) and by 0.8% for single coverage (−\$14). For family coverage, two opposing effects drive this change. The average cost per enrolled family decreases by 1.21%, but the average number of individuals covered under each policy rises by 0.52%, partially offsetting the cost reduction. For single coverage, the decline occurs because relatively older mothers leave individual plans and join family coverage with their children, lowering the average cost of single policies.

Consumption, medical expenditures, and welfare implications This section examines how changes in Medicaid eligibility affect household consumption, medical expenditures, and welfare. I distinguish between targeted and untargeted households. Targeted households include those in which at least one member loses Medicaid eligibility in the counterfactual relative to the baseline, while untargeted households experience no change in eligibility.

Among targeted families, average consumption falls by 1.8 percent, medical expenditures decline by 2.1 percent, and welfare drops by 2.8 percent. These averages conceal substantial heterogeneity across income levels and family types. The largest welfare losses occur among families with incomes below 150 percent of the FPL. For households below poverty, medical

expenditures fall by 5.9 percent, consumption declines by 2.6 percent, and welfare decreases by 4.8 percent. For families between 100 and 150 percent of the FPL, the reduction concentrates in consumption, which falls by 4.6 percent, while medical expenditures decline by 4.4 percent, leading to a 5.5 percent welfare loss.

These declines occur because families that lose Medicaid eligibility face higher out-of-pocket costs and must pay private insurance premiums when they switch to private plans. Both effects reduce disposable income and constrain spending on non-medical goods and health care. For higher-income households, the effects are smaller and operate mainly through lower consumption.

Across family types, welfare losses are largest among single pregnant women, who are less likely to be insured. Single adult households with children experience the second-largest declines. Together, these groups face an average reduction in consumption of 3.5 percent, a 5.1 percent drop in medical expenditures, and a 4.5 percent decline in welfare. Among two-adult households, effects are smaller: consumption falls by 1.2 percent, medical spending by 0.6 percent, and welfare by 1.3 percent.

For untargeted households—those whose eligibility does not change—the effects remain small and slightly positive. Modest reductions in private premiums generate minor gains in both consumption and welfare. On average, consumption increases by 0.03 percent, medical expenditures by 0.003 percent, and welfare by 0.036 percent.

Overall, aggregate welfare declines by 0.34 percent relative to the baseline. These results show that the post-1986 Medicaid expansions generated sizable welfare gains for low-income households who benefited directly from higher eligibility, while effects through higher premiums remain limited.

VI.2 Moving forward to 2010: Expanding Medicaid eligibility

In the second counterfactual, I update the eligibility rules for children, pregnant women, and adults to match their 2010 levels. Table 5 in Appendix K reports the resulting coverage effects.

Relative to the baseline, Medicaid eligibility rises by 14.4 percentage points (pp) for children and by 2.7 pp for adults (Panel A). Older children experience relative larger gains in eligibility. Eligibility increases by about 11.7 pp for ages 0–5 and by 15 pp for ages 12–18. Among adults, eligibility expands more evenly across ages, but younger adults (ages 19–24) gain the most, with a 5.9 pp increase (Panel B).

Across the income distribution, most of the gains for children occur among families with incomes above poverty and up to 300% of the FPL. For adults, the expansion mainly benefits

those below the poverty line and those between 100 and 200% of the FPL (Panels C and D).

In the counterfactual economy, Medicaid enrollment increases by 5.8 pp for children and by 2.5 pp for adults. Private coverage declines by 4.4 pp for children and by 1.8 pp for adults. Uninsurance rates fall by 1.4 pp for children and by 0.7 pp for adults. Among children, 2.6 pp of the 5.8 pp increase in enrollment comes from expanded eligibility for children and pregnant women. The remaining 3.2 pp reflects take-up by already-eligible children once their parents also become eligible.¹⁰ In the data, Medicaid enrollment over this time period rises by 16.8 pp for children and by 3.8 pp for adults, while private coverage falls by 9.0 pp and 6.2 pp, respectively.¹¹ The simulated eligibility changes explain roughly 20–60% of these shifts. For children, the model accounts for about 34% of the increase in Medicaid enrollment and 30% of the decline in private coverage. For adults, it explains about 65% of the increase in Medicaid participation and 20% of the reduction in private coverage. In the data, the uninsurance rate for children falls by 2.6 pp, and the model explains roughly half of that decline. For adults, the model cannot reproduce the 5 pp increase in uninsurance observed in the data, as it predicts a 0.7 pp decline. These results suggest that other, non-modeled factors contributed to the rise in uninsurance among adults during this period.

Private health insurance premiums The expansion generates modest equilibrium changes in private insurance premiums. Single-coverage premiums remain unchanged, while family premiums rise by about 0.3 percent (+\$18).

For single coverage, two offsetting effects neutralize each other. Relatively higher-cost men and lower-cost women both leave private coverage at similar rates, leaving premiums largely unchanged. For family coverage, the average cost per enrollee increases by 0.7 percent. However, as larger families exit private insurance, the average number of covered individuals per policy falls by 0.4 percent, offsetting most of the increase in cost. Overall, these equilibrium price effects remain quantitatively small. Between the baseline period and 2010, premiums deflated by the CPI for medical care rose by 46.3 percent for single coverage and 58.5 percent for family coverage, showing that the expansion contributes negligibly to the observed rise in premiums.

Consumption, medical expenditures, and welfare implications As in the previous exercise, I divide households into targeted and untargeted groups. A household is targeted if at least one member experiences a change in Medicaid eligibility in the counterfactual relative to the baseline. Untargeted households show no change in eligibility.

¹⁰I calculate a separate counterfactual that expands eligibility only for children and pregnant women, leaving adults' eligibility unchanged. That experiment increases children's Medicaid enrollment by 2.6 pp.

¹¹Based on MEPS-HC 2010, applying the same sample selection used in the model.

The expansion produces similar qualitative patterns in consumption, medical expenditures, and welfare as before. Among targeted low-income families—those with incomes below 200 percent of the federal poverty line (FPL)—both consumption and medical expenditures rise. These effects reflect two main mechanisms. First, crowd-out of private insurance increases disposable income, allowing families to allocate more resources to consumption. Second, lower effective treatment prices increase medical utilization and reduce out-of-pocket spending.

For relatively higher-income families, the expansion primarily raises consumption through increased disposable income, while effects on medical spending are modest.

Quantitatively, targeted households experience a 1.3 percent increase in consumption, a 2.2 percent rise in medical expenditures, and a 4.1 percent gain in welfare. These effects vary across family types. For single-adult households with children, consumption increases by 3.2 percent, medical expenditures rise by 2.9 percent, and welfare increases by 5.9 percent. For two-adult households with children, the gains are smaller: consumption rises by 0.8 percent, medical expenditures by 0.7 percent and welfare increase by 1.5 percent. As adult eligibility also increase, households without children also experience large welfare gains. For childless single-adult I find welfare gains of 5.4 percent, mostly driven by increased medical consumption.

For untargeted households—those whose Medicaid eligibility remains unchanged—the effects are negligible, as private insurance premiums change very little. At the aggregate level, overall welfare rises by 1.21 percent relative to the baseline economy.

Revenue-neutral policy Finally, I compute the new equilibrium assuming that the expansion is fully financed through higher taxes. I adjust the progressivity parameter of the income tax function so that the government budget constraint holds. Coverage outcomes and premiums remain virtually unchanged, so I omit them here. The expansion increases total Medicaid spending by 25.5 percent.

Financing the expansion with higher taxes makes welfare effects more heterogeneous. Untargeted households face higher tax liabilities and experience an average welfare loss of 0.24 percent. Among targeted households, the higher tax burden reduces the welfare gains from the expansion to 3.92 percent. The reduction is particularly pronounced for higher-income targeted families (above 200 percent of the FPL), for whom the tax increase offsets much of the direct welfare improvement.

Overall, welfare declines for high-income households, reflecting the additional tax burden, while low-income households retain most of their previous gains. Aggregate welfare rises by 0.98 percent, but now with meaningful trade-off as taxes are higher.

VI.3 Discussion

Across both counterfactuals, the expansions in Medicaid eligibility generates large shifts in coverage composition but only modest changes in private insurance premiums. Comparing the two exercises, the fraction of children enrolled in Medicaid increase by 12.1 percentage points and adult enrollment by 3.7 percentage points, private insurance declines by 9.1 percentage point for children and 3.3 pp for adults, while private premiums increase by only 0.8 percent for single coverage and 1.0 percent for family coverage. For family plans, the average cost per insured individual rises by 2.1 percent, but this increase is largely offset by a reduction in average plan size as larger families exit private coverage.

The results indicate that Medicaid expansions generate significant welfare gains for low-income households while exerting limited pressure on private premiums. This pattern suggests that the central trade-off of these policies is fiscal: Among families that are not targeted, the effect of this increase in premiums is small, while the effect from higher taxes could have larger implications as shown in the second counterfactual.

At the same time, the analysis abstracts from other important channels through which public insurance expansions may affect welfare. The model does not incorporate potential long-term gains from improved health outcomes, lower mortality, or the accumulation of health capital—effects that would likely magnify the welfare impact of the reforms.

The model also assumes a simple premium-setting mechanism tied to average medical costs. In practice, insurers and employers may adjust prices strategically in response to public insurance expansions—for example, by shifting costs to workers through higher cost sharing rules. Accounting for such employer or insurer responses could implied different equilibrium price effects and represents a promising direction for future research.

VII Conclusion

This paper studies how public health insurance expansions affect household coverage choices, private insurance markets, and welfare. Focusing on the Medicaid expansions for children and pregnant women between 1986 and 2010, I document large eligibility increases concentrated among dependents and a simultaneous decline in private coverage. I show that medical spending differs systematically across demographic and income groups. Children and young adults have lower expected expenditures and low-income households have higher likelihood of foregoing medical treatment. Both are important features for understanding both the direct and equilibrium effects of expansion of public health insurance.

I then develop and estimate a model of household insurance choice where individuals

face idiosyncratic health shocks and household members differ in their eligibility for public insurance. Families make joint coverage decisions, and private premiums adjust endogenously to shifts in the composition of the privately insured pool.

Counterfactual exercises reveal that expanding Medicaid eligibility generates large welfare gains among targeted low-income households. Welfare improvements arise from both higher consumption and greater medical utilization. Access to Medicaid lowers the effective price of care, reduces out-of-pocket expenses, and eliminates private premium payments. Among higher-income families, welfare gains are primarily driven by higher consumption, as these households already received medical treatment prior to the expansion.

The expansion induces large shifts from private to public coverage but has only modest effects on premiums, typically below one percent. For family plans, the average cost per insured individual increases above that threshold, but this effect is offset by a decline in the average number of covered members per policy.

Finally, I find that changes in eligibility rules can explain roughly 20–60% of the observed increase in Medicaid enrollment and decline in private coverage between the late 1990s and 2010, with limited impacts on premiums. About half of the increase in children’s Medicaid enrollment stems from the rise in adult eligibility, as many already-eligible children gained coverage once their parents became newly eligible.

References

- Daron Acemoglu, Amy Finkelstein, and Matthew J Notowidigdo. Income and health spending: Evidence from oil price shocks. *Review of Economics and Statistics*, 95(4):1079–1095, 2013.
- Angélique Acquatella. The demand elasticity of health care spending for low-income individuals. 2023.
- Naoki Aizawa and Chao Fu. Interaction of the labor market and the health insurance system: employer-sponsored, individual, and public insurance. 2020.
- Aviva Aron-Dine, Liran Einav, and Amy Finkelstein. The rand health insurance experiment, three decades later. *Journal of Economic Perspectives*, 27(1):197–222, 2013.
- Roland Benabou. Unequal societies: Income distribution and the social contract. *American Economic Review*, 91(1):96–129, 2000.
- Linda J Blumberg, Lisa Dubay, and Stephen A Norton. Did the medicaid expansions for children displace private insurance? an analysis using the sipp. *Journal of health economics*, 19(1):33–60, 2000.
- Margherita Borella, Mariacristina De Nardi, Michael Pak, Nicolo Russo, and Fang Yang. Fbbva lecture 2023. the importance of modeling income taxes over time: U.s. reforms and outcomes. *Journal of the European Economic Association*, 21(6):2237–2286, 09 2023. ISSN 1542-4766. doi: 10.1093/jeea/jvad053. URL <https://doi.org/10.1093/jeea/jvad053>.
- David W Brown, Amanda E Kowalski, and Ithai Z Lurie. Long-term impacts of childhood medicaid expansions on outcomes in adulthood. *The Review of Economic Studies*, 87(2):792–821, 2020.
- Francisco Bullano. Medicaid expansion and parental health insurance. 2023. URL https://franciscobullano.com/files/ME_PHI_2025.pdf.
- Elena Capatina and Michael Keane. Health shocks, health insurance, human capital, and the dynamics of earnings and health. Technical report, Federal Reserve Bank of Minneapolis, 2025.
- David Card and Lara D Shore-Sheppard. Using discontinuous eligibility rules to identify the effects of the federal medicaid expansions on low-income children. *Review of Economics and Statistics*, 86(3):752–766, 2004.

- James H Cardon and Igal Hendel. Asymmetric information in health insurance: evidence from the national medical expenditure survey. *RAND Journal of Economics*, pages 408–427, 2001.
- Michael Chernew, David Cutler, and Patricia Seliger Keenan. Charity care, risk pooling, and the decline in private health insurance. *American Economic Review*, 95(2):209–213, 2005a.
- Michael Chernew, David M Cutler, and Patricia Seliger Keenan. Increasing health insurance costs and the decline in insurance coverage. *Health services research*, 40(4):1021–1039, 2005b.
- David M Cutler and Jonathan Gruber. Does public insurance crowd out private insurance? *The Quarterly Journal of Economics*, 111(2):391–430, 1996.
- Mariacristina De Nardi, Eric French, and John B Jones. Why do the elderly save? the role of medical expenses. *Journal of political economy*, 118(1):39–75, 2010.
- Lisa Dubay and Genevieve M Kenney. Did medicaid expansions for pregnant women crowd out private coverage? *Health affairs*, 16(1):185–193, 1997.
- Lisa C Dubay and Genevieve M Kenney. The effects of medicaid expansions on insurance coverage of children. *The Future of Children*, pages 152–161, 1996.
- Liran Einav, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen. Selection on moral hazard in health insurance. *American Economic Review*, 103(1):178–219, February 2013a. doi: 10.1257/aer.103.1.178. URL <https://www.aeaweb.org/articles?id=10.1257/aer.103.1.178>.
- Liran Einav, Amy Finkelstein, Stephen P Ryan, Paul Schrimpf, and Mark R Cullen. Selection on moral hazard in health insurance. *American Economic Review*, 103(1):178–219, 2013b.
- Martin S Feldstein. The effects of taxation on risk taking. *Journal of Political Economy*, 77(5):755–764, 1969.
- John C Ham and Lara Shore-Sheppard. The effect of medicaid expansions for low-income children on medicaid participation and private insurance coverage: evidence from the sipp. *Journal of Public Economics*, 89(1):57–83, 2005.
- Jonathan Heathcote, Kjetil Storesletten, and Giovanni L Violante. Optimal tax progressivity: An analytical framework. *The Quarterly Journal of Economics*, 132(4):1693–1754, 2017.

- Kate Ho and Robin S Lee. Health insurance menu design for large employers. *The RAND Journal of Economics*, 54(4):598–637, 2023.
- R Glenn Hubbard, Jonathan Skinner, and Stephen P Zeldes. The importance of precautionary motives in explaining individual and aggregate saving. In *Carnegie-Rochester conference series on public policy*, volume 40, pages 59–125. Elsevier, 1994.
- Thomas G Koch. All internal in the family?: Measuring spillovers from public health insurance. *Journal of Human Resources*, 50(4):959–979, 2015.
- Richard Kronick and Todd Gilmer. Explaining the decline in health insurance coverage, 1979-1995: Rising health spending levels over the past two decades have created a growing pool of uninsured workers. *Health affairs*, 18(2):30–47, 1999.
- Su Liu and Deborah Chollet. Price and income elasticity of the demand for health insurance and health care services: a critical review of the literature. 2006.
- Lee M Lockwood. Health insurance and consumption risk. Technical report, Working Paper, 2024.
- Victoria R Marone and Adrienne Sabety. When should there be vertical choice in health insurance markets? *American Economic Review*, 112(1):304–342, 2022.
- Serdar Ozkan. Income differences and health disparities: Roles of preventive vs. curative medicine. Technical report, Federal Reserve Bank of St. Louis, 2024.
- Nicolo Russo. Health-dependent preferences, consumption, and insurance. Mimeo, November 2022. URL https://nicolorusso.github.io/jmp_current.pdf.
- Yu-Chu Shen and Sharon K Long. What’s driving the downward trend in employer-sponsored health insurance? *Health services research*, 41(6):2074–2096, 2006.
- Lara D Shore-Sheppard. Stemming the tide? the effect of expanding medicaid eligibility on health insurance coverage. *The BE Journal of Economic Analysis & Policy*, 8(2), 2008.
- Kenneth E Thorpe and Curtis S Florence. Health insurance among children: the role of expanded medicaid coverage. *Inquiry*, pages 369–379, 1998.
- Esel Y Yazici and Robert Kaestner. Medicaid expansions and the crowding out of private health insurance among children. *Inquiry*, pages 23–32, 2000.

Appendices for Online Publication

A The NHIS and the MEPS-HC: Sample selection

To ensure consistency across datasets, I apply a harmonized sample selection to both the NHIS and the MEPS-HC. Because health insurance eligibility often depends on family structure—which can differ from the household definitions used in these surveys—I construct households based on the following rules:

I first identify the head of household within each reporting unit and assign to that unit any spouse or dependent children under age 19, based on reported relationship variables. I retain households where both the head and spouse (if present) are between ages 19 and 64. I exclude extended relatives (e.g., parents, siblings, and other adults) and children under 19 who cannot be claimed as dependents. Cohabitant partner and their dependent children are assigned to their own household. Additionally, I drop households with resident grandchildren, as parent-child links cannot be reliably established to construct their own household units. Children older than 18 are assigned to their own household.

This selection yields a representative sample covering approximately 90% of the U.S. non-elderly population. Table 1 presents this data selection and its representativeness for the entire sample for the NHIS.

A.1 The MEPS-HC

The MEPS-HC has rich, high-quality information on health care consumption, as well as information on household demographics and income. The richness of this data is that it contains detailed information about individual medical expenditures, as well as information about health status for each member. It also contains information about insurance coverage source. My baseline measures of medical spending are inclusive measures of the types of services typically covered by health insurance, including hospital care, doctor visits, and prescriptions. I use the variable TOTEXP that measures the actual payments made for medical services, including from the insurer and the household. My income measures include income from all sources, including social insurance and means-tested programs, to reflect the net risk after such insurance. All monetary variables are converted to real 2009 dollars. Throughout, I use household weights to ensure that the estimates reflect the U.S. population.

Table 1: Sample Coverage and Age Statistics (NHIS)

Year	Age 0–17			Age 18–64		
	In Sample	Avg. Age In	Avg. Age All	In Sample	Avg. Age In	Avg. Age All
1984	0.89	8.59	8.54	0.89	38.2	37.6
1986	0.90	8.52	8.45	0.89	38.3	37.7
1989	0.90	8.34	8.26	0.90	38.3	37.9
1990	0.89	8.27	8.21	0.89	38.4	38.0
1991	0.81	8.26	8.20	0.84	38.8	38.2
1992	0.85	8.27	8.22	0.87	38.8	38.3
1994	0.89	8.36	8.30	0.89	38.8	38.3
1995	0.89	8.43	8.36	0.89	39.0	38.5
1996	0.89	8.48	8.42	0.89	39.1	38.7
1997	0.91	8.51	8.47	0.91	39.1	38.8
1998	0.91	8.50	8.48	0.91	39.2	39.0
1999	0.91	8.54	8.52	0.91	39.4	39.1
2000	0.90	8.53	8.54	0.90	39.5	39.3
2001	0.91	8.56	8.56	0.91	39.7	39.4
2003	0.91	8.59	8.60	0.90	40.0	39.7
2004	0.90	8.60	8.62	0.90	40.2	39.9
2005	0.91	8.59	8.60	0.90	40.3	40.0
2006	0.90	8.58	8.58	0.90	40.4	40.1
2007	0.90	8.56	8.57	0.90	40.5	40.2
2008	0.90	8.49	8.49	0.89	40.6	40.3
2009	0.89	8.45	8.44	0.89	40.6	40.4
2010	0.89	8.39	8.37	0.88	40.8	40.5

Notes: "In Sample" indicates the proportion of individuals in each age group included in the analysis. "Avg. Age In" is the average age in the sample; "Avg.Age All" is the average age in the full population average for that group.

B Medicaid expansions and eligibility

Before 1986, Medicaid primarily served the aged, disabled, medically needy, and poor families with dependent children. However, policy shifts in the late 1980s and early 1990s significantly altered eligibility rules for this last group. In particular, these changes focused on expanding coverage to children from low-income families and pregnant women.

In this section, I describe the pathways of Medicaid eligibility during my study period. The main sources of information are the BLK Medicaid Calculator ([Brown et al. \(2020\)](#)) documentation and source files, supplemented with additional historical data extending back to 1985.

Using these files, I construct monthly income eligibility thresholds by age (0-64) and pregnancy status, covering both children and adults. For non-targeted individuals, whose eligibility typically follows AFDC guidelines, I apply Medicaid eligibility rules for a family of three.

Since eligibility criteria existed prior to Medicaid's expansion, I started my analysis in 1985 when computing eligibility measures. This allows me to assess eligibility dynamics at least two years before changes in eligibility driven by this expansion.

B.1 Medicaid expansion in the late '80s and early '90s

The Omnibus Budget Reconciliation Act (OBRA) of 1986 was the first legislation during my analysis period to significantly change Medicaid eligibility. Before its enactment, Medicaid eligibility was attached to the State Aid to Families with Dependent Children (AFDC) eligibility. Beginning in April 1987, OBRA 1986 allowed states to increase income eligibility thresholds above AFDC levels to a maximum of 100% of the Federal Poverty Line (FPL) for pregnant women, infants, and children up to 5 years old. The expansion for children was implemented gradually through a phased-in process. Under OBRA 1986, the oldest cohort covered included children born on October 1, 1985.

Subsequent legislative changes, including OBRA 1987, the Medicare Catastrophic Care Amendments (MCCA) of 1988, OBRA 1989, and OBRA 1990, gradually expanded public coverage eligibility. OBRA 1987 extended states' optional authority to raise income thresholds for pregnant women and infants up to 185% of the FPL and accelerated OBRA 1986's phase-in timeline by October 1988.¹² Additionally, OBRA 1987 allowed states to increase income thresholds for children up to age eight to 100% of the FPL, with implementation

¹²States that exercised this option expedited the phase-in process, making children born on October 1, 1983, the oldest cohort covered.

following a phased-in approach.¹³

The MCCA of 1988 mandated a minimum coverage level for pregnant women and infants at 100% of the FPL. This expansion was originally designed for a two-year phase-in, covering those below 75% of the FPL by July 1989 and reaching full implementation by July 1990. However, OBRA 1989 accelerated this timeline and further expanded the minimum federal coverage guarantee. By April 1990, all states were required to cover, at a minimum, pregnant women and children up to six years old below 133% of the FPL. Lastly, OBRA 1990 mandated that states extend coverage to children up to age 19 who were born on or after October 1, 1983, with family incomes below 100% of the FPL, effective July 1991.

B.2 Medicaid expansion in the late 90 and early 2000

In October 1997, the federal government established the State Children’s Health Insurance Program (SCHIP) as part of the Balanced Budget Act of 1997. The purpose of SCHIP was to give states the option to expand coverage to uninsured low-income children in families with incomes above Medicaid eligibility thresholds but below 200% of the federal poverty level.

B.3 Mapping income eligibility rules to eligibility

To isolate policy-driven changes in Medicaid eligibility, I rely on simulated Medicaid eligibility, which maps the complexity of Medicaid eligibility rules into a simplified indicator of eligibility at national level.

Specifically, using state- and time-specific Medicaid income thresholds by age and pregnancy status, I evaluate whether individuals in my simulated sample satisfies income requirement for Medicaid eligibility. Then, I compute relevant measure of eligibility given those rules. Importantly, I apply those rules to the 1997-1998 MEPS-HC population, where I observe demographic characteristic and income. Since the public files of the MEPS-HC do not have state of residence, but rather region of residence, I replicate each household N_R times, and re-weight them based on the unconditional probability of belonging to state s in region R . I then aggregate those measure at national level. With the simulated sample in hand, I then compute the different measures of eligibility display in the main text. This procedure ensures that changes in simulated eligibility reflect changes in rules, rather than changes in demographic or economic outcomes.

¹³Children had to be born on or after October 1, 1983, to qualify.

C Household level health insurance contracts


Changes in coverage for adults and children have rich households dynamics. While trends in individual coverage outcomes, these patterns reflect underlying changes in household-level health insurance decisions. To capture this, I classify each household into one of six mutually exclusive health insurance contracts, defined by the combination of private and public coverage sources across household members. I begin by defining three mutually exclusive private insurance categories for adults: (1) no private coverage, (2) single coverage, and (3) family coverage. These categories carry different implications for children. In cases (1) and (2), children do not receive private coverage through the household, while in case (3), children are typically included under a family policy. In households with multiple adults, single coverage indicates that only one adult is privately insured.

These private coverage choices interact with public insurance through Medicaid. Specifically, any household member who is not privately insured may obtain coverage through Medicaid if eligible. This interaction gives rise to six distinct household insurance contracts: (a) No private – No public, (b) No private – Any public, (c) Single coverage – No public, (d) Single coverage – Any public, (e) Family coverage – No public, (f) Family coverage – Any public.

Figure 4 presents the evolution of these household health insurance contracts over time. Panel 4a shows the share of individuals living in households where no adult has private coverage. The red line represents the share in which at least one household member is covered by Medicaid, while the black line represents the share where no household member is insured. Panels 4b and 4c show the analogous trends for households with single private coverage and family private coverage, respectively.

Four novel patterns emerge from this exercise. First, the documented decline in private insurance coverage over this period was driven by a reduction in family coverage, rather than single adult coverage. Second, the fraction of individuals living in households with single coverage has shown an upward sloping trend, rather than a decline. Third, the rise in individuals living in households without private insurance is explained by an increase in Medicaid coverage among at least one household member (red line), while the share of individuals living in households with neither private nor public coverage has remained relatively stable over time. Fourth, mix-contracts, defined by differing insurance sources within households, have become more common.

Figure 4: Household health insurance contracts



Note: Own calculations using the National Health Interview Survey. Sample includes heads of household, spouses/cohabitants (ages 18–64), and their dependent children. Panels report the fraction of households with “No private coverage,” “Single coverage,” and “Family coverage.” Red lines depict contracts where at least one additional member is publicly insured. Black lines depict contracts where no other member is insured.

D Medical expenditures

The Household Component of the Medical Expenditure Panel Survey (MEPS-HC) provides a rich set of variables on medical charges and expenditures.

Total Health Care Charges (TOTTCCH) captures the total amount billed for all reported medical events. These charges reflect list prices and do not necessarily correspond to actual payments. Prescription drug charges are excluded.

Total Health Care Expenditures (TOTEXP) measure the actual payments made for medical services, including out-of-pocket spending and payments from private insurance, Medicaid, Medicare, and other third-party payers. This measure includes prescription drugs, but not over-the-counter medications.

Expenditures on Prescribed Medicines (RXEXP) reports total payments—both out-of-pocket and third-party—for prescription medications. Charge-based data for prescriptions are not available.

D.1 Medical expenditure and income

This section documents income gradients in the medical expenditure profiles of high- and low-income individuals, focusing on both extensive and intensive margins of care utilization.

To avoid selection concerns, the analysis is restricted to privately insured individuals. Since health shock processes are modeled independently of household insurance choices, I use these moments for estimating health-related parameters in the model.

D.1.1 Extensive margin

Zero medical expenditures and charges are common in the MEPS data. Moreover, this is common across other databases reporting medical expenditure (Marone and Sabety, 2022; Einav et al., 2013a; Ho and Lee, 2023). In this paper, I treat these zeros as reflecting the absence of any healthcare use during the reporting period. Let D_i be a binary variable indicating whether individual i has zero charges:

$$\log \left(\frac{\Pr(D_i = 0 \mid X_i)}{1 - \Pr(D_i = 0 \mid X_i)} \right) = X_i \beta \quad (6)$$


where X_i includes a third-order polynomial in age, fully interacted with sex and self-reported health status.

Individuals are classified as high- or low-income based on whether their family income (adjusted for household composition) falls above or below the median within their age group and survey year. Age groups are defined as: 0–5, 6–11, 12–17, 18–24, 25–34, 35–44, 45–54, and 55–64. Self-reported health status is grouped into two categories, with a binary indicator for bad health defined as reporting either “poor” or “fair” health. I estimate equation (6) separately by income group and compute the difference in predicted probabilities across the two groups.

Figure 5 shows the level of predicted probabilities, while Figure 6 shows the high-low


income gap. As illustrated, significant income-based differences in the likelihood of incurring any medical expenditure persist, even among the privately insured. This suggests that low-income individuals are more likely to forgo care, highlighting the presence of an extensive margin in healthcare use decisions.

Figure 5: Probability of any medical expenditure (Privately insured)



Note: Own calculations using the Household Component of the Medical Expenditure Panel Survey (MEPS). The sample includes heads of households, spouses/cohabitants (ages 18–64), and their dependent children. *Top left:* Men in good health. *Top right:* Men in bad health. *Bottom left:* Women in good health. *Bottom right:* Women in bad health. The red line corresponds to low-income households, while the black line corresponds to high-income households.

Figure 6: Gap in probability of any medical expenditure: High vs. low income (Privately insured)



Note: Own calculations using the Household Component of the Medical Expenditure Panel Survey (MEPS). The sample includes heads of households, spouses/cohabitants (ages 18–64), and their dependent children. *Top left:* Men in good health. *Top right:* Men in bad health. *Bottom left:* Women in good health. *Bottom right:* Women in bad health. 95% confidence interval depicted.

D.1.2 Expected medical consumption

I now turn to differences in expected medical expenditures across income groups. Let m_i denote individual medical spending. I estimate the following linear model:


$$E[m_i|X_i] = X_i\beta \quad (7)$$

where X_i includes a third-order polynomial in age, fully interacted with sex and self-reported health status. I estimate equation (7) separately for high- and low-income individ-

uals and compute the difference in predicted values to assess income-based gaps in expected spending.


Figures 7 and 8 present the results. The first panel shows average medical expenditures by age and income group. The second panel shows the estimated income gap. As shown, three of the four demographic groups exhibit significant differences in expected spending. On average, privately insured low-income individuals incur lower medical expenditures, indicating that they are less costly to insure.

Figure 7: Expected medical expenditure by age and income (Privately insured)



Note: Own calculations using the Household Component of the Medical Expenditure Panel Survey (MEPS). The sample includes heads of households, spouses/cohabitants (ages 18–64), and their dependent children. *Top left:* Men in good health. *Top right:* Men in bad health. *Bottom left:* Women in good health. *Bottom right:* Women in bad health. The red line corresponds to low-income households, while the black line corresponds to high-income households.

Figure 8: Gap in expected medical expenditure: High vs. low income (Privately insured)



Note: Own calculations using the Household Component of the Medical Expenditure Panel Survey (MEPS). The sample includes heads of households, spouses/cohabitants (ages 18–64), and their dependent children. *Top left:* Men in good health. *Top right:* Men in bad health. *Bottom left:* Women in good health. *Bottom right:* Women in bad health. 95% confidence interval depicted.

E Model details

E.1 Health insurance contracts and the role of Medicaid

A health insurance contract for household k is a household-specific object that determines which members are covered, the applicable out-of-pocket (OOP) cost schedule, and the total premium paid. Formally, a contract is defined as:

$$\kappa_k^j = \left\{ \{e_i^j\}_{i=1}^{N_k}, \text{OOP}_k^j(\cdot), P_k^j \right\},$$

where e_i^j denotes the source of coverage for individual i , $\text{OOP}_k^j(\cdot)$ is the out-of-pocket cost function implied by contract j , and P_k^j is the total premium paid by the household.

Contracts arise from the interaction between public and private health insurance plans, as determined by household characteristics and policy rules. These plans are defined independently of the household and are described below.

Private insurance plans: Private insurance is offered by competitive firms in the form of two standardized plans: individual coverage and family coverage. A private plan is denoted:

$$\delta^q = (\text{OOP}^q(\cdot), p^q), \quad q \in \{ind, fam\},$$

where $\text{OOP}^q(\cdot)$ is the plan-specific out-of-pocket cost function, and p^q is the premium. Individual plans ($q = ind$) cover only the policyholder, while family plans ($q = fam$) cover the policyholder and their dependents (e.g., spouse or children).

These plans interact with household composition to determine who is covered. Let $\delta_k^q = \delta^q \oplus X_k$ denote the implementation of plan q for household k , where the operator $\oplus$ maps household characteristics to a coverage assignment:

$$\delta_k^q = \left\{ \{e_i^q\}_{i=1}^{N_k}, \text{OOP}^q(\cdot), p^q \right\},$$

with $e_i^q \in \{0, 1, 2\}$ indicating whether individual i is not privately covered, covered by individual private insurance, or covered by family private insurance, respectively.

Access to private insurance occurs with probability $P(X_k)$, and the available private options are:

$$\Delta(X_k) = \begin{cases} \{\text{no private, } ind, fam\}, & \text{with probability } \pi_k^{priv} \\ \{\text{no private}\}, & \text{with probability } 1 - \pi_k^{priv} \end{cases}$$

where households can purchase at most one private plan.

Public insurance plans (Medicaid): The government offers a public plan $\delta^{med} = (\text{OOP}^{med}(\cdot), p^{med})$, which operates at the individual level. Eligibility is determined by household income and individual demographic characteristics, summarized by the rule set $\bar{Y}$. If the household chooses to apply for Medicaid, the implementation of the public plan is:

$$\delta_{k, \bar{Y}}^{med} = \delta^{med} \oplus^{med} \{X_k, \bar{Y}\} = \left\{ \{e_i^{med}\}_{i=1}^{N_k}, \text{OOP}^{med}(\cdot), p_k^{med} \right\},$$

where $e_i^{med} \in \{0, 1\}$ indicates whether individual i is covered by Medicaid. The set of public options is:

$$MED(X_k, \bar{Y}) = \{\text{no public, } med\}.$$

Available contracts: The set of all health insurance contracts available to household k is:

$$\Omega(X_k, \bar{Y}) = \Delta(X_k) \times MED(X_k, \bar{Y}),$$

and each element $j = (q, med) \in \Omega(X_k, \bar{Y})$ defines a contract κ_k^j as described above and emerges from the interaction between element $q \in \Delta(X_k)$ and $med \in MED(X_k, \bar{Y})$. The resulting contract is defined as:

$$e_i^j = \begin{cases} 0, & \text{if } e_i^q = 0 \text{ and } e_i^{med} = 0, \\ 1, & \text{if } e_i^q = 1 \text{ and } e_i^{med} = 0, \\ 2, & \text{if } e_i^q = 2 \text{ and } e_i^{med} = 0, \\ 3, & \text{if } e_i^{med} = 1. \end{cases}$$

where each household member is assigned to one of four mutually exclusive coverage categories where 0 = Uninsured, 1 = Private individual, 2 = Private family, and 3 = Medicaid. The household's out-of-pocket costs under contract j and medical expenditure $m_k = (m_1, \dots, m_{N_k})$,

$$\begin{aligned} \text{OOP}_k^j(m_k) = & \sum_{i: e_i^j=1} \text{OOP}^{ind}(m_i) + \text{OOP}^{fam} \left(\sum_{i: e_i^j=2} m_i \right) \\ & + \sum_{i: e_i^j=3} \text{OOP}^{med}(m_i) + \sum_{i: e_i^j=0} \text{OOP}^{unins}(m_i). \end{aligned}$$

The household's total premium under contract j is:

$$P_k^j = p^{ind} \cdot n_{k,ind}^j + p^{fam} \cdot 1\{n_{k,fam}^j > 0\} + p^{med} \cdot n_{k,med}^j,$$

where $n_{k,ind}^j$, $n_{k,fam}^j$, and $n_{k,med}^j$ are the number of household members covered under each type.

Potential contracts. The interaction of private access and Medicaid eligibility yields up to six possible contracts. Formally,

$$\Omega(X_k, \bar{Y}) = \begin{cases} \{(\text{npriv}, \text{np}), (\text{npriv}, \text{med}), (\text{ind}, \text{np}), (\text{ind}, \text{med}), (\text{fam}, \text{np}), (\text{fam}, \text{med})\}, & \text{w/prob. } \pi_k^{\text{priv}}, \\ \{(\text{npriv}, \text{np}), (\text{npriv}, \text{med})\}, & \text{w/prob. } 1 - \pi_k^{\text{priv}}, \end{cases}$$

where npriv denotes no private coverage, np denotes no public coverage, ind individual private coverage, fam family private coverage, and med Medicaid coverage for eligible members.

F Parameters set externally

Tax function

Following [Feldstein \(1969\)](#), [Benabou \(2000\)](#), and [Heathcote et al. \(2017\)](#), I model after-tax income using a non-linear tax function of the form:

$$Y(x; \eta) = (1 - v_\eta) \cdot x^{1-\tau_\eta},$$

where x denotes pre-tax income, and v_η and τ_η control the level and progressivity of the tax schedule, respectively. The parameters (v_η, τ_η) vary by the number of adults. I adopt externally estimated values from [Borella et al. \(2023\)](#), averaging across the years 1998, and 2000.

Medicaid eligibility rules

I compile Medicaid eligibility thresholds $\bar{Y}$ from a variety of administrative and historical sources. Detailed documentation is provided in Appendix B.

Transfer function

The transfer function captures the amount of support a household receives when its income net of medical spending falls below a minimum consumption threshold. Formally, the transfer is defined as:

$$\text{Tr}(y, \text{OOP}; \eta) = \max \{c_{\min} \cdot \zeta^\eta + \text{OOP} - y, 0\},$$

where y denotes after-tax income, OOP is out-of-pocket medical spending, $c_{\min}$ is the per capita minimum consumption level, and ζ^η is an equivalence scale that adjusts the consumption threshold to household composition η . Transfers are positive only when residual

income after medical expenses is insufficient to meet the minimum consumption floor. I set $c_{\min} = 0.5FPL$, or roughly \$5,000 for a single individual. This is in line with existing estimates, which range from \$7,000 (Hubbard et al., 1994) to \$2,600–\$3,500 ((De Nardi et al., 2010; Aizawa and Fu, 2020; Russo, 2022)).

Out-of-pocket functions and access to treatment

Out-of-pocket (OOP) medical expenses are characterized by three parameters: a deductible d , a coinsurance rate ρ , and a maximum out-of-pocket limit $OOP^{\max}$. The household pays the full cost of medical spending up to the deductible, a fraction ρ of spending above the deductible, and no additional cost once the out-of-pocket maximum is reached.

Formally, for total medical spending m , the out-of-pocket payment is given by:

$$OOP(m) = \begin{cases} m, & \text{if } m \leq d \\ d + \rho(m - d), & \text{if } d < m \leq \bar{m} \\ OOP^{\max}, & \text{if } m > \bar{m} \end{cases}$$

where $\bar{m} = d + \frac{OOP^{\max} - d}{\rho}$ is the level of medical spending at which the household reaches the out-of-pocket maximum.

Each private health insurance plan, Medicaid and uninsured is defined by a specific triplet $(d, \rho, OOP^{\max})$. I specify the following parameters

Table: Health Insurance Contract Parameters

Plan Type	Deductible d	Coinsurance ρ	OOP Max $OOP^{\max}$
Private (Single)	\$500	0.20	\$2,500
Private (Family)	\$1,500	0.20	\$7,500
Medicaid	\$0	0.10	\$2,500
Uninsured	\$0	1.00	∞

In this baseline, I set OOP^{unins} so that an uninsured individual pays the entire cost, then there is not uncompensated care.

Following Capatina and Keane (2025), I set the probability of not having access to treatment when uninsured to 0.35.

G Estimation: step 1

This appendix describes the estimation strategy for the health shock parameters and preference parameters. First, I show that medical expenditure of high income individual map to the health shocks. Second, I describe how the remaining parameters are estimated using Simulated Method of Moments (SMM).

Identification

Let household k face a vector of health shocks Λ_k , where each member i in the household draws a shock λ_i . The household chooses optimal treatment levels m_i for each member.

Result 1: If $\lambda_i = 0$, the household is indifferent between treatment, partial treatment and no treatment for individual i , and in all 3 cases the only choice is $m_i = 0$, and therefore generate the same level of utility.

Result 2: If $\lambda_i > 0$, the household compares the utility of treatment versus no treatment. Let m_{-i} denote the treatment choices for all household members other than i , and let y denote after-tax income net of the out-of-pocket cost associated with m_{-i} . Let $\text{OOP}(\lambda_i)$ denote the additional out-of-pocket expenditure required to treat individual i given shock λ_i , assuming a continuous and differentiable out-of-pocket function.

The household chooses full treatment for individual i when the utility loss from not treating the condition exceeds the consumption cost of treatment. This condition is given by:

$$w_x \left(\lambda_i + \frac{1}{\omega_x} \lambda_i^2 \right) > u(y) - u(y - \text{OOP}(\lambda_i)).$$

Applying the Mean Value Theorem to the right-hand side, there exists $\bar{c} \in (y - \text{OOP}(\lambda_i), y)$ such that:

$$u(y) - u(y - \text{OOP}(\lambda_i)) = \text{OOP}(\lambda_i) \cdot u'(\bar{c}).$$

Then, as $y \rightarrow \infty$, we have $\bar{c} \rightarrow \infty$, implying $u'(\bar{c}) \rightarrow 0$. This last result holds as long as λ_i is bounded from above, which holds by the log normality assumption. Since the left-hand side of the inequality does not depend on income, it follows that:

$$\lim_{y \rightarrow \infty} P(m_i = \lambda_i \mid \kappa^j) = 1 \quad \forall \lambda_i.$$

This result is intuitive: as income increases, the marginal utility of consumption approaches zero, reducing the utility cost of spending on medical care. Hence, full treatment becomes optimal for all positive shocks.

The result extends to the case of partial treatment as well. The logic remains the same:

for any positive shock, the utility loss from undertreatment dominates the cost of treatment in terms of forgone consumption when income is sufficiently high. Thus, as income approaches infinity, the household always finds it optimal to provide full treatment.

Result 3: The parameter w_x is identified from the difference in the probability of zero medical expenditure between high- and low-income individuals.

Consider a household facing an individual health shock λ_i that is arbitrarily small. The household chooses to treat the shock if:

$$w_x \left(\lambda_i + \frac{1}{\omega_x} \lambda_i^2 \right) > u(y) - u(y - \text{OOP}(\lambda_i)).$$

For small shocks, the out-of-pocket cost is approximately proportional to the shock size, so we can linearize the expression. Applying a first-order approximation:

$$u(y) - u(y - \text{OOP}(\lambda_i)) \approx u'(y) \cdot \text{OOP}(\lambda_i).$$

Assuming $\text{OOP}(\lambda_i) \approx \lambda_i$ for small shocks, we obtain:

$$w_x \left(\lambda_i + \frac{1}{\omega_x} \lambda_i^2 \right) > u'(y) \cdot \lambda_i.$$

Dividing both sides by λ_i and taking the limit as $\lambda_i \rightarrow 0$, we get:

$$\lim_{\lambda_i \rightarrow 0} w_x \left(1 + \frac{1}{\omega_x} \lambda_i \right) > u'(y) \quad \Rightarrow \quad w_x > u'(y).$$

This condition shows that the decision to treat even small shocks depends on the comparison between w_x and the marginal utility of consumption. Since $u'(y)$ decreases with income, the same small shock may be treated by high-income households but not by low-income ones. Thus, variation in treatment rates across income groups identifies w_x .

Result 4: The parameter ω_x controls the intensity of treatment.

The utility cost of not treating a health shock λ_i is given by:

$$w_x \left(\lambda_i + \frac{1}{\omega_x} \lambda_i^2 \right).$$

Importantly, ω_x determines the convexity of this cost function. When ω_x is small (approaching zero), the cost of even small shocks becomes large, making full treatment the optimal choice across income levels. In contrast, when ω_x is large, the cost function becomes less convex, reducing the utility penalty associated with partial or no treatment. In this case, for low-income households, the optimal decision may involve partial or no treatment, since the marginal utility of consumption is high and the treatment cost is less severe.

Health shock distribution

The parameters of the health shock distribution are identified based on the following result:

$$\lim_{y \rightarrow \infty} P(m_i = \lambda_i \mid \kappa^j) = 1 \quad \forall \lambda_i.$$

This result implies that, as income increases, the household chooses full treatment with probability one, for all realizations of the health shock.

Let $m_i = \lambda_i$. Since $\lambda_i \sim \text{ZILN}(\mu_i, \sigma_i^2, \varphi_i)$, we have:

$$\begin{aligned} P(m_i = 0) &= P(\lambda_i = 0) = \varphi_i, \\ E(m_i \mid m_i > 0) &= e^{\mu_i + \frac{1}{2}\sigma_i^2}, \\ \text{Var}(m_i \mid m_i > 0) &= (e^{\sigma_i^2} - 1) e^{2\mu_i + \sigma_i^2}. \end{aligned}$$

From the last two expressions, we can identify σ_i^2 as:

$$\sigma_i^2 = \log \left(1 + \frac{(E[m_i \mid m_i > 0])^2}{\text{Var}(m_i \mid m_i > 0)} \right).$$

Given σ_i^2 , the log-mean parameter μ_i is identified as:

$$\mu_i = \log(E[m_i \mid m_i > 0]) - \frac{1}{2}\sigma_i^2.$$

Preference parameters W and ω

The parameters $w(h_i)$ and $\omega(h_i)$ are identified from income-related variation in the probability and intensity of medical spending.

The parameter $w(h_i)$ controls the extensive margin of medical expenditure, i.e., the probability of positive medical spending across the income distribution. In contrast, $\omega(h_i)$ governs the intensity margin, capturing difference in expected medical expenditures conditional on being positive across these groups.

In combination, $w(h_i)$ and $\omega(h_i)$ are pinned down by matching both the probability of any spending and the mean of positive spending across income groups in the data. In the estimation I allow w and ω to vary by sex and health status. These moment conditions are included in the SMM estimation.

Data moments

Let each individual in the sample be indexed by i , and define their state vector as:

$$s_i = \{t_i, f_i, e_i, \text{age}_i, \text{sex}_i, \text{srhs}_i, \text{wg}_i\},$$

where t_i is the survey year, $f_i \in \{\text{low}, \text{high}\}$ indicates the individual's income group, e_i is the insurance coverage type, age_i , sex_i , and srhs_i are the individual's age, sex, and self-reported health status, and wg_i is the sampling weight.

I group individuals into cells s_{group} , defined by unique combinations of year t , income group f , age group, sex, and self-reported health status. An individual is assigned to a group if their characteristics match on t , f , sex, and srhs, and their age falls within the specified age group.

For each s_{group} , I compute the following empirical moments using sampling weights wg_i

1. Probability of zero medical spending:

$$\Pr(m = 0 \mid s_{\text{group}}) = \frac{\sum_{i \in s_{\text{group}}} \text{wg}_i \cdot 1(m_i = 0)}{\sum_{i \in s_{\text{group}}} \text{wg}_i} = \bar{\phi}_{\text{group}}$$

2. Expected medical spending, conditional on positive spending:

$$E(m \mid m > 0, s_{\text{group}}) = \frac{\sum_{i \in s_{\text{group}}} \text{wg}_i \cdot m_i \cdot 1(m_i > 0)}{\sum_{i \in s_{\text{group}}} \text{wg}_i \cdot 1(m_i > 0)} = \bar{m}_{\text{group}}$$

3. Conditional variance of medical spending:

$$\text{Var}(m \mid m > 0, s_{\text{group}}) = \frac{\sum_{i \in s_{\text{group}}} \text{wg}_i \cdot (m_i - \bar{m}_{\text{group}})^2 \cdot 1(m_i > 0)}{\sum_{i \in s_{\text{group}}} \text{wg}_i \cdot 1(m_i > 0)} = \bar{\sigma}_{m, \text{group}}^2$$

4. Implied variance of the log normal distribution:

$$\sigma_{\text{group}}^2 = \log \left(1 + \frac{\bar{m}_{\text{group}}^2}{\bar{\sigma}_{m, \text{group}}^2} \right)$$

5. Implied mean of the log normal distribution:

$$\mu_{\text{group}} = \log(\bar{m}_{\text{group}}) - \frac{1}{2} \sigma_{\text{group}}^2.$$

Where the last two set of moments exploits the properties of the log normal distribution.¹⁴

Each cell *group* is characterized by year, income group, coverage tier, mean age, sex, srhs, and two totals: W , the sum of sampling weights for all individuals in the group, and W_{pos} , the sum restricted to individuals with positive medical expenditures.

Finally, I compute data moments by running the following regressions, using W or W_{pos} as weights depending on the moment:

$$\begin{aligned}\log\left(\frac{\bar{\phi}_{\text{group}}}{1 - \bar{\phi}_{\text{group}}}\right) &= x_{\text{group}} \cdot \beta_{e,m}^{\phi} + \varepsilon^{\phi} \\ \bar{m}_{\text{group}} &= x_{\text{group}} \cdot \beta_{e,m}^{\bar{m}} + \varepsilon^{\bar{m}} \\ \log(\bar{\sigma}_{m,\text{group}}^2) &= x_{\text{group}} \cdot \beta_{e,m}^{\bar{\sigma}} + \varepsilon^{\bar{\sigma}} \\ \mu_{\text{group}} &= x_{\text{group}} \cdot \beta_{e,m}^{\mu} + \varepsilon^{\mu} \\ \log(\sigma_{\text{group}}^2) &= x_{\text{group}} \cdot \beta_{e,m}^{\sigma} + \varepsilon^{\sigma}\end{aligned}$$

The vector x_{group} includes a third-order polynomial in age, fully interacted with indicators for female, poor health, and low income, and there processes are estimated separately by insurance source (private, Medicaid or uninsured)

The estimated variance-covariance matrix of each moment condition can be calculated using the delta method, based on the variance-covariance matrix of the corresponding β_m coefficients from these regressions.

¹⁴Let $x \sim \text{LogNormal}(\mu, \sigma^2)$, so that $\log(x) \sim \mathcal{N}(\mu, \sigma^2)$. Then:

$$E[x] = e^{\mu + \frac{1}{2}\sigma^2}, \quad \text{Var}(x) = (e^{\sigma^2} - 1) e^{2\mu + \sigma^2}.$$

From this, the moment ratio

$$\frac{(E[x])^2}{\text{Var}(x)} = \frac{1}{e^{\sigma^2} - 1}$$

identifies the variance parameter:

$$\sigma^2 = \log\left(1 + \frac{(E[x])^2}{\text{Var}(x)}\right).$$

Given σ^2 , the mean parameter is identified as:

$$\mu = \log(E[x]) - \frac{1}{2}\sigma^2.$$

Applying these relationships to the empirical moments from each group, I compute:

$$\sigma_{\text{group}}^2 = \log\left(1 + \frac{\bar{m}_{\text{group}}^2}{\bar{\sigma}_{m,\text{group}}^2}\right), \quad \mu_{\text{group}} = \log(\bar{m}_{\text{group}}) - \frac{1}{2}\sigma_{\text{group}}^2.$$

Model implied moments and SMM objective

Let $\theta^a = \{\beta^\phi, \beta^\mu, \beta^\sigma, w, \omega\}$ denote the set of parameters governing the distribution of health shocks and medical consumption conditional on insurance choice.

For a given candidate $\tilde{\theta}^a$, I draw S independent realizations of health shocks Λ_k for each household k . For each draw $s = 1, \dots, S$, I compute optimal medical expenditures m_k^j by solving the second-stage problem, taking insurance choices as given. This generates S simulated samples of medical expenditures.

Using these simulated samples, I compute the same moments for each simulated sample s as I compute for the data. To reduce simulation error, I take the average of these moments over the S simulations. I set $S = 20$ ensuring that simulated error is of a order of magnitude lower than data uncertainty. With these estimates in hand, I construct the analogous $M_{\bar{N}}(\theta^a)$. And estimate

$$\theta^{SMM} = \arg \min_{\theta} \left(\hat{M}_N - M_{\bar{N}}(\theta) \right)' \hat{W}_N \left(\hat{M}_N - M_{\bar{N}}(\theta) \right)$$

where $\hat{M}_N$ are empirical moments, $M_{\bar{N}}(\theta)$ are simulated model moments, and $\hat{W}_N$ is the inverse of the diagonal variance-covariance matrix of the empirical moments.

Each demographic type h is defined over 8 age bins, 2 sexes, and 2 self-reported health status categories, yielding $8 \times 2 \times 2 = 32$ distinct groups. Focusing exclusively on privately insured individuals $e_i = \{1, 2\}$, I construct $3 \times 32 = 96$ empirical moments for medicaid expenditures, including the probability of zero medical expenditure, the implied variance of the log normal distribution and the implied mean of the log normal distribution. I also include the difference in the probability of zero medical expenditure between high and low income individuals, adding additional 32 moments. That is I have 128 moments. Let denote these moments by $\hat{M}_N$. Let $\hat{W}_N$ denote the inverse of the diagonal of the empirical variance-covariance matrix of the moments.

H Estimation: step 2

Given $\hat{\theta}^a = \{\theta^{aSMM}\}$, I estimate $\theta^b = \{\Gamma, \Psi, \Sigma\}$ by maximum likelihood. I solve the household insurance choice problem conditional on $\hat{\theta}^a$ and construct the predicted choice probabilities over the available set of contracts. Let $\pi_{kj}(\theta^b; \hat{\theta}^a)$ denote the model-implied probability that household k selects contract j under parameters $(\theta^b; \hat{\theta}^a)$, where $j \in \{1, \dots, 6\}$. Let d_{kj} be an indicator equal to one if household k chooses contract j in the data.

θ^b is estimated by Maximum Likelihood, solving:

$$\hat{\theta}^b = \arg \max_{\theta^b} \sum_k \sum_{j=1}^6 d_{kj} \log \pi_{kj}(\theta^b; \hat{\theta}^a),$$

or equivalently by minimizing the negative log-likelihood function:

$$\mathcal{L}(\theta^b) = - \sum_k \sum_{j=1}^6 d_{kj} \log \pi_{kj}(\theta^b; \hat{\theta}^a).$$

In this step, I estimate a total of 25 parameters. This includes household type specific risk aversion, application cost for Medicaid, and mixing public and private health insurance. I also estimate 2 scale parameters for the type 1 E.V. shock, one for single adults households and one for two adults households.

I Estimated parameters

Table 2: Estimated Parameters of the Health Shock Process

Covariate	β^μ (Log-Mean)	β^σ (Log-Variance)	β^ϕ (Zero Spending)
Constant	5.807	1.570	-4.032
Age	1.778	-1.850	13.991
Age ²	-5.388	5.326	-23.150
Age ³	5.574	-3.771	10.598
Const $\times$ Bad Health	1.456	0.072	-0.807
Age $\times$ Bad Health	-5.759	1.076	-1.516
Age ² $\times$ Bad Health	14.804	-4.148	5.811
Age ³ $\times$ Bad Health	-9.721	3.054	-5.559
Const $\times$ Female	-0.659	0.024	1.182
Age $\times$ Female	6.576	-1.151	-11.142
Age ² $\times$ Female	-7.915	1.591	16.709
Age ³ $\times$ Female	1.965	-0.703	-7.489
Const $\times$ Bad $\times$ Female	-0.158	-0.341	-0.585
Age $\times$ Bad $\times$ Female	0.232	3.183	9.482
Age ² $\times$ Bad $\times$ Female	-0.867	-6.464	-22.901
Age ³ $\times$ Bad $\times$ Female	0.992	3.605	14.565

Table 3: Estimated Parameters

Parameter	Group	Estimate
<i>First Step Estimates</i>		
<u>w (Utility weight on treatment)</u>		
	Men, good health	1.30E-09
	Men, bad health	1.28E-09
	Women, good health	2.01E-09
	Women, bad health	2.41E-09
<u>ω (Curvature of treatment utility)</u>		
	Men, good health	1188.1
	Men, bad health	690.0
	Women, good health	1069.5
	Women, bad health	233.4
<i>Second Step Estimates</i>		
<u>γ_r (Risk aversion)</u>		
	Single-adult w/o children	3.25
	Single-adult w/children	5.26
	Two-adult w/o children	4.48
	Two-adult w/children	1.4
<u>ψ_{med} (Cost of applying for Medicaid)</u>		
	Single-adult w/o children	1767.89
	Single-adult w/children	8.29
	Two-adult w/o children	1915.24
	Two-adult w/children	1993.14
<u>$\psi_{\text{med,priv}}$ (Cost of mixing public and private coverage)</u>		
	Single-adult w/children	4316.90
	Households without children	4850.45
	Households with children	2755.42

Notes: Parameters w and ω are estimated in the first stage using the Simulated Method of Moments (SMM). Parameters γ_r , ψ_{med} , and $\psi_{\text{med,priv}}$ are estimated in the second stage by maximum likelihood. Full details are provided in Appendix H.

J Model Fit

J.1 Coverage choices


Figure 9 compares the model-predicted contract choice probabilities to their empirical counterparts. The model successfully replicates the distribution of households across six contract types, defined by combinations of private and Medicaid coverage.

Figure 10 shows individual-level coverage shares (private, Medicaid, and uninsured) for children (ages 0–18) and adults (ages 19–64). The model closely matches the observed distribution across these categories.

Figure 11 disaggregates insurance coverage by age. The model captures key life-cycle patterns. Among children, it reproduces higher public coverage at younger ages and a rising share of private insurance as they get older. Among adults, the model matches the lower private coverage, higher Medicaid participation, and higher uninsurance rates observed for younger individuals.


Figure 12 breaks down insurance coverage by family income. The model captures the increase in private coverage as income rises, the corresponding decline in Medicaid participation, and the reduction in uninsurance rates at higher income levels.

Figure 9: Choice probabilities across insurance contracts




Note: NP-NP = No Private – No Public, NP-AP = No Private – Medicaid, SC-NP = Individual coverage – No Public, SC-AP = Individual coverage – Medicaid, FC-NP = Family coverage – No Public, FC-AP = Family coverage – Medicaid.

Figure 10: Private, Medicaid, and uninsured coverage by age group




(a) Children (0-18), Adults (19-64), and All Ages (0-64).



(b) Adult subgroups: all adults (19-64), adults with children, and adults without children.


Note: The figure shows model-predicted coverage rates for private insurance, Medicaid, and uninsured categories. Panel (a) presents results for children, adults, and all ages, while Panel (b) focuses on adult subgroups.

Figure 11: Life-cycle patterns in coverage: Private, Medicaid and uninsured



Note: Own calculations using the Household component of the Medical Expenditure Panel Survey 1998-2000. Sample includes heads of household, spouses/cohabitants (ages 18-64), and their dependent children. The figure shows model-predicted and empirical coverage rates break down by age-groups.

Figure 12: Individual coverage by family income



Note: Own calculations using the Household component of the Medical Expenditure Panel Survey 1998-2000. Sample includes heads of household, spouses/cohabitants (ages 18-64), and their dependent children. The figure shows model-predicted and empirical coverage rates by Family income. Family income is measured relative to the federal poverty line.


J.2 Medical expenditure

Figure 13 presents the model fit for medical expenditures over the life cycle. Panel (a) shows mean individual total medical expenditures. Panel (b) reports the probability that medical expenditures are below \$100 for low- and high-income individuals, which I use as a proxy for treatment. Panel (c) depicts mean medical expenditures conditional on health status, and Panel (d) displays mean expenditures by source of coverage.


Figure 14 summarizes the model fit for differences in medical expenditures by income. Panel (a) shows the probability that medical expenditures are below \$100 for low- and high-

income individuals, and Panel (b) reports mean individual medical expenditures by family income group.


Figure 13: Medical expenditures by age




(a) Mean total medical expenditure



(b) Prob. of $m < 100$




(c) By health status



(d) By insurance choice

Note: Own calculations using the Household component of the Medical Expenditure Panel Survey 1998-2000. Sample includes heads of household, spouses/cohabitants (ages 18-64), and their dependent children. *Top left:* Mean medical expenditure by age *Top right:* Probability medical expenditures below \$ 100. *Bottom left:* Mean medical expenditure by age and self-reported health status. *Bottom right:* Mean medical expenditure by age and source of insurance. 95% confidence interval reported.

Figure 14: Medical expenditures and difference by income



(a) Prob. of $m < 100$


(b) By income family income

Note: Own calculations using the Household component of the Medical Expenditure Panel Survey 1998-2000. Sample includes heads of household, spouses/cohabitants (ages 18-64), and their dependent children. *Left:* Probability medical expenditures below \$ 100. *Right:* Mean individual medical expenditures by family income. 95% confidence interval reported.

J.3 Premiums

Figure 15 compares equilibrium premiums from the model with empirical premiums. The model captures the variation in premiums across contract types and the role of adverse selection in pricing.

Figure 15: Model fit: Equilibrium premiums across contract types



Note: Own calculations using the Household component of the Medical Expenditure Panel Survey 1998-2000. Sample includes heads of household, spouses/cohabitants (ages 18-64), and their dependent children.

K Counterfactual Exercises

Table 4: Counterfactual Exercise 1: Coverage

Group	Eligibility			Medicaid			Private			Uninsured		
	Baseline	Exc.1	Δ	Baseline	Exc.1	Δ	Baseline	Exc.1	Δ	Baseline	Exc.1	Δ
Panel A: Aggregate Groups												
Children (0–18)	28.0	9.7	-18.4	14.4	8.2	-6.2	77.6	82.2	4.6	8.0	9.6	1.6
Adults (19–64)	9.1	7.6	-1.5	6.6	5.5	-1.2	79.8	81.4	1.5	13.6	13.2	-0.4
Adults with children	10.0	7.0	-2.9	8.2	5.8	-2.3	83.7	87.1	3.4	8.1	7.0	-1.1
Adults without children	8.4	8.0	-0.4	5.6	5.2	-0.4	77.1	77.3	0.3	17.4	17.5	0.1
All Ages	14.4	8.2	-6.2	8.8	6.2	-2.6	79.2	81.6	2.4	12.0	12.2	0.2
Panel B: By Age Groups												
0–5	36.3	11.9	-24.4	20.7	9.9	-10.9	70.4	78.3	7.9	8.8	11.8	3.0
6–11	28.5	9.8	-18.6	13.8	8.5	-5.3	78.3	82.3	4.0	7.9	9.2	1.3
12–18	21.1	7.9	-13.1	10.2	6.8	-3.4	82.5	85.1	2.6	7.3	8.1	0.8
19–24	27.8	21.6	-6.2	15.0	9.7	-5.3	50.1	54.9	4.8	34.9	35.5	0.6
25–34	9.5	6.7	-2.8	7.1	4.9	-2.1	78.5	81.8	3.3	14.4	13.3	-1.1
35–44	6.4	5.8	-0.6	5.3	4.8	-0.5	85.7	86.6	0.9	9.1	8.6	-0.5
45–54	5.5	5.5	-0.0	4.7	4.7	-0.0	86.8	86.9	0.1	8.5	8.4	-0.0
55–64	7.3	7.3	-0.0	6.3	6.3	-0.0	77.8	77.8	0.0	15.9	15.8	-0.0
Panel C: By Family Income (Children)												
0–100	98.2	62.3	-35.9	65.6	51.5	-14.1	20.0	27.0	7.0	14.4	21.5	7.1
100–200	67.4	5.9	-61.5	31.4	5.9	-25.5	55.8	76.6	20.8	12.8	17.5	4.7
200–300	14.2	0.1	-14.1	3.7	0.1	-3.6	89.0	92.2	3.2	7.3	7.8	0.4
300–400	1.3	0.0	-1.3	0.2	0.0	-0.2	94.9	95.2	0.2	4.9	4.8	-0.0
400–500	0.0	0.0	0.0	0.0	0.0	0.0	95.9	95.9	0.0	4.1	4.1	-0.0
Panel D: By Family Income (Adults)												
0–100	66.0	60.9	-5.1	44.6	40.7	-3.9	24.2	28.3	4.1	31.3	31.1	-0.2
100–200	13.8	8.3	-5.5	12.7	8.0	-4.7	59.7	66.6	6.9	27.7	25.4	-2.2
200–300	1.4	0.1	-1.3	0.9	0.1	-0.8	86.5	87.4	0.9	12.7	12.5	-0.1
300–400	0.1	0.0	-0.1	0.0	0.0	-0.0	91.0	91.1	0.1	8.9	8.9	-0.0
400–500	0.0	0.0	0.0	0.0	0.0	0.0	92.6	92.6	-0.0	7.4	7.4	0.0

Notes: This table reports baseline and counterfactual (Exercise 1) values for Medicaid eligibility, Medicaid enrollment, private coverage, and uninsurance. Δ denotes the difference between the counterfactual and baseline economies. Panel A aggregates all individuals; Panels B–D disaggregate by age and income. Income is reported up to 500 % of the FPL. All numbers are expressed in percentage points.

Table 5: Counterfactual Exercise 2: Coverage

Group	Eligibility			Medicaid			Private			Uninsured		
	Baseline	Exc.2	Δ	Baseline	Exc.2	Δ	Baseline	Exc.2	Δ	Baseline	Exc.2	Δ
Panel A: Aggregate Groups												
Children (0–18)	28.0	42.4	14.4	14.5	20.3	5.8	77.5	73.1	-4.4	8.1	6.6	-1.4
Adults (19–64)	9.1	11.7	2.7	6.6	9.1	2.5	79.8	78.0	-1.8	13.6	12.9	-0.7
Adults with children	10.0	12.7	2.8	8.1	11.0	3.0	83.8	81.2	-2.6	8.1	7.8	-0.3
Adults without children	8.4	11.0	2.6	5.5	7.7	2.2	77.1	75.8	-1.3	17.4	16.5	-0.9
All Ages	14.4	20.4	6.0	8.8	12.2	3.4	79.2	76.6	-2.5	12.0	11.2	-0.9
Panel B: By Age Groups												
0–5	36.3	47.9	11.7	20.7	26.6	5.9	70.3	66.2	-4.1	9.0	7.2	-1.8
6–11	28.5	44.0	15.6	13.9	20.2	6.3	78.1	73.4	-4.7	8.0	6.4	-1.6
12–18	21.1	36.2	15.1	10.3	15.4	5.2	82.4	78.2	-4.2	7.3	6.4	-1.0
19–24	27.8	33.7	5.9	14.9	19.6	4.7	50.2	47.2	-3.0	34.9	33.3	-1.7
25–34	9.5	12.9	3.5	7.0	10.3	3.3	78.5	75.9	-2.6	14.5	13.7	-0.7
35–44	6.4	8.7	2.2	5.2	7.4	2.2	85.6	83.9	-1.7	9.1	8.7	-0.4
45–54	5.5	6.7	1.3	4.7	5.9	1.3	86.8	85.9	-0.9	8.5	8.2	-0.3
55–64	7.3	9.5	2.2	6.3	8.5	2.2	77.8	76.6	-1.2	15.9	14.9	-1.0
Panel C: By Family Income (Children)												
0–100	98.3	100.0	1.7	65.4	69.2	3.9	19.9	17.7	-2.2	14.8	13.1	-1.7
100–200	67.4	98.9	31.5	31.7	51.1	19.4	55.3	40.5	-14.8	13.0	8.4	-4.6
200–300	14.2	47.4	33.2	3.9	12.5	8.6	88.7	81.3	-7.4	7.4	6.2	-1.2
300–400	1.3	8.8	7.5	0.2	1.3	1.1	94.9	94.0	-0.9	4.9	4.7	-0.2
400–500	0.0	0.0	0.0	0.0	0.0	0.0	95.9	95.9	-0.0	4.1	4.1	0.0
Panel D: By Family Income (Adults)												
0–100	66.1	72.6	6.5	44.4	49.2	4.8	24.2	21.0	-3.1	31.4	29.8	-1.6
100–200	13.8	26.3	12.5	12.5	24.6	12.1	59.8	50.8	-9.0	27.7	24.6	-3.1
200–300	1.4	2.8	1.4	0.9	2.2	1.3	86.5	85.0	-1.5	12.7	12.8	0.1
300–400	0.1	0.1	-0.0	0.0	0.0	-0.0	91.0	91.0	0.0	9.0	9.0	-0.0
400–500	0.0	0.0	0.0	0.0	0.0	0.0	92.6	92.5	-0.0	7.4	7.5	0.0

Notes: This table reports baseline and counterfactual (Exercise 2) values for Medicaid eligibility, Medicaid enrollment, private coverage, and uninsurance. Δ denotes the difference between the counterfactual and baseline economies. Panel A aggregates all individuals; Panels B–D disaggregate by age and income. Income is reported up to 500 % of the FPL. All numbers are expressed in percentage points.