What Prevents Health Insurance Take-up?: Quantifying Frictions and Policy Responses*

Soheil Ghili[†], Igal Hendel[‡], Michael Sullivan[§], Michael D. Whinston[¶]

September 10, 2026

Abstract

Most consumers eligible for plans listed on Affordable Care Act (ACA) health insurance exchanges remained uninsured despite substantial premium subsidies and, historically, fines for lacking coverage. Several impediments, including adverse selection, informational frictions, uncompensated care, and hassle costs, could explain low participation. We use rich data on consumers' medical conditions, medical expenses by payer, and insurance status to estimate a model of exchange equilibrium to assess these impediments. We find that uncompensated care depresses take-up by more than adverse selection does, and by considerably more than informational frictions. Uncompensated care depresses consumers' willingness to pay for insurance to the extent that the exchange market fully unravels without premium subsidies.

*We thank Amy Finkelstein and seminar participants at Princeton, Emory, Stanford, UC Berkeley, UCLA, and the University of Pennsylvania for comments; Benjamin Vatter, Fernando Pereira Cordeiro, Anran Li, and Russ Yoon for excellent research assistance; and Ray Kuntz and Emma Evans for their help in accessing restricted MEPS data at the Agency for Healthcare Research and Quality. We thank the financial support of NSF grant SES-1758201.

[†]Yale University; Email Address: soheil.ghili@yale.edu

[‡]Northwestern University; Email Address: igal@northwestern.edu

[§]University of British Columbia; Email Address: michael.sullivan@sauder.ubc.ca

[¶]Massachusetts Institute of Technology; Email Address: whinston@mit.edu

1 Introduction

At the time of the passing of the Affordable Care Act (ACA) in 2010, about 50 million Americans were not covered by comprehensive health insurance (Rae et al., 2021). The ACA included several measures intended to expand coverage. The most prominent measures were the expansion of Medicaid, a federal program providing insurance for low-income Americans, and the creation of health insurance exchanges intended to serve those without coverage but with incomes too high to qualify for Medicaid. The exchanges are organized platforms on which private insurers list plans that adhere to regulated plan features. The ACA stipulates generous premium subsidies and initially featured a penalty for failure to purchase insurance called the individual mandate; this penalty was as high as 2.5% of household income for some consumers. Lawmakers enacted the individual mandate with the goal of preventing the market from unraveling due to adverse selection (Akerlof, 1978). Adverse selection is especially problematic as the ACA bans pricing based on pre-existing medical conditions. Despite these inducements, about 25 million Americans remain uninsured almost 15 years after the passage of the ACA.[1]

Middling take-up of ACA exchange insurance suggests the presence of major impediments to participation. Many consumers—including those with significant medical expenses—decline highly subsidized insurance. A standard model of insurance choice with risk-averse consumers and correct beliefs about medical expenses cannot rationalize such low participation.

The health insurance literature has documented low levels of consumer willingness to pay for insurance and discussed its possible causes (Finkelstein et al., 2019). Among them is uncompensated care — i.e., the failure of uninsured consumers to pay the entirety of their medical bills. Bills often go unpaid as consumers declare or threaten to declare bankruptcy to avoid paying their bills (Mahoney 2015). In addition, consumers may rely on charitable contributions as well as state and local programs to pay their bills. The ability of consumers to avoid footing their medical bills naturally lowers their willingness to pay for insurance.

Informational frictions may also explain limited take-up of exchange insurance. Abaluck and Gruber (2011) and Abaluck and Gruber (2016) show that consumers may not fully internalize future out of pocket expenses (OOPs), which leads them to undervalue coverage. We use the term *myopia* to refer to incomplete internalization of OOPs in consumers' insurance choices. This "myopia" could result from behavioral biases, or instead reflect financial constraints that make consumers more concerned about immediate payments than about future ones. Hassle costs (Shepard and Wagner 2022, McIntyre et al. 2021), other types of informational frictions (Goldin et al. (2019), Domurat et al. (2021), and Myerson et al. (2022)) and ideologically motivated opposition to the mandate may also compromise insurance uptake.

In this paper, we estimate the role of various impediments to insurance take-up in suppressing exchange participation. To do so, we develop and estimate a model of exchanges that augments that of Handel et al. (2015) to incorporate these frictions. In estimating the model, we use data

[1] See here: `https://www.kff.org/uninsured/issue-brief/key-facts-about-the-uninsured-population/`.

directly measuring uncompensated care and, to identify myopia, the co-variation of expected medical costs and participation decisions. With the estimated model in hand, we compute the extent to which each impediment depresses participation with the aim of determining which are most important. In light of the results of this horserace, we evaluate the impact of various policies intended to bolster take-up on participation and welfare.

Our primary data source is the Medical Expenditure Panel Survey (MEPS), which provides information on panelists' source of insurance (if any), medical conditions, and medical events (e.g., emergency room visits or prescription drug purchases). MEPS, unlike other sources of information on insurance choice (such as the American Community Survey), includes data on whether each panelist held insurance through an exchange.

Our main empirical innovation is in measuring uncompensated care and its impact on willingness to pay for insurance. MEPS reports payments for medical services by *type of payer* (e.g., out-of-pocket, insurance, or government program) and *by medical event* (e.g., a particular inpatient stay or emergency room visit), thus permitting the separate measurement of two forms of uncompensated care. The first form, *charitable care*, encompasses payments for medical services made by entities other than the patient or an insurer (e.g., charities or government programs); the second, *true uncompensated care* comprises non-recouped expenses that are implicitly borne by providers. Whereas the former is recorded directly in our data, we infer the latter. To do so, we estimate the gap between payments received by providers from treating an insured consumer and those received from treating an uninsured one *conditional on the consumer's medical utilization* (i.e., their mix of medical events and services received therein). This gap reflects implicit subsidization of uninsured consumers by medical providers. We find that the mean amount of true uncompensated care in the entire uninsured population is \$806, somewhat lower than the mean level of charitable care of \$991.

Taking both forms of uncompensated care into account, insured individuals pay 16.9% of their medical expenses on average, against 29.1% for the uninsured: more than the insured pay, but still a fraction of the bill. The uncompensated share also rises with medical utilization—uninsured consumers pay 30.5% out of pocket in the bottom nine deciles of utilization but only 13.7% in the top decile—so uncompensated care functions as a *de facto* catastrophic plan, offered at zero premium to the uninsured.

To quantify the equilibrium impact of uncompensated care relative to other frictions, we formulate a model of the exchanges. Ingredients of the model include, in addition to these frictions: Adjusted Clinical Group (ACG)-type risk scores mapping consumers' pre-existing diagnostics to expected expenses, the distribution of health costs conditions on risk score, and distributions of OOPs conditional on insurance status and health costs. In the model, consumers with constant absolute risk aversion (CARA) utility make a binary choice of coverage to maximize their expected utility taking expectations over these OOP distributions, which embed uncompensated care. In addition, the consumer choice model features myopia à la Abaluck and Gruber (2011) that varies with income and age, as well as heterogeneity in tastes for insurance unre-

lated to plan financial characteristics across geography and income groups. We estimate the binary choice model using the MEPS data and a method-of-moments estimator that matches model-predicted participation patterns to those in the data.

The estimated model permits simulating equilibrium participation and welfare under counterfactual regimes featuring varying frictions and regulations. Our first counterfactual exercise runs a horserace between impediments weighing on participation, including uncompensated care, myopia, and adverse selection. We find that uncompensated care is a first-order impediment to take-up: within the ACA's community-rating regime it is the largest of the frictions we consider. Under the baseline regime with status quo regulations (premium subsidies but no individual mandate), the model predicts a 32.0% participation rate among the exchange-eligible population. We estimate that participation would increase to 84.1% upon the elimination of uncompensated care. Based on our estimates, myopia plays an important but more modest role: if consumers were to weight OOPs dollar-to-dollar with premiums when deciding whether to buy insurance, participation would be 58.3%.

We additionally use our model to quantify the effects of policies intended to raise exchange participation. At their current level, we estimate that premium subsidies are much more effective than was the mandate penalty prior to its elimination in 2019. We find that absent subsidies the exchange market would fully unravel to zero participation as willingness to pay for insurance is below the cost of insuring any share of the population due largely to uncompensated care. Moreover, abolishing both subsidies—which boost participation—and uncompensated care—which depresses participation—simultaneously leads to 19.4%. Heavy subsidization of premiums is necessary to prevent market unraveling due to the presence of uncompensated care. We estimate that the mandate penalty had a smaller 6.6 percentage point impact on participation, a finding consistent with the fact that exchange participation did not significantly fall with the individual mandate's elimination in 2019. Moreover, these penalties only work when paired with subsidies: the market unravels without subsidies even with the individual mandate in effect.

A distinct factor limiting participation—and the one most discussed in the debate over enactment of the ACA—is adverse selection. To assess the role of adverse selection on participation, we compute equilibria under a risk-based pricing regime in which insurers price based on observed consumers' pre-existing medical conditions. Healthier individuals pay lower premiums under risk-based pricing, which consequently raises participation to 75.4% from 32.0% under the baseline community rating regime. Adverse selection is therefore a quantitatively important impediment, worth 43.4 percentage points against the 52.1 points attributable to uncompensated care.

Having characterized the participation effects of frictions and inducement policies, we conduct welfare analysis to understand the normative impact of two policies. First, we study the welfare impact of the expanded subsidies that were in effect from 2021 to 2025. The model predicts quite accurately the impact on participation. We find that as participation doubles attracting healthier individuals, premiums decline by 25.6%, while total welfare increases. The per capita

welfare increase is of the order of 28.2% of premiums. The second exercise we conduct is formalizing uncompensated care. We simulate a policy that replaces the current inefficient uncompensated care system with formal insurance with similar actuarial value as the current uncompensated care but which is more certain and has the same non-pecuniary benefits/costs (our z's) as exchange insurance. If such policy were (implementable and) offered to consumers participation in the exchanges would totally unravel. Yet, consumer surplus would increase, in the order of 22.9% of the baseline premiums, while total per capita surplus would increase more modestly, roughly 27.9% of that amount.

The paper is organized as follows: In Section 2, we describe our data. In Section 3, we describe the extent and patterns of uncompensated care revealed in the MEPS data. In Section 4, we describe the key features of the ACA exchanges and comment on how these features help us identify the role of various frictions in determining consumers' insurance purchase decisions. In Section 5, we describe the details of our model. In Section 6, we discuss the details of our estimation method and present our parameter estimates. In Section 7, we use these estimates and our model to examine the role of various frictions on exchange participation, welfare, and various potential government policies. Section 8 concludes.

1.1 Related literature

Our paper's first contribution is to a literature that identifies and quantifies frictions inhibiting health insurance take-up. First among the frictions studied in this literature are informational frictions. Several studies have used experiments to measure these frictions. These include Goldin et al. (2019), who study the effects of a randomized governmental outreach to individuals paying the individual mandate penalty on insurance take-up and health, and Domurat et al. (2021) and Myerson et al. (2022), who study a randomization in subsidy information sent to qualifying individuals in California. Each study finds a positive, albeit modest, effect of its informational intervention on participation in health insurance. The particular form of informational friction on which we focus is myopia, the umbrella term that we use for consumers' underweighting of OOPs relative to insurance premiums when choosing whether to purchase coverage. Myopia was identified by Abaluck and Gruber (2011) and Abaluck and Gruber (2016) in the context of insurance plan choice in Medicare Part D, a federal program that helps Medicare beneficiaries buy prescription drugs. Abaluck and Gruber (2011) find that consumers value a dollar of OOPs at 0.15–0.19 of a dollar of premiums, which they interpret as a form of myopia. Like Abaluck and Gruber (2011) and Abaluck and Gruber (2016) we also use a revealed preference approach linking consumers' insurance choice and medical risk. We build upon these studies by estimating myopia in the context of comprehensive health insurance, to then quantify the equilibrium implications of myopia. We capture other informational frictions (related, e.g., to lack of awareness of exchange insurance) by specifying consumers' (dis)taste for insurance as a function of year, income, and region, variables that may influence their level of information about exchange insurance.

Shepard and Wagner (2022) and McIntyre et al. (2021) document the impact of hassle costs as an impediment to insurance. They show the impact of auto-enrolment, suggesting that frictions lower enrollment. Drake et al. (2022) compare enrollment of those paying no premiums to those paying a dollar to sign up. The dollar premium lowers enrollment, which likely reflects hassle costs rather than monetary costs. One conclusion from this literature is that because such minor frictions reduce participation, willingness to pay for coverage must be low.

Uncompensated care has been considered an explanation for the low willingness to pay for insurance. Finkelstein et al. (2019) find that, in the context of the Massachusetts exchanges, individuals' willingness to pay for health insurance falls below the costs of providing coverage, attributing the gap primarily to uncompensated care. In doing so, they make reference to prior findings in the literature suggesting that uninsured individuals pay 20–35% of their costs of care (including 20% in the Oregon Health Insurance Experiment). Kaiser Family Foundation (2014) found that, in 2013, OOPs represented 23% of expenses for uninsured individuals. Insured individuals received some uncompensated care, but much less (under 5%). Finkelstein et al. (2018) also studies uncompensated care, and Brevoort et al. (2020) unpaid medical bills. Mahoney (2015) documents how the cost of bankruptcy relates to insurance take up, the latter being lower in regions and for individuals who find it less costly to declare bankruptcy. These two strands of the literature suggest uncompensated care is likely the reason low-income people have low willingness to pay for insurance. Camilleri (2018) estimates the effects of Medicaid expansion on uncompensated care and, in doing so, provides further evidence of uncompensated care: Camilleri (2018) finds that uncompensated care fell at hospitals that treat a disproportionately high number of low-income patients.

Our paper also adds to a literature studying policies that encourage participation in health insurance exchanges. Saltzman (2019) estimates demand in the California and Washington exchanges, reporting a modest response of enrollment to penalties. Saltzman (2021) studies government interventions such as risk adjustments and the individual mandate impact on a market with adverse selection and moral hazard, to assess the participation-underinsurance trade-off. Namely, policies can impact participation as well as the level of insurance of those who purchase coverage. His estimates imply a modest 2% reduction in enrollment from mandate repeal. Frean et al. (2017) find that coverage responded moderately to subsidies whereas penalties did little to impact participation. Tebaldi et al. (2023) studies the design of ACA premium subsidies given their role in limiting the negative participation impacts of adverse selection. We similarly identify a role for subsidies in stymieing adverse selection, although we focus on the interaction between uncompensated care and subsidies rather than on how the design of the subsidy program may be optimized to positively select healthy consumers as studied by Tebaldi et al. (2023).

Last, our article relates to the literature on health insurance choice models. These studies generally take one of two approaches. The first approach involves assuming that consumers are expected utility maximizers with specified utility index functions (e.g., CARA or constant

relative risk aversion) and distributions of health risk (Cardon and Hendel (2001), Handel et al. (2015), Ghili et al. (2024)). The other approach is to use a standard discrete choice model (DCM), capturing differentiation across plans indirectly through plan characteristics (e.g., Tebaldi et al. 2023, Tebaldi 2025). Whereas the latter approach is well suited to the analysis of changes in the set of insurance plans available to consumers, it does not permit analysis of changes in the distribution of health risk facing consumers given that the DCM preference parameters capture both consumers' perceptions of the underlying distributions of risk and consumer preferences over these distributions. Given our focus on studying the effects of changes in the distribution of health risk facing consumers rather than changes in plan characteristics, we use the former approach that involves explicitly specifying utility indices and risk distributions. Furthermore, we use data from MEPS given its rich information on panelist health conditions and payment by source, which allow us to estimate consumer responses to OOP distributions in quantifying myopia and also to directly measure uncompensated care. Quantifying these frictions is the primary goal of this article. MEPS does not feature detailed information on plan characteristics, which prevents us from incorporating such characteristics in our analysis to the extent of recent articles taking the DCM approach to insurance choice estimation.

2 Data

In this section, we describe the data sources that we use in our analysis.

2.1 MEPS

MEPS is a panel of US households administered by the Agency for Healthcare Research and Quality (AHRQ) that is intended to be representative of the civilian non-institutionalized population of the US. MEPS panelists are drawn from households responding to the National Health Interview Survey (NHIS) of the National Center for Health Statistics, which chooses households to participate in the NHIS by drawing from a database of US residential addresses.[2]

A new MEPS panel begins each year, and AHRQ tracks the constituent panelists over a two-year period. Each panelist's two-year tenure in MEPS is subdivided into five rounds, each of which ends with an interview conducted by MEPS administrators. MEPS compensates its respondents with $50 per completed interview. In addition to interviewing panelists in the "Household Component" of the survey, MEPS interviews medical providers in the "Medical Provider Component" to obtain information on medical visits, medical payments and their sources, and panelist medical conditions.

[2]The sample design of NHIS is described in detail here: `https://ftp.cdc.gov/pub/Health_Statistics/NCHS/Dataset_Documentation/NHIS/2019/srvydesc-508.pdf` (see page 9). The procedure that MEPS uses to randomly select NHIS respondents is described here: `https://meps.ahrq.gov/data_files/publications/mr33/mr33.pdf?_gl=1*1mfqzg8*_ga*MTY3NTAxODAzNi4xNzU2ODg3MjM2*_ga_45NDTD15CJ*czE3NTczNjAwNzkkbzEkZzEkdDE3NTczNjAxNDckajYwJGwwJGgw`. In 2019, the total household response rate to NHIS was 61.1%. The MEPS response rates were 40%-56% in the 2012–2019 time period.

MEPS offers several datasets containing distinct types of information on panelists. We use MEPS data from 2013 to 2019 on: panelists' demographic characteristics, including income, family size, and census region; insurance status; medical conditions applying to each panelist in the form of three-digit International Classification of Diseases (ICD) codes; and medical events.[3] We use the data on medical conditions to identify the health risks that an individual faces when considering her insurance decision.

The key information in MEPS that enables us to measure uncompensated care is that medical expenses are reported by payer (insured, insurance company, charities, etc.). This allows us to observe the share of total expenses (the amounts received by providers) that is paid by a consumer - the consumer's "out-of-pocket" (OOP) expenses.[4] The distributions of OOP expenses (net of expenses paid by other sources) that a consumer faces given her health status when insured versus when uninsured represent the risks faced by the individual as a consequence of her insurance decision.

MEPS includes panelist-level weights that are recommended for use in cross-panelist analyses. We use these weights in all of our panelist-level analysis; e.g., when examining the extent of uncompensated care, when forming the moments that we use in estimating our choice model, and when we investigate counterfactuals.

2.2 Premiums data

We use data from the Kaiser Family Foundation (KFF) on premiums by state for insurance plans offered on the ACA exchanges. In particular, we use their data on premiums for the second-lowest-cost Silver plan available in the state for a 40-year-old non-smoker. The Silver tier of insurance was the most popular tier during the sample period in large part because only Silver plans were eligible for income-based cost-sharing reductions (see Section 4.4).[5] Given that the publicly available MEPS data report panelists' census regions but lack state identifiers, we aggregate the KFF state-level premiums to the census region level: for each year, we compute average state-level premiums within each census region, weighting states by their populations. We then set the pre-subsidy premium faced by a 40-year-old consumer to the average premium for their region/year pair. To obtain premiums for other ages, we apply the ACA age schedule that we will describe in Section 4.2. Since (as we describe below) premium subsidies depend on

[3]MEPS provides data on the characteristics of medical events belonging to the following categories: outpatient hospital visits, inpatient hospital visits, emergency room visits, office-based visits, dental visits, home health care services, pharmaceutical prescriptions, and other events, a category that includes medical items and services falling outside the other categories (e.g., ambulance services, disposable supplies, and long-term medical equipment). The visit-based files provide information on the characteristics of visits (including procedures administered and timing) and on payments received by providers by source (more on this below). The pharmaceutical file instead reports characteristics of prescriptions filled including quantities dispensed and payments received for prescriptions by source. We do not incorporate information on dental services and expenses in our analysis given that such services are typically covered by specialized dental insurance rather than comprehensive health insurance plans of the sort offered on the ACA exchanges.

[4]In no part of our analysis do we use the "charges" variables provided by MEPS, which may refer to list prices or billed amounts rather than actually exchanged dollars.

[5]In the 2017 Open Enrollment period, 71% of consumers chose a Silver tier plan. See `https://www.cms.gov/newsroom/fact-sheets/health-insurance-marketplaces-2017-open-enrollment-period-final-enrollment-report-novembe`

income, we use information on premiums and consumers' income to compute the post-subsidy premium that each individual in the sample faced.

3 Uncompensated care

The possibility that the uninsured do not pay the full cost of their medical care — what the literature has called "uncompensated care"[6] — provides one possible explanation for a low willingness of consumers to pay for exchange insurance. MEPS allows us to quantify two components of uncompensated care. First, it measures directly the share of consumers' total expenses that are covered by insurance, OOPs, and other payers. We call payments made by these other payers *charitable care*. Second, MEPS includes data on total payments received for medical services along with detailed information on characteristics of these services and the patients receiving these services. This information allows us to assess the extent to which medical providers receive greater payments from insured consumers, which would indicate implicit subsidization of uninsured consumers by medical providers. We call this implicit subsidization *true uncompensated care*.

3.1 Charitable care

We assess the extent of charitable care by decomposing medical payments by payer. MEPS classifies payments made by various payers: out-of-pocket, insurance, state/local programs, Veterans Affairs, workers' compensation, other federal programs (which includes the Indian Health Service, military treatment facilities, and other care provided by the federal government), other public insurance, and other sources (which includes automobile, homeowner's, and liability insurance). Figure 1 plots the share of medical payments received by providers across these payment sources, separately for the uninsured and the exchange insured. The figure shows that about half of uninsured consumers' payments are covered by OOPs with charitable care accounting for the remainder. In contrast, insurance payments constitute almost 80% of insured individuals' total expenses, with OOPs comprising the overwhelming majority of other payments. Although the uninsured pay a larger share of total expenses than do the insured, they are far from covering the total payment amounts.

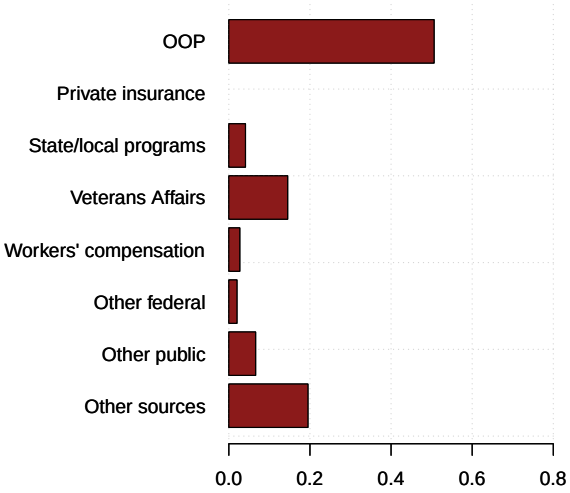
Next, we assess the relationship between an individual's total expenses and the amount of charitable care they receive. This relationship is critical for the insurance value of charitable care to the uninsured and thus consumers' willingness to pay for formal insurance. Charitable care, for example, would provide greater insurance value and more significantly lower the willingness to pay for formal insurance if it disproportionately covered large expenses.

Figure 2 shows this relationship. Restricting to panelist-years in which the individual was uninsured all year and had positive total expenses, we sort these panelist-years into deciles of total expenses and summarize, within each decile, the distribution of the charitable-care share

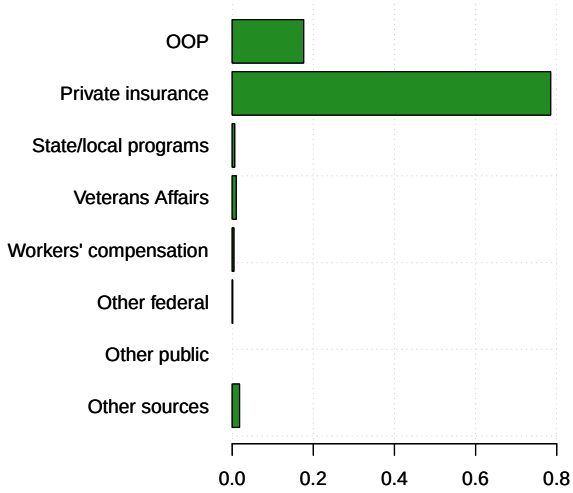
[6] See, e.g., Hadley et al. (2008), Coughlin et al. (2014), Dranove et al. (2016), Garthwaite et al. (2018), and Finkelstein et al. (2019).

Figure 1: Sources of payments for medical services (2013–2019)

(a) Uninsured



(b) Exchange insured



of expenses. The figure shows that charitable care represents a much larger source of payments for uninsured panelists with higher total expenses, especially in the top decile. It also indicates a large degree of cross-panelist variation in access to charitable care conditional on their total expenses.

3.2 True uncompensated care

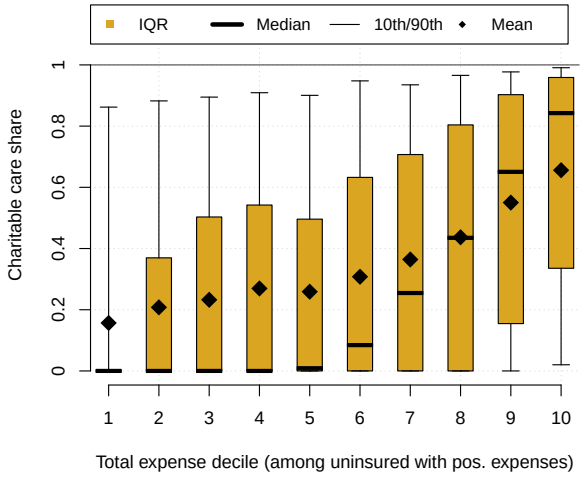
As we noted above, in addition to charitable care, providers may receive lower compensation from the uninsured, which we refer to as *true uncompensated care*. To quantify the extent of provider under-collection for services rendered to the uninsured we start by creating a measure of *utilization* for every individual in our sample and we then compare payments collected from insured versus uninsured patients.

We use MEPS data on the total payments providers receive when treating the insured to map observed annual medical events into predicted payments to providers for insured consumers. Specifically, we estimate using nonlinear least squares the following model on a sample of panelist-years in 2014–2019 for which the panelist was insured for the entire year (even if under a non-exchange plan), had an income at least 138% of the FPL, and was between ages 25 and 64:

$$\text{TOT}_{jt} = e^{\psi_t + \phi_{\text{region(j)}}}(x'_{jt}\beta + u_{jt}),$$

where j indicates a panelist, t indicates year, TOT_{jt} are total medical expenses (excluding dental expenses) and region(j) is panelist j's census region. The vector x_{jt} consists of the the number of medical events in each of seven categories (inpatient, outpatient, emergency room, office-based, prescription fills, home health, and other) together with the total visit count, counts

Figure 2: Charitable care as share of total expenses among the uninsured



Notes: This figure plots the interquartile range and median of the share of total expenses (payments received by providers) constituted by payments other than OOPs (i.e., charitable care). We conduct the analysis at the panelist-year level and include only panelists who are uninsured for the entire year and have positive total expenses.

of selected procedures performed at those events (surgeries, MRIs, anesthesia administrations, and deliveries), and the number of inpatient nights. The event counts and surgery counts are interacted with indicators for four age groups (25–34, 35–44, 45–54, and 55–64), so that each is priced separately by age; the remaining components of x_{jt} are priced uniformly.

Having estimated the mapping from medical events into payments on the sample of insureds, we generate our *utilization measure* m_{jt}. We do so by entering medical events for each panelist into the estimated mapping:

$$m_{jt} := e^{\hat{\psi}_t + \hat{\phi}_{\text{region(j)}}} \times x'_{jt}\hat{\beta}$$

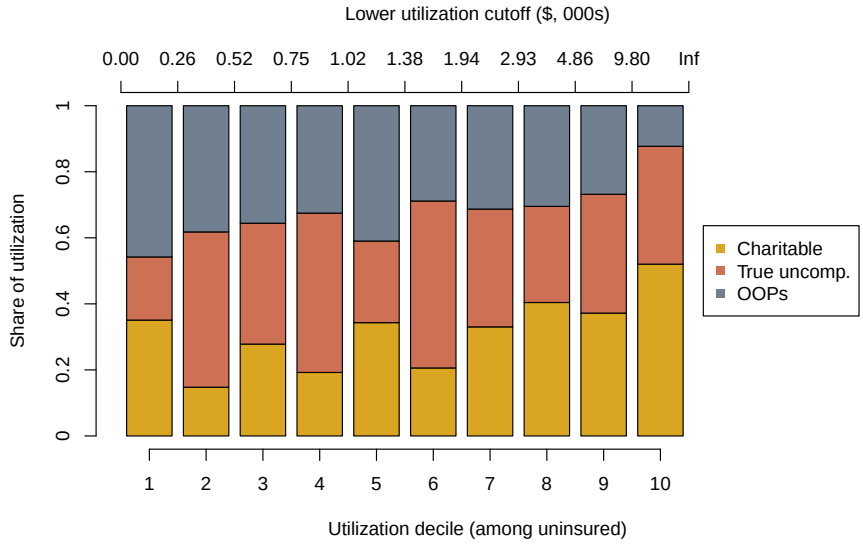
Figure 3 shows, for each decile of utilization, the share of utilization dollars accounted for by charitable care, true uncompensated care, and OOPs respectively.[7] (The top of the figure shows the utilization cutoffs for the deciles.) The figure shows that OOPs represent about 46% of payments in the lowest decile of utilization, but only 12% of payments in the highest. Overall, charitable care represents 55% of all uncompensated care: averaged over the entire uninsured population, charitable care amounts to $991 per panelist-year against $806 of true uncompensated care.[8]

Figure 4 shows characteristics of the resulting OOP distributions that the uninsured and ex-

[7]To average out the prediction error in utilization, this figure reports for each decile the ratio (true uncompensated care payments/aggregate utilization) and similarly for charitable care, with the remaining share being OOPs.

[8]Roughly half of uninsured panelist-years record no medical expenses at all and so receive neither form of uncompensated care. These means are taken over the entire uninsured population, including those panelist-years; among panelist-years with positive expenses, which are the ones the figure sorts into utilization deciles, the corresponding means are $1,915 and $1,539.

Figure 3: Share of utilization payments for the uninsured accounted for by charitable care, true uncompensated care, and OOPs



Notes: This figure decomposes payments received by providers for the care of uninsured panelists into out-of-pocket payments, charitable care, and true uncompensated care, each expressed as a share of utilization within deciles of utilization. The sample pools 2014–2019 and includes panelists aged 25 to 64 with family incomes of at least 138% of the federal poverty line who were uninsured for all twelve months of the year and had positive medical expenses. Dollar amounts are deflated to 2019 dollars using the medical-care consumer price index.

change insured face for the various utilization deciles, displaying the mean, median, interquartile range, and 10/90 range for the exchange insured in panel (a) and the uninsured in panel (b). (The deciles in both panels are defined for the overall sample, with the cutoffs shown at the top of each panel's figure.) Average OOPs are very similar at high utilization levels, but are about double for the uninsured at low utilization levels. Throughout the range of utilization. there is much more variation in the share of utilization represented by OOPs for the uninsured.[9] Averaging these ratios across panelist-years, the uninsured pay 29.1% of their medical expenses out of pocket against 16.9% for the exchange insured—more than the insured pay, but still a minority of the bill. Pooling the bottom nine deciles, the uninsured pay 30.5% of utilization out of pocket, against 13.7% in the top decile. Uninsurance looks like a catastrophic insurance plan, but with a highly uncertain level of coverage conditional on medical utilization.[10]

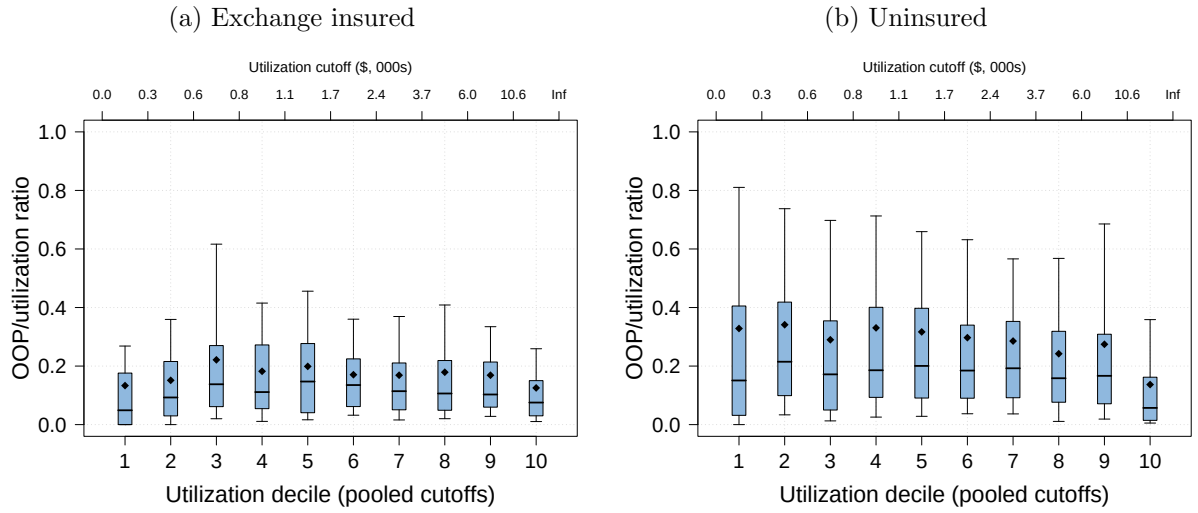
4 Features of the ACA exchanges

In this section, we describe institutional features of the ACA exchanges that are relevant for our modeling and estimation.

[9] In our estimation we will allow OOPs to depend on a consumer's observed income; our estimates show that the uninsured face greater OOP uncertainty for any given level of utilization even conditioning on income.

[10] Mahoney (2015) also notes how uninsurance looks like catastrophic insurance based on the average level OOPs paid by the uninsured displayed in his Figure 1 (Panel A), our Figure 4 adds to this observation the fact that this coverage is highly uncertain.

Figure 4: Out-of-pocket payments as a share of utilization, by utilization decile



Notes: Each panel summarizes, within each decile of total utilization, the distribution across panelist-year observations of OOPs as a share of utilization. The boxes indicate interquartile ranges, the horizontal lines indicate medians, the whiskers indicate 10/90 ranges, and the diamonds indicate means. Deciles are defined over the pooled sample, so the cutoffs (in thousands of dollars, shown across the top of each panel) are common to both panels. The sample includes exchange-eligible panelists pooled across 2014–2019; the insured panel covers panelists with twelve months of exchange coverage and the uninsured panel those uninsured for all twelve months. Dollar amounts are nominal, and the ratio is winsorized at its weighted 98th percentile within each panel.

4.1 Eligibility

The ACA exchanges serve individuals who are ineligible for Medicaid, a public program that provides health insurance for individuals with low incomes. Medicaid covers those with household incomes below 138% of the Federal Poverty Line (FPL) in states that expanded the program under the ACA ("expansion states") and below 100% of the FPL in non-expansion states. In 2019, the most recent year that we study, the FPL for a single-member household was $12,490 and increased $4,420 for each additional household member; these base and incremental amounts are adjusted for inflation each year.[11] As we will describe in Section 4.3, eligibility for premium subsidies on the ACA exchanges also depends on household income relative to the FPL.

We summarize each panelist's resources with an individualized income measure y_{jt} that can be compared directly to the single-person FPL to determine eligiblity for Medicaid and premium subsidies, both of which depend on household income relative to the FPL. Letting $FPL_t(n)$ denote the FPL level for a household of size n in year t, we define the individualized income for a MEPS panelist j with household size n and household income in year t equal to Y_{jt} to be

$$y_{jt} = Y_{jt} \times \frac{FPL_t(1)}{FPL_t(n)}.$$

By construction, individualized income y_{jt} relative to the single-person FPL equals household income Y_{jt} relative to the household's own FPL.

[11] Appendix A provides these amounts for each year in the sample period.

We restrict the sample of exchange-eligible consumers based on age and income. Specifically, we exclude individuals younger than 25 or older than 64, as well as those with incomes below 138% of the FPL. Individuals under 25 are generally covered by parental plans, whereas those 65 and older are eligible for Medicare. As noted above, individuals with incomes below 138% of the FPL were eligible for Medicaid in expansion states whereas those with incomes under 100% of the FPL were eligible in non-expansion states. We use the 138% threshold to ensure that Medicaid-eligible consumers are consistently excluded from the sample. In estimating our choice model we additionally drop individuals who belong to any insurance plan other than a plan that was directly purchased from an insurer. That is, our sample consists only of the uninsured and those who purchased insurance on an exchange.

4.2 Age-adjusted community rating

The ACA enacted a modified version of community rating that prevents insurers from pricing based on a consumer's observable pre-existing health conditions but in which premiums vary with a consumer's age according to a regulated schedule $\bar{p}(p_{25}, a_{jt})$. Figure 5 displays the ratio of the premiums charged to consumers of various ages relative to those charged to a 25 year old. The schedule rises by about three times between age 25 and 64. As noted in Handel et al. (2015), community rating exacerbates adverse selection because insurers are not allowed to price based on health conditions that a consumer knows.

Figure 5: Schedule of relative premiums by age

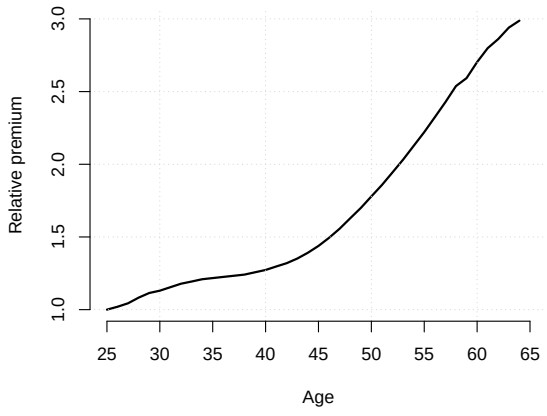
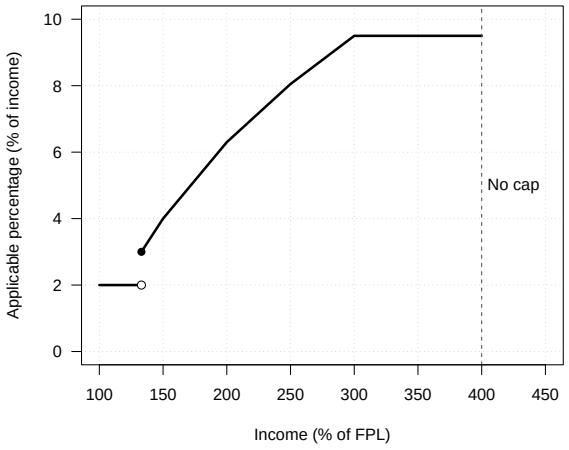


Figure 6: Applicable-percentage schedule for premium subsidies



4.3 Premium subsidies

Under the ACA, consumers with incomes below 400% of the FPL are eligible for premium subsidies that effectively cap the premium that they pay for exchange insurance at a fraction of household income. That fraction, which is set by a statutory schedule of applicable percentages, increases with household income relative to the FPL. Figure 6 displays this schedule. Premium subsidies are determined at the level of the tax family, which consists of an individual, their

spouse (if married filing jointly), and any dependents claimed on the individual's tax return This determination involves first calculating the tax family's pre-subsidy premium as the sum of the age-rated premiums of those family members who enroll in exchange coverage. The combined premium that these members pay is then the minimum of this pre-subsidy amount and the applicable percentage multiplied by tax-family income. Enrollment itself, however, is an individual decision: not every family member need enroll, and the premium one member effectively faces depends on which other members do. In our analysis, we follow this tax-family-level determination but abstract from this interdependence, since our model treats take-up as an individual choice. To assign a premium to an individual, we compute their tax family's post-subsidy premium as if all of its exchange-eligible members (those aged 25–64 with no non-exchange coverage during the year) enroll and, assuming family members share it evenly, divide by the number of exchange-eligible members.

We end our sample period in 2019, prior to the COVID-19 pandemic. In response to the pandemic, in 2021 Congress passed legislation that increased and expanded to higher income levels the premium subsidies, but these subsidy enhancements were rolled back at the end of 2025.

4.4 Cost-sharing reductions

Plans marketed under the ACA Silver tier also feature cost-sharing reductions (CSRs) that reduce coinsurance rates for consumers with incomes below 250% of the FPL. Table 1 reports the cost-sharing features of the Silver tier for income groups that are defined by FPL cutoffs. Group IV (above 250% of the FPL) is not subsidized; hence, its actuarial value is 70%, the designated level for Silver plans before CSRs. Each level of cost-sharing reductions has a specified actuarial value, deductible, and cap on OOP expenses ("Maximum OOP").

Table 1: Cost-sharing levels under the ACA Silver plan

	(I)	(II)	(III)	(IV)
Lower income cutoff (% of FPL)	100%	150%	200%	250%
Actuarial value	94%	87%	73%	70%
Deductible	\$0	\$250	\$1,750	\$2,000
Maximum OOP	\$1,000	\$2,000	\$4,000	\$5,500

As we explain in Section 6.4, in our modeling and estimation we capture how OOPs for insured consumers vary with their income and utilization without imposing the specific form described in Table 1.

4.5 Individual mandate

The ACA originally introduced a penalty for remaining uninsured known as the individual mandate. Under the individual mandate, the federal government taxed individuals who failed to purchase insurance an amount that, between 2016 and 2018, was equal to 2.5% of the

consumer's income above the government's tax-filing threshold ($10,650 in 2018) but no less than $695 and no greater than the premium charged for a benchmark exchange insurance plan before premium subsidies. Figure 7 plots the dollar amount of the individual mandate across various levels of consumer income in the year 2018. Federal legislation eliminated the individual mandate penalty beginning in 2019.

The mandate involved several exemptions relating to affordability, the tax-filing threshold, and short gaps in coverage. First, the mandate penalty was waived if the cheapest post-subsidy bronze exchange premium exceeded 8% of household income, in which case it was considered non-affordable. Second, individuals with income below the tax-filing threshold were exempt as well. Third, the first uninsured spell, if shorter than 3 months, was exempt from penalty.

We adjust the mandate penalty a consumer faces to account for these exemptions and we also consider that collection was incomplete even after accounting for these exemptions.[12] Incomplete collection reflected the IRS's limited wherewithal to collect penalties: by statute, the IRS could not use liens, levies, or criminal prosecution to collect penalty amounts and could recover them only by offsetting a filer's tax refund.[13] Consistent with incomplete collection, Treasury Inspector General for Tax Administration (2017) finds that the IRS did not pursue roughly 3.3 million "silent" returns reporting neither coverage, an exemption, nor a payment; and for tax year 2014 only about 7.5 million filers made a payment ($1.5 billion, averaging roughly $200), while some 12 million claimed an exemption (Koskinen, 2015).

When we compute from MEPS the implied penalties owed in 2015 by uninsured individuals after taking account of the affordability, tax-filing threshold, and short-gap exemptions, we project that of the 40.8 million returns with an uninsured individual, 25.5 million had no exemption and owed $18.9 billion. Instead, the IRS collected $3.1 billion in 2015[14], only 16% of the amount owed according to our MEPS data. Given this discrepancy, we include a non-compliance parameter in our estimation.

4.6 Exchange participation and turnover

When we estimate our choice model, we focus on individuals who made insurance decisions at the start of the calendar year during the January-March open enrollment period. Specifically, we focus on the sample of consumers who, during the January–March open enrollment window, were either exchange-insured in March (having entered during the window) or uninsured in all of January, February, and March. The former chose insurance while the latter chose uninsurance. Figure 8 shows the participation rates among this "Open Enrollment" sample for the years 2015-2019 in panel (a). Figure 9 shows participation rates in 2019 for age groups, quartiles of income, quartiles of risk score, and gender in its panels (a)-(d).

In our modeling and estimation, we treat the insurance decision of this population as being for

[12] One additional source of exemption that we cannot observe are those due to demonstrated financial hardship.
[13] See I.R.C. §5000A(g)(2), which waives criminal penalties and bars liens and levies for nonpayment.
[14] Congressional Research Service 2018

Figure 7: Individual mandate penalty by income (2018)

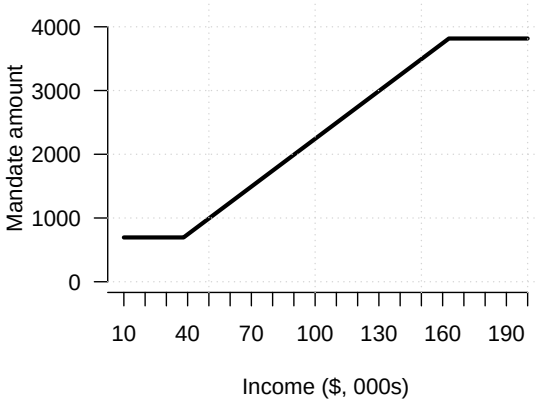
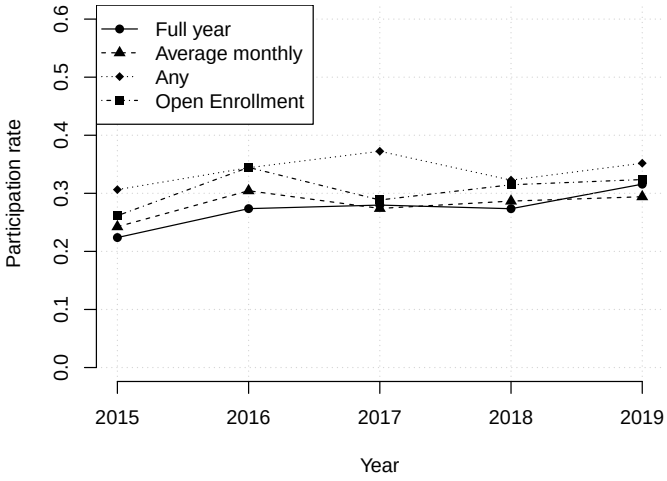


Figure 8: Comparison of definitions of participation rate

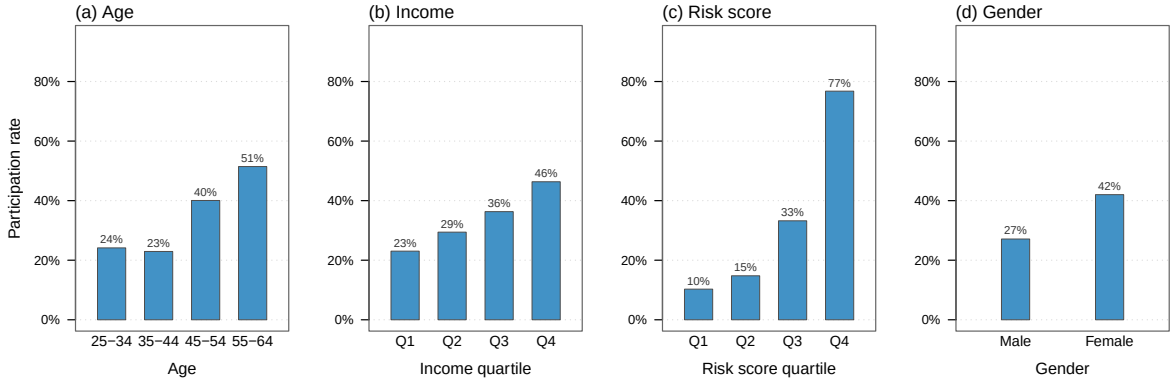


Notes: This figure displays exchange participation rates for 2015–2019 among eligible consumers under four definitions. The "Full year" and "Any" rates are computed among panelists who were exchange-insured or uninsured in every month of the year. "Full year" is the share who were insured in all twelve months whereas "Any" is the share who were insured in at least one month. The "Average monthly" rate averages, over the twelve months of the year, the share of consumers who were insured among those who were insured or uninsured in the given month (but not necessarily for the entire year). Last, "Open Enrollment" measures participation at the close of the annual Open Enrollment window, whose elected coverage takes effect by March: among consumers either insured in March or uninsured throughout January–March, it is the share insured in March.

the full calendar year. That is, we match the model's predictions for these consumers, assuming they are considering full-year decisions, to observed participation rates in MEPS for the subset of consumers who were eligible during open enrollment.

However, under the ACA, consumers who have "qualifying life events" (e.g., job loss, marriage,

Figure 9: Participation by subpopulation, 2019



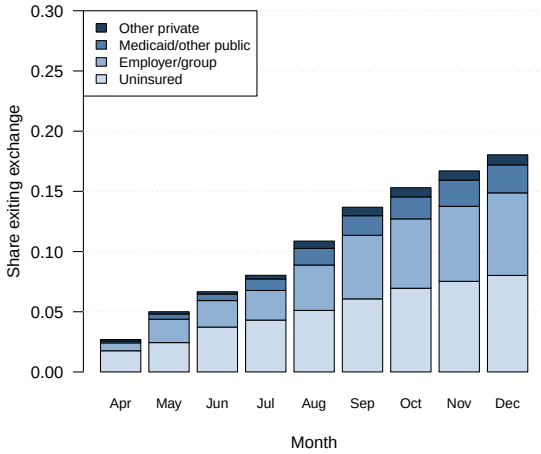
and change of residence) can enroll later in the year. As well, individuals who are initially enrolled in an ACA plan can choose to drop out later in the year and stop paying premiums. Moreover, because individuals all face the same annual deductibles and out-of-pocket maximums regardless of whether they are enrolled for the full year, the value of ACA insurance is lower for a consumer who anticipates being enrolled for less than the full year. Here we argue that, despite this fact, turnover levels in the exchange make our approach reasonable.

Figure 10 examines cumulative attrition from April to December from the insurance/uninsurance decision made during open enrollment. The figure shows that a little less than 20% of those enrolling in exchange insurance during open enrollment subsequently drop out, with about 8% to uninsurance (possibly because they run into financial troubles or simply forget to pay their premiums) and about 6% to employer/group insurance. For those choosing not to enroll during open enrollment despite being eligible, almost none subsequently enroll in exchange insurance. Overall, a little less than 25% of consumers become insured by December with the vast majority being covered by employer/group insurance or to Medicaid/Medicare insurance. Although there is clearly some movement away from the decisions made during Open Enrollment, we believe these levels to be sufficiently low to make our assumption that a consumer makes their Open Enrollment decisions based on an expectation of holding to their decision for the full year to be reasonable.

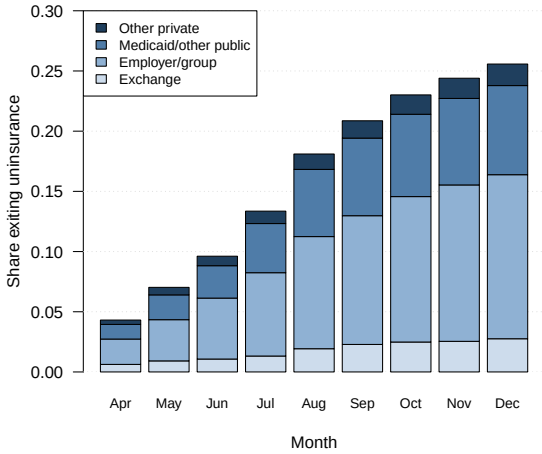
When we conduct counterfactuals in Sections 7 and 8, we focus on participation rates for this same set of consumers who were eligible during open enrollment. Figure 8 above also shows three other participation rates for the years 2015-2019 based on different samples and definitions of the participation rate. For "Full year" and "Any", the sample is limited to consumers who were either exchange-insured or uninsured in every month of the year. The "Full year" participation counts participation only when a consumer remained in ACA insurance every month of the year, while "Any" counts individuals as participating if they were enrolled for at least one month of the year regardless of when they enrolled. Comparing the two gives a sense of how much of a difference in participation rates the turnover in the exchanges can generate. The "Average monthly" curve is instead the average of monthly participation rates. For each month of each

Figure 10: Attrition from Open Enrollment decisions

(a) Attrition from exchange insurance



(b) Attrition from uninsurance



Notes: Each panel decomposes cumulative attrition from an Open Enrollment origin state by the coverage state into which consumers transition. The sample pools the 2014–2019 MEPS data, including consumers of age 25–63 and with incomes above the FPL. Panel (a) follows consumers holding exchange coverage in March; panel (b) follows consumers uninsured in every month of January through March. For each month from April through December, the height of the bar is the cumulative share of the origin group that has left the origin state by that month; each consumer who exits is held in the destination of their first departure from that month onward. "Other private" includes non-group, non-exchange private insurance.

year, we identify all panelists who were exchange-eligible in a given month—that is, who satisfied the age- and income-based conditions for eligibility and who were either exchange insured or uninsured in that month—irrespective of their exchange eligibility in other months of the year. We compute the monthly participation rate as the share of those panelists who were exchange insured, and then average across the twelve monthly participation rates for each year to get that year's annual rate. Overall, these various measures are close to each other, indicating that our focus on open enrollment participation rates in Sections 7 and 8 seems reasonable.

5 Model

We develop a model that captures the key factors that the academic and non-academic literature hypothesize as shaping participation and welfare in the ACA exchanges. In addition to adverse selection (due to community rating and private information), the model features (i) public premium subsidies, (ii) mandate penalties, (iii) uncompensated care, (iv) preference or distaste for insurance, and (v) the possibility that consumers may fail to anticipate the magnitudes of their healthcare expenses ("myopia"). We embed these forces in an equilibrium model of interactions among consumers and insurers regulated by the government.

5.1 Consumers

There is a population of consumers who may purchase health insurance in the individual market. Each consumer j is between 25 and 64 years old. In each year t, consumer j has a concave utility index $U_j(\cdot)$ over total consumption. We specify $U_j(\cdot)$ as a constant absolute risk aversion (CARA) utility function

$$U_j(c_{jt}, \gamma_j) = -e^{-\gamma_j \times c_{jt}}, \tag{1}$$

where γ_j is the consumer's risk aversion coefficient and c_{jt} is j's consumption level in year t.

Consumption is a function of the consumer's insurance choice $I_{jt} \in \{0, 1\}$ at the beginning of the year, where $I_{jt} = 1$ and $I_{jt} = 0$ respectively represent being insured and uninsured during year t:[15]

$$c_{jt} = \begin{cases} y_{jt} - p^c_{jt} - OOP^c_{jt}(I_{jt} = 1) + z_{jt}, & I_{jt} = 1 \quad \text{(Insured)} \\ y_{jt} - OOP^c_{jt}(I_{jt} = 0) - \phi \ell_{jt}, & I_{jt} = 0 \quad \text{(Uninsured)} \end{cases} \tag{2}$$

In equation (2), y_{jt} is consumer j's income during year t and p^c_{jt} is the premium net of government subsidies that is paid by consumer j if insured in year t. (When the gross premium charged by insurers for consumer j is p_{jt}, the government covers the gap between p_{jt} and p^c_{jt} as a subsidy.) $OOP^c_{jt}(\cdot)$ captures the individual's *perceived* out-of-pocket medical expenses, which depend on whether she is insured. As in Abaluck and Gruber (2011), we allow for the possibility that consumers may discount out-of-pocket expenses (relative to premiums), generating a gap between expenses as perceived by the consumer OOP^c_{jt} and *actual* out-of-pocket expenses OOP_{jt}. We formulate this difference by letting

$$OOP^c_{jt} = \mu_{jt} \times OOP_{jt}, \tag{3}$$

where $\mu_{jt} \in [0, 1]$ captures the fraction of OOP_{jt} that consumer j takes into account when making insurance choices. We will estimate μ_{jt}, allowing it to depend on income y_{jt} and age a_{jt}.

A consumer's actual OOPs are random and the reason why insurance may be socially valuable. It depends on her utilization of medical services m_{jt}, and a random element ξ_{jt}: $OOP(m_{jt}, \xi_{jt} | I_{jt}, y_{jt})$. Notice OOPs are conditional on insurance status I_{jt}, income y_{jt}. When insured, the consumer's OOPs can depend on income because of the cost-sharing reductions we described in Section 4.4. When uninsured, a consumer's income may affect her ability to access uncompensated care. Since we observe realized OOPs, we will be able to directly estimate this relationship. (We will describe the specific functional form for $OOP_{jt}(\cdot)$ that we estimate in Section 6.4.)

The term z_{jt} governs consumer tastes for insurance that are unrelated to the pecuniary elements of premiums and coverage of medical expenses. In Section 6.5, we will model this taste as depending on observable consumer characteristics, specifically a consumer's income, census

[15] Medical utilization m_{jt} in the model perfectly treats an individual's medical needs, so utility is only impacted by the share of m_{jt} an individual must pay out of pocket.

region, and year. Non-pecuniary costs include any hassle costs in signing up for exchange insurance (Shepard and Wagner 2022, McIntyre et al. 2021) as well as any politically-motivated distaste for the ACA exchanges. On the other hand, if medical services when insured are superior to those that the uninsured receive or if gaining access to uncompensated care requires costly effort, consumers may prefer insurance over and above any financial benefits it offers.

Finally, ℓ_{jt} denotes any mandate penalty that consumer j is supposed to pay in year t if uninsured. We introduce the coefficient ϕ to capture non-compliance and unobserved exemptions, as we discussed in Section 4.5.

As we will describe in more detail in Section 6, we allow for the possibility that at the time of purchasing insurance the consumer may have more information than do insurers about her likely medical needs, and hence OOP expenses. Specifically, while both insurers and consumers know the health risks implied by the consumer's past diagnoses and demographics, captured in a *risk score* λ_{jt} that we estimate in Section 6.3, consumers may observe an additional signal ε^0_{jt} about their health status.[16]

The optimal insurance decision I^*_{jt} by consumer j at time t is the choice that maximizes her expected utility:

$$I^*_{jt}(p^c_{jt}, \ell_{jt}, \mu_{jt}, z_{jt}, \gamma_j, \lambda_{jt}, \varepsilon^0_{jt}) = \arg \max_{I \in \{0,1\}} \mathbb{E}\left[U_{jt}(c_{jt}) | I, p^c_{jt}, \ell_{jt}, \mu_{jt}, z_{jt}, \gamma_j, \lambda_{jt}, \varepsilon^0_{jt}\right], \quad (4)$$

where the expectation is over the random realizations of consumption generated by the consumer's ex ante uncertain medical needs.

5.2 Equilibrium among insurers

Insurers partially protect their enrollees against uncertainty in realized medical costs m_{jt}. We assume that insurers are risk neutral and that (i) insurers are homogeneous in all aspects of their services, (ii) they comply with the regulations set in place by the government, and (iii) the premiums $\mathbf{p}^* = \{p^*_{jt}\}_{jt}$ they charge yield a common and fixed proportional markup over expected costs. We set the markup factor $\tau \geq 1$ based on the ACA's Minimum Loss Ratio regulation that stipulates that insurers pay out at least 80% of premium revenue to consumers as claims or other benefits (so $\tau = (1/0.8) = 1.25$). This markup over expected insurance costs can capture a combination of administrative costs and market power.[17]

When we consider pricing under the ACA community-rating rules, equilibrium determines the premium p_{25} for an age-25 consumer, with premiums p_{jt} for all older consumers determined by the regulated age schedule that we discussed in Section 4.2: $p_{jt}(p_{25}) = \bar{p}(p_{25}, a_{jt})$. The

[16] We allow for private information but not moral hazard (consumer's utilization m_{jt} depends directly on whether they are insured) in our analysis in the main text.

[17] Loss ratios remained near the 80% floor after the market stabilized following the ACA's early years. Fehr et al. (2020) report average individual market medical loss ratios of 82% in 2017, 70% in 2018, and 79% in 2019. The ACA's MLR regulation requires insurers to issue rebates to enrollees that offset shortfalls relative to the 80% threshold.

equilibrium premium p^*_{25} then satisfies:

$$\int \mathbb{E}\left([p_{jt}(p^*_{25}) - \tau(m_{jt} - OOP_{jt})] \cdot I^*_{jt}(p^c_{jt}, \ell_{jt}, \mu_{jt}, z_{jt}, \gamma_j, \lambda_{jt}, \varepsilon^0_{jt}) | p^*_{25}, \ell_{jt}, \mu_{jt}, z_{jt}, \lambda_{jt}\right) dj = 0, \tag{5}$$

where the expectation is over realizations of each consumer j's private information ε^0_{jt}, risk aversion γ_j, medical utilization m_{jt}, and OOP shock ξ_{jt}. In (5), we integrate over consumers who have different risk scores λ_{jt} and who condition their insurance decisions on this information. Because under community rating insurers cannot condition premiums on a consumer's risk score, consumer's decisions based on their knowledge of λ_{jt} effectively add to the adverse selection present due to ε^0_{jt} (Handel et al., 2015).[18] There may be no value of p^*_{25} such that (5) holds, in which case the market unravels. [19]

Our notion of equilibrium, in which both demand and pricing depend on market primitives and regulation, allows us to capture adverse selection: the demand model implies that high premiums discourage healthy individuals from purchasing insurance, which in turn further increases premiums through the pricing model and thus exacerbates the exodus of healthy individuals from the market. Note, however, that in addition to this classical channel, participation in the insurance market is shaped by a number of other factors: regulatory factors such as subsidies, age-adjusted community rating, and mandate penalties, and possibly also frictions such as a distaste/preference for insurance, myopic discounting of out-of-pocket expenses (Abaluck and Gruber, 2011), and uncompensated care (Finkelstein et al., 2019).

6 Estimation

6.1 Overview

We estimate the model in four steps. The first step, already described in Section 3.2, delivers a measure of medical utilization m_{jt}. The goal is to have a measure of utilization comparable across insured and uninsured, not contaminated by uncompensated care nor arbitrary list prices. The measure is based on the medical events of each individual throughout that same year. The second step is the estimation of an ACG-type of model of health risk. The individual risk has one component based on prior-year diagnostics and demographic information, denoted λ_{jt}, and a measure of the residual uncertainty over and above this score. The role of this step is to capture the health risk perceived by each panelist at the time of purchasing insurance. The third step estimates the conditional distribution of out-of-pocket expenses given utilization and income, separately for insured and uninsured consumers, which translates a consumer's health risk into the financial risks she would face if insured and if uninsured. The fourth step estimates the parameters of the consumer choice model (4) – consumer myopia and non-pecuniary tastes

[18] When we consider later risk-based pricing, a version of condition (5) must hold for each value of λ_{jt}. There is still adverse selection then, but only due to ε^0_{jt}.

[19] In principle, there may also be multiple values of p^*_{25} satisfying condition (5). Our algorithm for computing equilibria, which we describe in Appendix B, selects the lowest premium and highest participation equilibrium, although we do not find multiple equilibria in practice.

for insurance – using the prior steps as inputs. We first describe these four steps in more detail and then discuss identification.

6.2 Step 1: Utilization measure

Utilization m_{jt} is the panelist/year-level medical utilization measure constructed in Section 3.2. We refer the reader to that section for details; the rest of this section takes m_{jt} as observed.

6.3 Step 2: Health-risk model

We summarize consumer j's health status in a year by an index $h_{jt} = \lambda_j(d_{j,t-1}) + \varepsilon_{jt}$, where $d_{j,t-1}$ is the consumer's vector of pre-period diagnostic indicators (including demographics) and ε_{jt} is an idiosyncratic residual. The risk score

$$\lambda_{jt}(d_{j,t-1}) = e^{\psi_t + \phi_{\text{region}(j)}} d'_{j,t-1}\beta \tag{6}$$

is a function of $d_{j,t-1}$, as well as year and region fixed effects.[20] The residual is assumed normal $\varepsilon_j \sim N(0, \sigma_j^2)$ and possibly heteroskedastic:

$$\log(\sigma_j) = \sigma_{\text{cons}} + \sigma_\lambda \, \lambda_j. \tag{7}$$

We model utilization with a two-part structure that captures both the extensive and intensive margins:

$$m_{jt} = \begin{cases} 0, & \text{with probability } 1 - \rho_{jt}, \\ g(h_{jt}), & \text{with probability } \rho_{jt}, \end{cases} \tag{8}$$

with intensive margin

$$g(h_{jt}) = \log(1 + e^{h_{jt}}) + \nu \left[\log(1 + e^{h_{jt}})\right]^q, \tag{9}$$

and logit extensive margin

$$\rho_{jt} = \frac{e^{(\alpha + h_{jt})/\chi}}{1 + e^{(\alpha + h_{jt})/\chi}}. \tag{10}$$

The functional form $g(\cdot)$ allows the conditional mean of positive utilization to depart from a pure exponential while maintaining monotonicity in h_{jt};[21] the logit form for ρ_{jt} ties the probability of any utilization to the same latent health index, scaled by χ.

We allow for the consumer to have private information about their health status by supposing that they observe some information about the residual component ε_{jt} prior to their insurance

[20] The region fixed effects capture geographic variation in medical intensity (Finkelstein et al., 2016). The risk score λ is similar to the ACG measure used in Handel (2013) and Handel et al. (2015).

[21] Despite its appearance g's functional form is pretty straightforward. The first term is basically linear in h for positive h, and 0 otherwise. While the second term is just the power of the first term, to enable a non-linear relation. The role of the second term is to fit the heavy right tail of medical expenses.

Table 2: Health-risk model estimates

Coef.	Estimate	Std. error
α	-2.893	(0.050)
χ	0.199	(0.304)
σ_{cons}	0.950	(0.062)
σ_λ	-0.072	(0.015)
ν	0.660	(0.147)
q	4.950	(0.243)

Notes: The "Std. error" columns provides analytical GMM standard errors.

decision. Specifically, as we describe further in Section 6.5, we decompose ε_{jt} into ε_{jt}^0 and ε_{jt}^1, the former representing private information known to the consumer at the time of choice and the latter representing the remaining uncertainty to the consumer about their health status realization. (There is additional uncertainty on top of that about whether the consumer will need any positive level of medical services.)

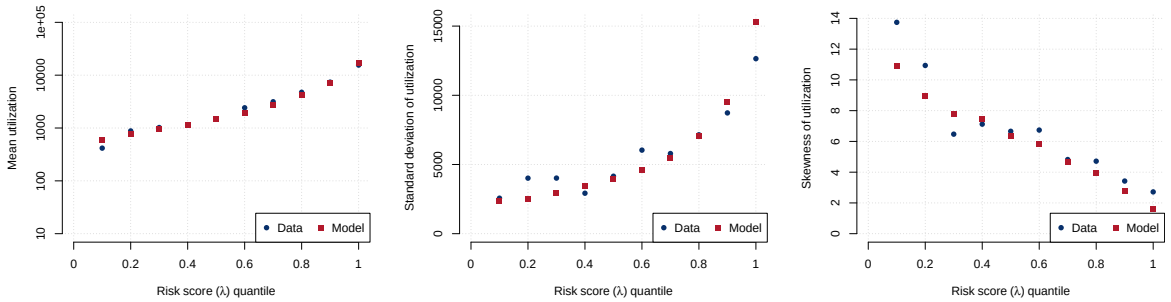
The vector of pre-period diagnostics $d_{j,t-1}$ that we consider includes 756 variables: 692 indicators for fine ICD categories, 32 indicators for coarse ICD chapters, counts of distinct diagnoses and their powers, and 26 indicators for MEPS priority conditions. It also includes a cubic in age, an indicator for female gender, and interactions of the age cubic terms and gender. To avoid overfitting, we select among them using 10-fold cross-validated Lasso at the penalty minimizing cross-validated error, leaving year, region and the cubic age-by-sex terms unpenalized. The selection procedure retains 88 of the 756 diagnostic variables.

With the selected diagnostics in hand, we estimate the parameters of (8)–(10) jointly by GMM. The moment conditions are: (i) the covariances with utilization m_{jt} of each selected regressor, comprising the retained diagnostics and the demographic terms; (ii) mean utilization by year and region; (iii) mean utilization by decile of the risk score λ_{jt}; (iv) the variance of utilization by decile of λ_{jt}; and (v) the share of consumers with strictly positive utilization within each decile of λ_{jt}. Together, these moments pin down the slopes of the two-part model, the heteroskedasticity in residual utilization, and the extensive-margin coefficients. We use two-step efficient GMM; in estimating the optimal weighting matrix for the second step, we use an identity-weighted first step.

Estimates and fit. Selected parameter estimates and their standard errors are reported in Table 2. Figure 11 displays the in-sample fit of the model along three moments of the conditional distribution of utilization within deciles of the risk score: mean, standard deviation, and skewness. The model captures the level and dispersion of utilization across the distribution of health risk well, and additionally fits well the right tail of the utilization distribution as captured by skewness.

Figure 11: Fit of health-risk model

(a) Mean (b) Standard deviation (c) Skewness



Notes: Each panel compares a moment of annual medical utilization in the data (red circles) with the same moment simulated from the estimated risk model (blue squares), within deciles of the estimated health-risk score. Points are plotted at the upper bound of each decile. Panel (a) reports the mean on a logarithmic scale. Panels (b) and (c) report the standard deviation and the skewness, computed after dropping observations above the 99.5th percentile of observed utilization, with the same dollar threshold applied to the simulated draws.

6.4 Step 3: Conditional out-of-pocket expense distributions

We estimate the conditional distribution of out-of-pocket expenses given utilization, separately for exchange-insured and uninsured consumers. We model OOPs under a given insurance status I_{jt} (which we suppress in the notation for simplicity) as follows: $OOP_{jt}(m_{jt}, \xi_{jt}|y_{jt}) = 0$ when $m_{jt} = 0$, and otherwise

$$\frac{OOP_{jt}(m_{jt}, \xi_{jt}|y_{jt})}{m_{jt}} = \Lambda\left(\kappa_0 + NCS(\log(m_{jt}); \kappa_m) + \kappa_y \log(\tilde{y}_{jt}) + \xi_{jt}\right), \tag{11}$$

where $\Lambda(x) = e^x/(1+e^x)$ is the logistic transformation, $NCS(\log(m); \kappa_m)$ is a natural cubic spline in $\log(m)$, $\tilde{y}_{jt} = y_{jt}/FPL(1)$ is income relative to the FPL, ξ_{jt} is a disturbance that is iid conditional on observables (explained below), and $(\kappa_0, \kappa_m, \kappa_y)$ are parameters.

The natural cubic spline $NCS(\cdot; \kappa_m)$ is specified to flexibly capture the relationship between the utilization level m_{jt} and the share of utilization paid out of pocket. In general, a natural cubic spline is a piecewise-cubic function whose pieces join smoothly at points called *knots* and that is linear beyond boundary knots, points chosen at extreme values of the function's domain. We specify one interior knot, at the median value of $\log(m_{jt})$, and two boundary knots at the 1st and 99th percentiles of $\log(m_{jt})$. These quantiles are computed separately for each insurance-status group.

A notable feature of the empirical distribution of the OOP ratio OOP_{jt}/m_{jt} is that it is highly right-skewed with a mean that substantially exceeds its median. To match this data pattern, it is necessary to specify a disturbance ξ_{jt} with a skewed distribution. We specify ξ_{jt} as a

log-normal random variable shifted to have a mean of zero:[22]

$$\begin{aligned} \xi_{jt} &= \tilde{\xi}_{jt} - \mathbb{E}[\tilde{\xi}_{jt} \mid m_{jt}] \\ \log\left(\tilde{\xi}_{jt}\right) &\sim N(\vartheta(m_{jt}), \zeta^2) \\ \vartheta(m_{jt}) &= \vartheta_0 + \vartheta_1 \log(m_{jt}). \end{aligned}$$

We specify $\vartheta(m_{jt})$ as a function of m_{jt} to render ξ_{jt} heteroskedastic, and hence reflect the dependence of variance in OOP_{jt}/m_{jt} on m_{jt}.[23]

We estimate the parameters $(\kappa_0, \kappa_m, \kappa_y, \vartheta_0, \vartheta_1, \zeta)$ by GMM. The moments that we specify match the following quantities between the data and the model predictions: (i) the mean, median, and width of the interquartile range by decile of utilization m_{jt} and (ii) the covariance between OOP_{jt}/m_{jt} and log relative income $\log(\tilde{y}_{jt})$. To limit the influence of outliers, we winsorize OOP_{jt}/m_{jt} at the 98th percentile within each insurance-status group, applying the same threshold to the data and to the model's predictions when computing moments. Because the estimation is done separately for exchange-insured and uninsured panelists, the conditional OOP distribution is allowed to depend not only on utilization and income, but also on coverage status.

Table 3: Estimates of OOP-ratio model

(a) Exchange insured

Coefficient	Estimate	SE
κ_0	-2.376	(0.203)
κ_{m1}	-0.624	(0.372)
κ_{m2}	-1.302	(0.210)
κ_y	0.474	(0.073)
ϑ_0	2.373	(0.359)
ϑ_1	-0.135	(0.034)
ζ	0.323	(0.085)

(b) Uninsured

Coefficient	Estimate	SE
κ_0	-1.097	(0.232)
κ_{m1}	-1.640	(0.376)
κ_{m2}	-2.258	(0.209)
κ_y	0.780	(0.092)
ϑ_0	0.738	(0.227)
ϑ_1	-0.087	(0.016)
ζ	1.048	(0.127)

Notes: The table contains estimates of the OOP model of equation (11), which we estimate separately for exchange-insured and uninsured panelists. The κ_{m1} and κ_{m2} parameters are the coefficients on the two basis functions of the natural cubic spline (NCS) in log utilization.

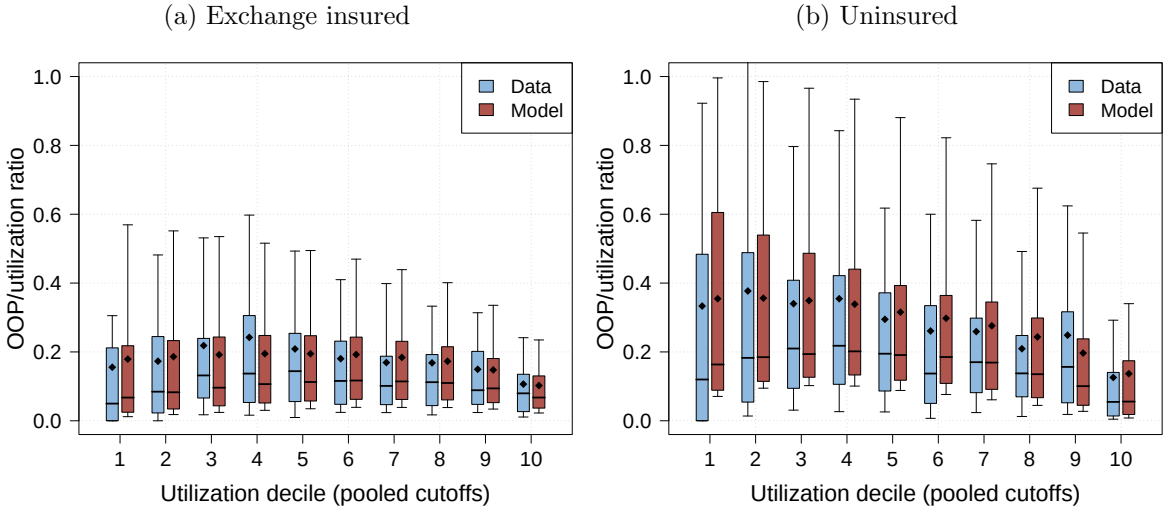
Estimates. Table 3 reports the parameter estimates. For both the exchange-insured and the uninsured, the positive κ_y parameters indicate that OOPs increase in income, as expected, due to cost sharing for the insured and uncompensated care for the uninsured. Given that the other parameters are not directly interpretable, we discuss the estimates in the context of Figure 12, which displays the distribution of the OOP_{jt}/m_{jt} ratio by decile of utilization. In particular, the figure shows the 10-90th percentile range, the interquartile range, the mean, and the median within each utilization decile and insurance status group. The figure shows that the

[22]The conditional mean of ξ_{jt} is zero by construction whereas its median conditional on m_{jt} is $\exp(\vartheta(m_{jt})) \times (1 - \exp(\zeta^2/2))$, which is negative to an extent that depends on ζ. The parameter ζ governs the skewness of ξ_{jt}, which is $(\exp(\zeta^2) + 2) \times \sqrt{\exp(\zeta^2) - 1}$.

[23]The variance of ξ_{jt} depends on $\vartheta(m_{jt})$: $\text{Var}(\xi_{jt} \mid m_{jt}) = \exp(2\vartheta(m_{jt}) + \zeta^2) \times (\exp(\zeta^2) - 1)$.

uninsured face a level of OOPs that is generally higher than that of the insured, although the difference dissipates at higher levels of utilization; within the top decile of utilization, OOPs constitute an average of 10% of utilization for the insured and 13% for the uninsured under the estimated model. Moreover, the uninsured face greater variance in OOPs than the insured. For both insurance-status groups, the mean of the OOP_{jt}/m_{jt} distribution is consistently above its median on account of skewness. The model fits the mean, median, and spread of this distribution well.

Figure 12: Fit of the distribution of the OOP ratio by utilization decile



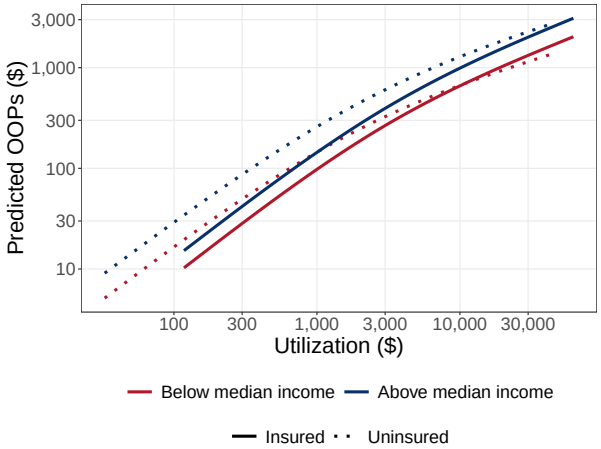
Notes: This figure juxtaposes the distribution of the ratio of OOPs to utilization in the data with the distribution implied by the estimated model, by decile of utilization, separately for exchange-insured and uninsured panelists. The utilization deciles are computed on the pooled sample rather than separately by insurance-status group. Boxes span the interquartile range, the horizontal bar marks the median, whiskers extend to the 10th and 90th percentiles, and the diamond marks the mean. The ratio is winsorized at its 98th percentile within each group, in both the data and the model's predictions. The model's distributions integrate over the estimated disturbance at each panelist's observed utilization and income. Panelists with zero utilization are dropped whereas those with zero OOPs are retained.

Figure 13 illustrates how the conditional median of OOPs varies with utilization, income, and insurance status under our estimates. Each curve plots conditional predicted OOPs as a function of utilization; the curves plot this relationship separately for below and above median income groups—particularly, conditional on the average below median and average above median income—and for insured and uninsured consumers. The figure reflects that, under our estimates, OOPs rise with income. Furthermore, it shows that the gap between insured and uninsured consumers' OOPs shrinks as utilization rises.

6.5 Step 4: Choice model

Consumer j purchases insurance according to the decision rule $I^*_{jt}(\cdot)$ that we described in Subsection 5.1. As described there, we allow μ_j to depend on a consumer's age a_{jt} and income, with the latter a function of the consumer's income relative to the FPL, $\tilde{y}_{jt}(y_{jt})$; specifically, it depends on the combination of whether the consumer is high- or low-income (above or below

Figure 13: Predicted OOPs by utilization and income



Notes: The figure plots predicted OOPs implied by the estimated model as a function of medical utilization, on logarithmic axes, for insured and uninsured consumers. Consumers are split into groups with income above and below the median, where income is measured relative to the FPL. Each line displays the median OOP amount conditional on the indicated level of utilization and on the mean income of the corresponding income group. Utilization ranges from the 1st to the 99th percentile of each insurance group's utilization distribution.

300% of the FPL) and above or below age 45, giving rise to four possible values: $\mu_{\text{low/young}}$, $\mu_{\text{low/old}}$, $\mu_{\text{high/young}}$, and $\mu_{\text{high/old}}$. Nonpecuniary taste for insurance z_{jt} is a linear function of income, census region, and year in addition to an intercept term.

We estimate the choice model by indirect inference on the sample of exchange-eligible panelists, matching participation rates and regression coefficients computed identically in the data and in simulations of the model. We decompose the residual health shock ε_{jt} (with variance σ_j^2) in the health-risk model into a component ε_{jt}^0 that the consumer observes prior to making the insurance decision and a component ε_{jt}^1 that is realized only after the decision is made, so that $\varepsilon_{jt} = \varepsilon_{jt}^0 + \varepsilon_{jt}^1$. The share of the variance of ε_j attributable to ε_{jt}^0 is a parameter ϑ_j to be estimated; it captures the information the consumer has about her future health at the time of choice, above and beyond observable demographics and prior diagnostics. Specifically, we assume that $\varepsilon_{jt}^0 \sim N(0, \sigma_{j0}^2)$ and $\varepsilon_{jt}^1 \sim N(0, \sigma_{j1}^2)$ where $\sigma_{j0}^2 = \vartheta \times \sigma_j^2$ and $\sigma_{j1}^2 = (1 - \vartheta_j) \times \sigma_j^2$. We specify $\vartheta_j \in [0, 1]$, which governs the extent of selection on unobservables, to depend on the consumer's gender (male or female) and age group (above or below age 45), thus taking one of four values: $\vartheta_{\text{m/young}}$, $\vartheta_{\text{f/young}}$, $\vartheta_{\text{m/old}}$, and $\vartheta_{\text{f/old}}$. The rationale for allowing ϑ_j to vary by gender and age is that a need for care related to pregnancy and childbirth—which is concentrated among younger women—constitutes an important source of private information. Indeed, among exchange-eligible women aged 25–44, pregnancy, childbirth, and fertility care account for 19.7% of total medical expenses, but less than 0.1% for older women and for men.[24]

[24]We compute this share from the MEPS condition–event link files, which assign to each medical event the corresponding panelist's associated medical conditions. In making the calculation, we treat expenses for events linked to more than one condition by splitting the expenses equally across these conditions; assigning instead the full amount to any event carrying a pregnancy-related link moves the share by 0.3 percentage points.

For each panelist we draw ten values of ε^0_{jt} and ten values of risk aversion γ_j to evaluate (4) by simulation. We calibrate the distribution of γ_j as discussed below.

Moments. We estimate the model parameters $(\{\mathit{ɔ}\}, \{\mu\}, \{\psi\}, \phi)$ on the second-year insurance decisions in MEPS (using first-year diagnostics as predetermined) by matching the following data features:

1. Chiappori–Salanié (Chiappori and Salanié, 2000) style selection coefficients: the coefficients b_1 in the Poisson quasi-maximum-likelihood regression

$$\log\left(\mathbb{E}[m_{jt} \mid \mathbf{1}\{\text{insured}\}_{jt}, x_{jt}]\right) = b_0 + b_1\, \mathbf{1}\{\text{insured}\}_{jt} + x'_{jt}\beta, \tag{12}$$

   with x_{jt} controlling for the risk score λ_{jt}, log relative income, year, and region. To limit the influence of the utilization tail, we winsorize m_{jt} at its survey-weighted 95th percentile, applying the same threshold to the data and to the model's simulated utilization. We estimate this regression separately for four age/gender groups (men and women, below and above age 45), yielding four moments. These coefficients target the degree of selection on the component ε^0_{jt} of the health shock that is private to the consumer at the time of choice, as governed by the group-specific $\mathit{ɔ}$ parameters;

2. The coefficient on the mandate penalty in the least-squares regression

$$\mathbf{1}\{\text{insured}\}_{jt} = b_0 + b_1\, \ell_{jt} + x'_{jt}\beta + u_{jt}, \tag{13}$$

   with x_{jt} controlling for year, age, income, and the risk score λ_{jt}; this moment targets the non-compliance coefficient ϕ on the mandate;

3. Participation rates by income group, census region, and year, which target the income/region/year-based tastes in z_j;

4. Participation by quartiles of the risk score λ_j within each of four cells defined by income and age: high- versus low-income (above or below 300% of FPL) interacted with older versus younger (above or below age 45). These sixteen moments target the cell-specific myopia coefficients, which govern the relationship between λ_j and participation.

Appendix D reports the value of these moments on the estimation sample and as predicted by the estimated model.

Risk aversion. Handel et al. (2015) use data on medical claims and insurance plan choices from a large employer to estimate a model of CARA risk preferences, and we use their estimates in our analysis. In particular, we adopt the "Robustness 3" specification of their Table 2, under which risk aversion takes the form $\gamma_{ja} = \gamma_0 + \gamma_{\text{age}} a + \tilde{\gamma}_j$, where a designates age and $\tilde{\gamma}_j \sim_{\text{iid}} N(0, \sigma^2_\gamma)$ is drawn once per consumer; their estimates are $\gamma_0 = 2.54 \times 10^{-4}$, $\gamma_{\text{age}} =$

3.99×10^{-6}, and $\sigma_\gamma = 1.20 \times 10^{-4}$.[25] We use their Robustness 3 specification because their other specifications involve variables that are not present in MEPS. The estimates imply that a 45-year-old consumer with mean risk aversion would be indifferent between accepting and rejecting a 50–50 gamble in which they win \$1,000 or lose \$695.

6.6 Identification

The parameters of the health-risk model in Step 2 and the conditional OOP distributions in Step 3 are identified directly by the moments of the utilization and OOP data they target. We therefore focus the discussion here on the choice-model parameters of Step 4.

Selection on private information. The share ϑ_g of the health-shock variance that consumers in age/gender group g observe at the time of choice is identified by the group-specific Chiappori–Salanié coefficients in (12). In our model, utilization does not respond to insurance status, and tastes and risk aversion are independent of the health shock. Under these assumptions, a positive coverage coefficient conditional on the risk score λ_{jt} and the other controls can arise only through selection on the privately observed component ε^0_{jt}: consumers who observe a high draw anticipate greater utilization and are more likely to insure. The magnitude of each group's coefficient therefore identifies that group's ϑ_g, with a larger coefficient implying that a larger share of the health-shock variance is in the consumer's information set at the time of the insurance decision.

Myopia. The myopia parameters $\{\mu\}$ are identified from the gradient of participation with respect to the risk score λ_j within demographic cells. Under expected utility with $\mu = 1$, the participation–health gradient is determined entirely by the conditional OOP distributions, risk aversion, and the premium schedule, all of which are taken from prior steps or external sources. Departures from the predicted gradient, by demographic cell, identify the cell-specific μ. Intuitively, when μ is small the consumer's perceived payoff from insurance depends weakly on health risk, flattening the participation–λ profile; when μ is close to one, the profile is steeper.

Mandate response. The coefficient ϕ is identified from the responsiveness of the participation rate to variation in the individual-mandate penalty ℓ_{jt} across income brackets and years, after partialling out region and year fixed effects.

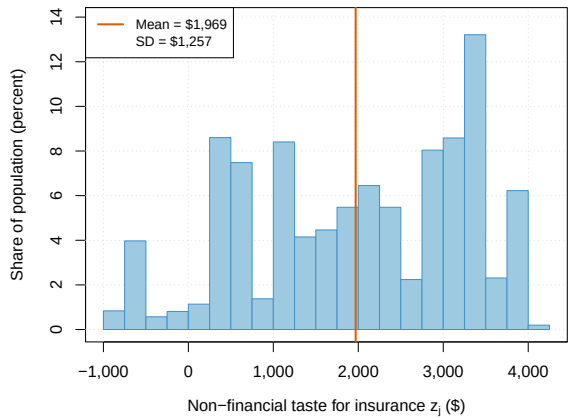
Tastes for insurance. The income, region, and year coefficients in z_j are identified from the average participation rates within the corresponding cells, after the other parameters have accounted for variation due to λ, premiums, and the mandate.

6.7 Results

Table 4 reports the parameter estimates of the choice model. Three findings deserve emphasis.

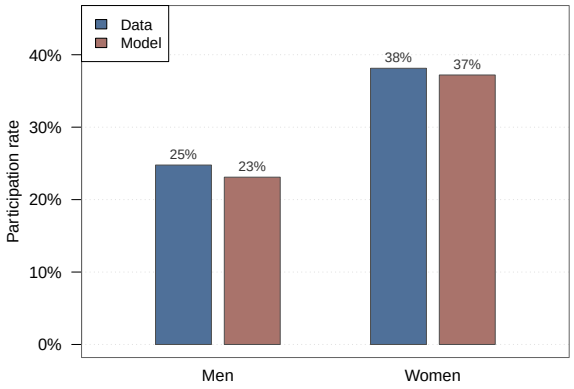
[25] Handel et al. (2015) allow managers and non-managers at the firm they study to have different intercepts γ_0 in their Robustness 3 specification. We use the intercept for non-managers in our analysis.

Figure 14: Distribution of non-financial tastes for insurance z_j



Notes: This figure plots the distribution of z_j, the variable capturing non-financial tastes for insurance, in the choice model sample.

Figure 15: Participation fit by gender



Notes: This figure compares observed and model-predicted exchange participation rates, pooled across ages, separately for men and women in the choice model sample.

First, the model implies that the share ϑ_j of the variance of the health shock that is observed by the consumer at the time of choice ranges from 0.13 to 0.38. Notably, it is highest for young women who may hold private information related to pregnancy. The estimated magnitudes are consistent with substantive but not overwhelming adverse selection on information observed by consumers but not insurers.

Second, tastes for insurance, captured by z_j, are increasing in income and grow over time. Both patterns are consistent with the existing literature: liquidity-constrained consumers may face binding upfront-premium constraints, and exchange participation grew over the sample as information about the exchanges spread and operational frictions in the exchanges declined. Notably, for most consumers non-pecuniary factors *favored insurance purchase* over and above the financial benefits – see Figure 14 which shows the distribution of these non-pecuniary effects among consumers. Note that z_j has a mean of \$1,1969 with a standard deviation of \$1,257. This is consistent with the medical care consumers receive when insured being of higher quality and/or there being effort costs involved in securing uncompensated care when uninsured. In Section 8.2, where we consider the effects of replacing uncompensated care with more formal (but less comprehensive than ACA) insurance, we will do so in part by eliminating this non-pecuniary preference for ACA insurance.

Third, the mandate coefficient as estimated on the basis of revealed preference is 18%, compared to 16% of expected mandate revenue actually collected in 2015 as computed based on an accounting approach.

Fit. Figures 15 and 16 compare the participation rates predicted by the model to those observed in the data, by gender, age and by decile of the risk score λ_j. Note that participation by gender is not targeted by an estimation moment, yet is matched well by the estimated model.

Table 4: Choice-model parameter estimates

Parameter	Estimate
$\vartheta_{\text{male/young}}$	0.24
$\vartheta_{\text{male/old}}$	0.13
$\vartheta_{\text{female/young}}$	0.38
$\vartheta_{\text{female/old}}$	0.14
$\mu_{\text{low/young}}$	0.45
$\mu_{\text{low/old}}$	0.38
$\mu_{\text{high/young}}$	0.36
$\mu_{\text{high/old}}$	0.59
$\psi^{\text{inc}}_{\text{1.50-2.00}}$	0.41
$\psi^{\text{inc}}_{\text{2.00-2.50}}$	1.20
$\psi^{\text{inc}}_{\text{2.50-3.00}}$	2.07
$\psi^{\text{inc}}_{\text{3.00-3.50}}$	2.57
$\psi^{\text{inc}}_{\text{3.50-4.00}}$	2.83
$\psi^{\text{inc}}_{\text{4.00-5.00}}$	3.12
$\psi^{\text{inc}}_{\text{5.00+}}$	3.07
$\psi^{\text{year}}_{\text{2016}}$	0.87
$\psi^{\text{year}}_{\text{2017}}$	1.14
$\psi^{\text{year}}_{\text{2018}}$	1.27
$\psi^{\text{year}}_{\text{2019}}$	1.63
$\psi^{\text{region}}_{\text{Northeast}}$	-0.72
$\psi^{\text{region}}_{\text{Midwest}}$	-0.92
$\psi^{\text{region}}_{\text{South}}$	-0.94
$\psi^{\text{region}}_{\text{West}}$	-0.94
Mandate coefficient	0.18

The fit by age is likewise strong.

6.8 Out-of-sample check: Enhanced subsidies

Another test of the model is provided by the enhanced premium subsidies passed into law in 2021 as part of the federal response to the COVID-19 pandemic. The enhanced subsidies considerably expanded the income range over which consumers received subsidies and reduced the post-subsidy premium for those who already received them, as summarized in Table 5. According to the Kaiser Family Foundation, participation on the ACA exchanges roughly doubled by 2025 in response to this policy. Feeding the new subsidy schedule into our estimated model – holding all other parameters fixed – yields a predicted increase in participation of roughly 2.02×, in line with the realized response.

6.9 The role of hidden selection in OOP distributions

A final piece of evidence on the importance of consumer-private information about ε_{j0} comes from comparing the distribution of OOPs for insured and uninsured consumers conditional on the observable risk score λ_j and income y_j. Panel A of Figure 17 shows that, conditional only

Figure 16: Choice-model fit: participation by age and by decile of the risk score, data versus model

(a) Participation by age

(b) Participation by risk score λ_j

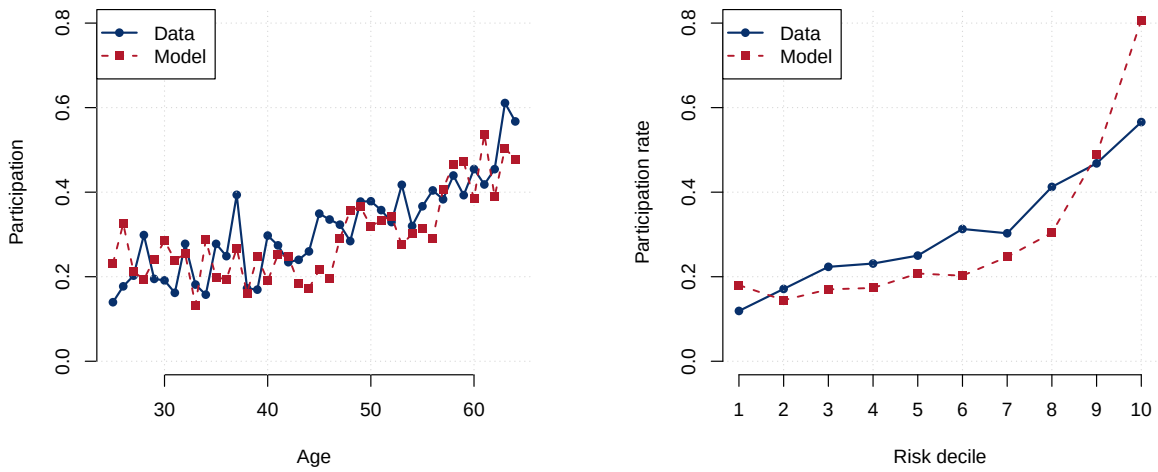


Table 5: Subsidized premiums as a share of income: baseline vs. enhanced

Income (% of FPL)	Baseline	Enhanced
Under 150%	2–4%	0%
150–200%	4–6.3%	0–2%
200–250%	6–8.05%	2–4%
250–300%	8–9.5%	4–6%
300–400%	9.5%	6–8.5%
Over 400%	No cap	8.5%

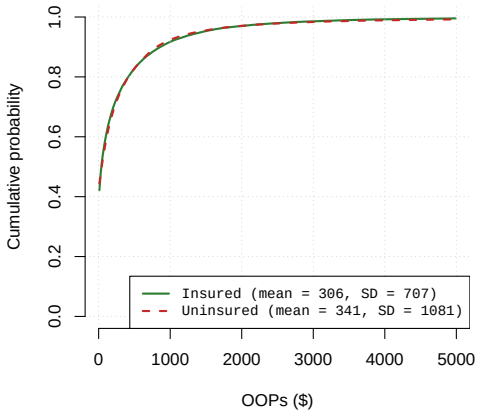
on λ and y, the OOP CDFs of insured and uninsured consumers appear *reversed* relative to the textbook prediction: at the 50th percentile of λ and of y, the uninsured appear to face a weakly more favorable OOP distribution than the insured.[26] Conditioning in addition on $\varepsilon^0_{jt} = 0$–thereby holding constant the component of the health shock observed by the consumer at the time of choice—restores the expected ordering, as shown in Panel B of Figure **??**: the insured face a more favorable OOP distribution. The same pattern obtains at the 75th and 95th percentile of λ. This exercise illustrates why ignoring the consumer-private component of the health shock can generate a misleading picture of the value of insurance, and motivates our inclusion of ε^0_{jt} in the choice model.

[26]The median household income is 298% of the FPL.

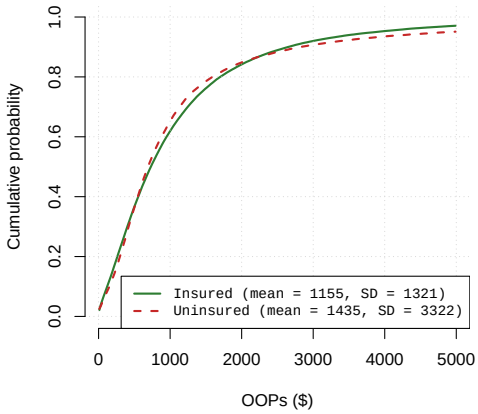
Figure 17: OOP distributions by risk score, insured versus uninsured, with and without the correction for hidden selection

Panel A: conditioning on λ_j and y_j only

(a) 50th percentile λ

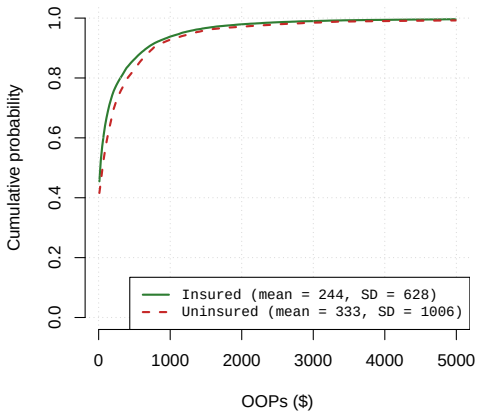


(b) 95th percentile λ

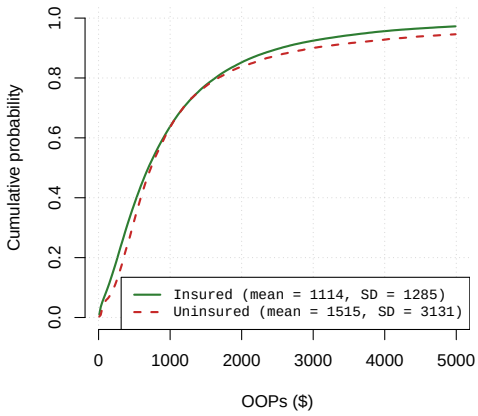


Panel B: additionally conditioning on $\varepsilon_{j0} = 0$

(c) 50th percentile λ



(d) 95th percentile λ



Notes: Each panel plots the simulated distribution of OOPs for the insured and the uninsured, holding the risk score at the indicated percentile and income at its median. Panel A conditions on the risk score and income only, so the two distributions differ both because coverage differs and because of selection on private information. Panel B additionally sets the private-information shock to zero, removing selection on private information.

7 Analysis of participation

In this section, we analyze the relative importance of the various factors affecting participation in the ACA exchanges.

7.1 Construction of population used in counterfactual analysis

We compute exchange equilibria on the subsample of MEPS panelists who meet the age- and income-based exchange eligibility criteria of Section 4.1 and who are either exchange-insured or uninsured according to our Open Enrollment classification. We use second-year observations from 2019 panel and fix the policy environment (i.e., the premium-subsidy schedule, cost-sharing rules, and the absence of a mandate penalty) at its 2019 level under the baseline regime that we analyze, using the 2019 FPL in computing subsidy eligibility. To introduce heterogeneity in risk preferences, we replicate each panelist 10 times and assign each replicate a unique risk-aversion coefficient draw independently from the distribution outlined in Section 6.5.

Our counterfactual eliminates uncompensated care by making uninsured consumers bear the full cost of their medical utilization: we set their OOPs equal to their utilization m_{jt}. Insured consumers are unaffected.

7.2 Impediments to participation

In this section, we examine the implications of our estimates for the role that three frictions—myopia, uncompensated care, and adverse selection—play in reducing exchange participation among eligible consumers. To do so, we simulate what equilibrium would be in each of their absences. The results are in Table 6.

Table 6: Participation rates (%) upon friction elimination

Regime	Current regime (no penalty)	Risk-based pricing
Baseline	32.0	75.4
No myopia	58.3	90.5
No uncompensated care	84.1	91.1
Neither	94.7	94.6

Notes: This table provides participation rates in percentage terms under regimes various combinations of the myopia and uncompensated care frictions. The first column provides rates for the current regime with a community rating whereas the second column provides rates under risk-based pricing.

The column labeled "Current regime (no penalty)" in Table 6 describes how the removal of each of myopia and uncompensated care affect exchange participation rates starting from a baseline ACA regime under community-rating featuring current levels of premium subsidies, cost-sharing, and no mandate penalty. In this baseline, 32.0% of eligible consumers purchase exchange insurance.

Eliminating myopia—setting $\mu_j = 1$ so that consumers weight OOPs dollar-for-dollar with premiums—raises participation from 32.0% to 58.3%. The corresponding figure for uncompensated care is considerably higher, at 84.1%, and removing both frictions together yields 94.7%. Thus, among the frictions operating within the ACA's community-rating regime, uncompensated care is the primary impediment to exchange participation.

The second column in Table 6 labeled "Risk-based pricing" describes a similar exercise conducted under the alternative assumption of risk-based pricing, which removes the adverse selection caused by community rating. We assume firms can price based on the observed type, λ, but naturally cannot price the unobserved type ε_0. Thus, adverse selection on unobservables still persists. Were there no frictions at all (including no markup, no adverse selection, and no distaste for insurance, as in the classic textbook insurance setting), participation in the presence of premium subsidies and cost sharing would be 100%. Table 6 shows that with risk-based pricing we estimate that participation in the absence of both myopia and uncompensated care would be 94.6% (note the model still features insurer markups, a distaste for insurance among some consumers, and private information). This 94.6% participation rate compares to a rate of 75.4% when both frictions are present, 90.5% absent myopia but with uncompensated care, and 91.1% with myopia but without uncompensated care. Two conclusions emerge from this. First, myopia and uncompensated care continue to depress participation under risk-based pricing, though less strongly than under community rating: eliminating them raises participation by 15.1 and 15.7 percentage points respectively, against 26.3 and 52.1 points under community rating. The comparison with adverse selection itself points the same way: at the baseline, eliminating uncompensated care is worth 52.1 percentage points, more than the 43.4 points gained by moving to risk-based pricing. Second, absent uncompensated care and myopia, adverse selection due to the ACA community rating would not weigh on participation at all: we estimate participation of 94.7% under community rating against 94.6% under risk-based pricing once the two frictions are gone.[27]

7.3 Impacts of ACA policies on participation

The ACA included several provisions aimed at encouraging exchange participation, the most notable among them being premium subsidies and the individual mandate penalty. In this section, we investigate the impacts of these two policies in offsetting the identified frictions.[28]

Table 7 presents the results. The current situation with subsidies but no mandate is shown in the second row of the "Baseline" column and has a participation rate of 32.0% (matching Table 6). Adding the mandate penalty raises participation only by 6.6 percentage points. However, subsidies have a much larger impact on participation than the mandate penalty. The market almost fully unravels without the subsidies even if there is a mandate penalty. Notice that at their current magnitude, the premium subsidies more than offset the friction from uncompensated care: the second column labeled "No uncomp. care" shows that if uncompensated care is not available but subsidies were also removed, participation would decline from 84.1% to 19.4%—below the 32.0% observed at the baseline, where uncompensated care is present but subsidies are paid.

[27] Adverse selection could still be a large problem for the *choice* between plans of different actuarial values, as shown in Handel et al. (2015).

[28] While cost-sharing reductions (CSRs) also aimed to improve exchange participation by lowering out-of-pocket costs for low-income enrollees, their budgetary footprint is far smaller than that of premium subsidies—roughly \$10–15 billion annually versus \$90–110 billion Congressional Budget Office 2023, 2017.

Table 7: Participation rates (%) by policy regime

Policy regime	Baseline	No uncomp. care	Risk-based pricing
Mandate and subsidies	38.6	89.6	72.1
Subsidies without mandate	32.0	84.1	75.4
Mandate without subsidies	0.1	20.8	47.6
Neither	0.1	19.4	56.4

Notes: this table provides participation rates in percentage terms under regimes with various combinations of ACA participation inducement policies. These policies are the individual mandate penalties (at 2018 levels) and premium subsidies. The first column provides rates under the baseline regime with uncompensated care and a community rating. The second column provides rates under a community rating regime without uncompensated care. The third column provides rates under risk-based pricing with uncompensated care.

Figures 18 and 19 depict respectively the equilibria without and with subsidies (the latter being our baseline) and are very instructive to understand the role of ACA subsidies. Figure 18, as in Einav et al. (2010), displays each individual in our sample, ordered from highest to lowest willingness to pay for insurance, accounting for the estimated frictions, including myopia and uncompensated care. Willingness to pay for insurance ranges from almost \$20k to virtually nothing. The red line displays the average cost of insuring different shares of the population. The costs are uniformly higher than the estimated willingness to pay, illustrating why the market fully unravels absent subsidies.

Figure 19 presents instead the post-subsidy willingness to pay. Namely, the market premium at which an individual is willing to purchase insurance given that they do not pay the full premium. The first thing to notice, on the left side of the figure, is that over 30% of the population enjoy a premium cap that is lower than how much they value insurance. These individuals participate regardless of the market premium (they are willing to accept an arbitrarily large premium). The horizontal part of the red line reflects the average cost of insuring these individuals (plus the markup τ). The rest of the graph presents individuals whose willingness to pay is less than their premium cap. The ACA subsidies (more precisely, the premium cap) do not directly affect these individual's decision of whether to participate. In the depicted equilibrium, some of them participate in the market. Adverse selection means that average costs for this group decrease as more of these individuals join the market, but overall average costs are an average of this group and the 30% who always participate and so initially rise because of less healthy individuals who, despite not being subsidized, are willing to join. Subsidies therefore induce participation in two ways: a group of individuals are led to participate regardless of the premium, while others despite not being subsidized are willing to purchase because average costs are driven down by the subsidized participants.

Figure 18: Unravelling absent subsidies

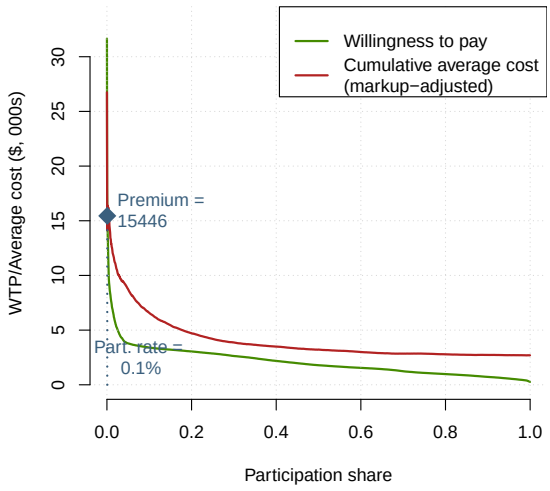
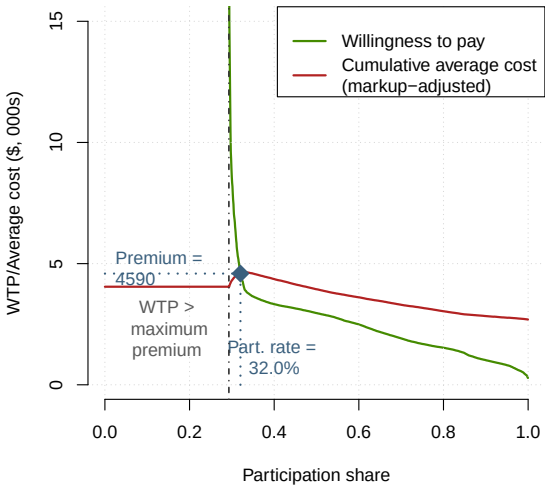


Figure 19: Exchange equilibrium, baseline



8 Welfare analysis of policy changes

In this section, we examine two possible policy changes. In the first, we study the *actual* expanded subsidies that existed between 2021 and 2025. As noted above, our model predicts well the participation changes that accompanied this policy change. Here we examine its welfare consequences. In the second, we consider the positive and normative effects of replacing the current inefficient uncompensated care system with formal insurance whose coverage is, on average, the same as current uncompensated care but which is more certain and has the same non-pecuniary benefits/costs (our z's) as exchange insurance.

Before doing so, we first introduce our welfare measures.

8.0.1 Welfare measures

We evaluate a potential change in policy from the status quo using a money-metric measure of consumer welfare change together with a corresponding social welfare measure. For the consumer measure, we compute the compensating constant annual payment from a consumer to the government that, combined with the alternative policy, would leave the consumer with the same lifetime discounted expected utility as in the status quo. Our aggregate welfare change combines these compensating payments for consumers with the changes in government revenue, insurer profits, and uncompensated care to get an aggregate welfare change.

In calculating a consumer's compensating payment we take the perspective of an unborn consumer whose age-25 risk aversion is $\gamma_{j,25}$ but whose realizations of income and health are uncertain. Specifically, let c^0_{ja} denote consumer j's (uncertain) consumption at age a under her optimal age-a insurance choice in the equilibrium associated with the status quo regime and let c^1_{ja} denote the same thing under the alternative policy regime. The compensating payment T_j

for consumer j whose age-25 risk aversion is γ_{ja} satisfies

$$\sum_{a=25}^{64} \delta^{a-25}\, \mathbb{E}[u_{ja}(c_{ja}^1 - T_j)|\gamma_{j,25}] = \sum_{a=25}^{64} \delta^{a-25}\, \mathbb{E}[u_{ja}(c_{ja}^0)|\gamma_{j,25}], \tag{14}$$

where $u_{ja}(\cdot)$ is CARA utility with risk-aversion parameter γ_{ja}, myopia parameter μ_{ja}, and non-pecuniary preferences z_{ja}. We set the discount factor δ to be 0.975. The expectations compute the consumer's optimal insurance decision at each age, integrating over the consumer's distribution of possible risk-score and income realizations at that age. As we describe in more detail in Appendix C, the consumer's risk-score trajectory is stochastic, which we estimate using our MEPS panelist data. Given our focus on insurance of health risks and not income risks, we specify a simple income transition process: we assume that a consumer's position in their cohort's income distribution stays constant throughout their life, so their income trajectory is pinned down by the stochastic realization of their income at age 25.

The change in per capita consumer welfare is the average of these compensation amounts across the population,

$$T = \mathbb{E}[T_j]. \tag{15}$$

We arrive at the change in aggregate per-capita welfare ΔW by adding T to the per-capita increase in profits accruing to insurers ΔS^i , the per-capita improvement in the government's budget ΔS^g, and the per-capita reduction in uncompensated care costs ΔS^u:[29]

$$\Delta W = T + \Delta S^i + \Delta S^g + \Delta S^u. \tag{16}$$

Several observations about this welfare measure deserve mention. First, the changes in policy we consider in this section can affect not only the insurance of health risks, but also the insurance of age-25 income risk (i.e., they can be redistributive). In addition, they can alter the smoothing of consumption across time. Our welfare measure captures all of these effects. Second, how we construct the compensating payments can matter for the welfare change by determining in which states of the world (with which marginal utilities of consumption) compensating payments are made.[30] Finally, since equilibrium participation is always computed under consumers' myopic behavior, the welfare measure depends on the stance the planner takes toward myopia—that is, whether myopic behavior should be corrected or respected. Here we take a paternalistic view and replace the perceived out-of-pocket expense OOP^c with the true expense OOP^c/μ when calculating the consumption levels c_{ja}^0 and c_{ja}^1 in (14).

[29] Here we treat insurer markups as representing pure profit – i.e., we assume that there are no administrative costs that are changed by the policy.

[30] Compensating payments need not be constant over time or over states of the world; constant compensation is only one possible form they could take.

8.1 Expanded subsidies

Table 8 shows the welfare effects of the enhanced subsidies that were in place from 2021-2025.[31] The enhanced subsidies raise exchange participation from 32.0% to 64.8%, causing the age-25 premium to fall from $4,590 to $3,413. Per-capita ex-ante consumer welfare rises by $918 per year. While the government deficit increases by an average of $595 per consumer per year, insurer profit rises by an average of $198 and average uncompensated care costs fall by $774. Overall, the enhanced subsidies raise aggregate welfare by $1,294 per consumer per year. It would take a cost of government funds equal to 217% to make enhanced subsidies reduce aggregate welfare.

Table 8: Effects of enhanced subsidies

	Baseline	Enhanced subsidies
Participation rate (%)	32.0	64.8
Age-25 premium ($)	4590	3413
Consumer welfare ($)	0	918
Insurer profit ($)	527	725
Government budget ($)	-1851	-2446
Uncompensated care ($)	1382	608
Total welfare ($)	-2706	-1412
Change vs baseline ($)	–	1294

Notes: The table reports per-capita outcomes by regime, holding the estimated parameters fixed. "Baseline" is community rating with current premium subsidies and no mandate penalty. The "Enhanced subsidies" regime departs from the baseline by instead featuring the 2021–2025 enhanced premium subsidy schedule. Consumer welfare is the paternalistic (myopia-correcting) measure, normalized to zero at the baseline. "Total welfare" is social welfare, equal to consumer welfare plus insurer profit plus the government budget less uncompensated care costs.

8.2 Formalizing uncompensated care

The literature has typically treated uncompensated care as a friction that undermines formal insurance. Yet it can equally be viewed as a valuable, if informal, form of *social insurance*, shielding the uninsured from the full cost of medical care. Its optimal level depends on its efficiency, which is plausibly lower than that of formal insurance for two reasons: (i) uncompensated care may require effort to secure and may be of lower quality, and (ii) it is riskier, as we documented in Section 3. This raises a natural question: what would happen if we formalized uncompensated care, delivering the same average level of protection more efficiently? Such improvements would directly raise the welfare of those who are now uninsured but, by making current ACA exchange insurance policies less attractive, could undermine exchange participation to some degree, leading to premium increases for those policies and reducing welfare for

[31] Our analysis should be viewed as capturing the welfare effects once participation fully adjusts; in reality, it took several years for participation to reach its 2025 level after the enhanced subsidies were enacted.

consumers now covered under those policies.[32]

We formalize uncompensated care through two changes, each of which makes uninsurance resemble a formal plan. First, we endow uninsurance with the non-pecuniary benefit z_j that insurance confers. In the 2019 population on which we simulate counterfactuals, z_j has a mean of \$2,583 and a standard deviation of \$1,093. Second, we adjust the distribution of the uninsured OOP shock ξ_j for each consumer j such that, conditional on consumer j's utilization level m_j and income y_j: (i) mean uninsured OOPs are unchanged and (ii) the variance of OOPs when uninsured equals that of OOPs when insured. We implement these changes by replacing ξ_j with $\xi_j' = A_j + B_j\xi_j$, where the parameters $\{(A_j, B_j)\}_j$ are chosen to ensure conditions (i) and (ii) above hold for each consumer j. This change reduces the mean standard deviation of uninsured OOPs conditional on utilization and income from \$908 to \$526.[33] Table 9 reports the effects of this change.

Table 9: Welfare effects of formalizing uncompensated care

	Baseline	Formal UC
Participation rate (%)	32.0	0.0
Age-25 premium (\$)	4590	–
Consumer welfare (\$)	0	1052
Insurer profit (\$)	527	0
Government budget (\$)	-1851	-3464
Premium subsidies (\$)	-1851	0
Formal UC (\$)	–	-3464
Uncompensated care (\$)	1382	–
Total welfare (\$)	-2706	-2411
Change vs baseline (\$)	–	294

Notes: per-capita outcomes from formalizing uncompensated care, holding estimated parameters fixed. "Formal UC" both (i) endows uninsurance with the non-pecuniary benefit of insurance and (ii) sets the uninsured out-of-pocket variance equal to the insured level. Consumer welfare is the paternalistic measure, normalized to zero at the baseline. In the formalization column, the cost of the now-formalized care is counted in the government budget, so the uncompensated-care line is reported only for the baseline.

Formalizing uncompensated care in this way leads the exchange to unravel completely, but nonetheless enhances welfare. The exchange collapses on account of the enhanced appeal of uninsurance, which further limits consumer willingness to pay for insurance. The formalized risk protection and non-pecuniary benefits of uninsurance, however, are of substantial value to consumers: per-capita consumer welfare rises by \$1,052. The fiscal consequences run in opposing directions. The government assumes the \$3,464 per-consumer cost of the formalized coverage, care that costs \$1,382 at the baseline, when it was borne by providers and charities rather

[32] In late 2018 the first Trump administration redefined allowable "short-term limited duration" health plans from a maximum of three months duration to 12 months with extensions up to three years; critics worried that they would undermine ACA exchange participation. See "'Short-term' Health Insurance? Up to 3 Years Under New Trump Policy," *New York Times*, August 1, 2018.

[33] Thus, we model the market as having a free lower-coverage plan (the formal uncompensated care) and a premium plan that consumers can instead opt to purchase, paying the full cost of their care.

than solely by the government. However, it no longer spends $1,851 on exchange subsidies. Therefore, its net spending rises by $1,613. On aggregate, total social welfare rises by $294 absent an efficiency cost of public spending. Because the policy raises government expenditures, a marginal cost of public funds above 18% would undo the gain.[34]

Furthermore, our welfare calculation treats the non-pecuniary benefit z_j as costless to extend to the uninsured. Whether that is correct depends on what z_j stands for. To the extent that it reflects the hassle and uncertainty of obtaining uncompensated care, formalization could deliver the benefit at little resource cost. To the extent that it instead reflects the higher quality of medical care received by the insured, conferring it on the formalized system would likely require real resources. The welfare gain leaves little room for such costs: a resource cost of providing z_j equal to just 11% of its value to consumers would eliminate the gain entirely, even before any excess burden of the public funds required to finance it.

9 Conclusion

We developed and estimated an equilibrium model of exchange insurance markets featuring frictions that rationalize limited uptake of ACA exchange insurance. The primary conclusion of the analysis is that uncompensated care is the main driver of limited insurance take-up. Uncompensated care plays an even larger role in depressing participation than adverse selection in an insurance market in which risk-based pricing is prohibited. Myopia is present but plays a more modest role, depressing participation by 26.3 percentage points in the baseline specification, and 15.1 under risk pricing. This highlights the interaction between myopia and adverse selection: myopia leads healthier individuals to drop coverage, thus exacerbating adverse selection under a community rating.

We find that premium subsidies at ACA levels have a larger impact on participation than the pre-2019 ACA mandate penalty: even with the individual mandate in place, the exchange market unravels upon the withdrawal of subsidies. This is due to the profound impact of uncompensated care on consumers' willingness to pay for exchange insurance. Absent subsidies the community-rating market virtually unravels. Removing uncompensated care restores participation only to 19.4%, still below the 32.0% that the subsidized baseline sustains; risk-based pricing restores more, 56.4%, but only by abandoning community rating altogether. Within the ACA's own pricing rules, then, it is the subsidies that sustain the exchange.

We additionally analyze the welfare effects of enhanced subsidies and of formalizing uncompensated care into a form of public insurance. First, subsidy enhancement raises aggregate welfare by $1,294 per consumer per year, a gain that survives any plausible cost of public funds. This gain reflects that the enhanced subsidy limits adverse selection and reduces the externality of uncompensated care that the uninsured exert upon society. Formalizing uncompensated care,

[34] We have not provided an analysis of eliminating uncompensated care, which we and others (cf. Einav and Finkelstein 2023) regard as unrealistic.

which we implement by endowing uninsurance with the certainty and non-pecuniary benefits of formal coverage, also raises welfare. However, the gain of $294 per consumer is fragile: it is undone by a marginal cost of public funds above 18%, or by a resource cost of providing the non-pecuniary benefit equal to 11% of its value to consumers.

References

Abaluck, Jason and Jonathan Gruber (2011) "Choice inconsistencies among the elderly: evidence from plan choice in the Medicare Part D program," *American Economic Review*, 101 (4), 1180–1210.

——— (2016) "Evolving choice inconsistencies in choice of prescription drug insurance," *American Economic Review*, 106 (8), 2145–2184.

Akerlof, George A (1978) "The market for "lemons": Quality uncertainty and the market mechanism," in *Uncertainty in economics*, 235–251: Elsevier.

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

Camilleri, Susan (2018) "The ACA Medicaid Expansion, Disproportionate Share Hospitals, and Uncompensated Care," *Health Services Research*, 53 (3), 1562–1580, https://doi.org/10.1111/1475-6773.12702.

Cardon, James H and Igal Hendel (2001) "Asymmetric information in health insurance: evidence from the National Medical Expenditure Survey," *RAND Journal of Economics*, 408–427.

Chiappori, Pierre-André and Bernard Salanié (2000) "Testing for asymmetric information in insurance markets," *Journal of political Economy*, 108 (1), 56–78.

Congressional Budget Office (2017) "The Budgetary Treatment of Cost-Sharing Reductions,"Technical Report 53961, Congressional Budget Office, `https://www.cbo.gov/publication/53961`, Accessed: 2025-05-09.

——— (2023) "Federal Subsidies for Health Insurance: 2023 to 2033,"Technical Report 59613, Congressional Budget Office, `https://www.cbo.gov/publication/59613`, Accessed: 2025-05-09.

Congressional Research Service (2018) "The Individual Mandate for Health Insurance Coverage: In Brief," CRS Report R44438, Congressional Research Service, `https://www.congress.gov/crs_external_products/R/PDF/R44438/R44438.12.pdf`, Updated March 19, 2018.

Coughlin, Teresa A, John Holahan, Kyle Caswell, and Megan McGrath (2014) "An estimated $84.9 billion in uncompensated care was provided in 2013; ACA payment cuts could challenge providers," *Health Affairs*, 33 (5), 807–814.

Domurat, Richard, Isaac Menashe, and Wesley Yin (2021) "The Role of Behavioral Frictions in Health Insurance Marketplace Enrollment and Risk: Evidence from a Field Experiment," *American Economic Review*, 111 (5), 1549–74, 10.1257/aer.20190823.

Drake, Coleman, Conor Ryan, and Bryan Dowd (2022) "Sources of inertia in the individual health insurance market," *Journal of Public Economics*, 208, 104622, https://doi.org/10.1016/j.jpubeco.2022.104622.

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

Einav, Liran and Amy Finkelstein (2023) *We've Got You Covered: Rebooting American Health Care*, New York: Portfolio/Penguin.

Einav, Liran, Amy Finkelstein, and Mark R. Cullen (2010) "Estimating Welfare in Insurance Markets Using Variation in Prices," *The Quarterly Journal of Economics*, 125 (3), 877–921, 10.1162/qjec.2010.125.3.877.

Fehr, Rachel, Daniel McDermott, and Cynthia Cox (2020) "Individual Insurance Market Performance in 2019," issue brief, Kaiser Family Foundation, https://www.kff.org/private-insurance/individual-insurance-market-performance-in-2019/.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams (2016) "Sources of geographic variation in health care: Evidence from patient migration," *The quarterly journal of economics*, 131 (4), 1681–1726.

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

Finkelstein, Amy, Neale Mahoney, and Matthew J. Notowidigdo (2018) "What Does (Formal) Health Insurance Do, and for Whom?" *Annual Review of Economics*, 10 (Volume 10, 2018), 261–286, https://doi.org/10.1146/annurev-economics-080217-053608.

Frean, Molly, Jonathan Gruber, and Benjamin D. Sommers (2017) "Premium subsidies, the mandate, and Medicaid expansion: Coverage effects of the Affordable Care Act," *Journal of Health Economics*, 53, 72–86, https://doi.org/10.1016/j.jhealeco.2017.02.004.

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

Ghili, Soheil, Ben Handel, Igal Hendel, and Michael D Whinston (2024) "Optimal Long-Term Health Insurance Contracts: Characterization, Computation, and Welfare Effects," *The Review of Economic Studies*, 91 (2), 1085–1121, 10.1093/restud/rdad054.

Goldin, Jacob, Ithai Z Lurie, and Janet McCubbin (2019) "Health Insurance and Mortality: Experimental Evidence from Taxpayer Outreach," Working Paper 26533, National Bureau of Economic Research, 10.3386/w26533.

Hadley, Jack, John Holahan, Teresa Coughlin, and Dawn Miller (2008) "Covering the uninsured in 2008: current costs, sources of payment, and incremental costs," *Health Affairs*, 27 (5), w399–w415.

Handel, Ben, Igal Hendel, and Michael D Whinston (2015) "Equilibria in health exchanges: Adverse selection versus reclassification risk," *Econometrica*, 83 (4), 1261–1313.

Handel, Benjamin R (2013) "Adverse selection and inertia in health insurance markets: When nudging hurts," *American Economic Review*, 103 (7), 2643–2682.

Kaiser Family Foundation (2014) "Uncompensated Care for the Uninsured in 2013: A Detailed Examination," report, Kaiser Family Foundation, https://www.kff.org/uninsured/report/uncompensated-care-for-the-uninsured-in-2013-a-detailed-examination/.

Koskinen, John A. (2015) "Letter to Members of Congress: Preliminary Results from the 2015 Filing Season Related to Affordable Care Act Provisions," July, https://www.irs.gov/pub/newsroom/CommissionerLetterlwithcharts.pdf.

Mahoney, Neale (2015) "Bankruptcy as Implicit Health Insurance," *American Economic Review*, 105 (2), 710–46, 10.1257/aer.20131408.

McIntyre, Adrianna, Mark Shepard, and Myles Wagner (2021) "Can Automatic Retention Improve

Health Insurance Market Outcomes?" *AEA Papers and Proceedings*, 111, 560–66, 10.1257/pandp.20211083.

Myerson, Rebecca, Nicholas Tilipman, Andrew Feher, Honglin Li, Wesley Yin, and Isaac Menashe (2022) "Personalized Telephone Outreach Increased Health Insurance Take-Up For Hard-To-Reach Populations, But Challenges Remain," *Health Affairs*, 41 (1), 129–137.

Rae, Matthew, Cynthia Cox, Gary Claxton, Daniel McDermott, and Anthony Damico (2021) "How the American Rescue Plan Act Affects Subsidies for Marketplace Shoppers and People Who Are Uninsured,"Technical report, Kaiser Family Foundation, https://www.kff.org/affordable-care-act/issue-brief/how-the-american-rescue-plan-act-affects-subsidies-for-marketplace-shoppers-and-people-who-are-ur

Saltzman, Evan (2019) "Demand for health insurance: Evidence from the California and Washington ACA exchanges," *Journal of Health Economics*, 63, 197–222, https://doi.org/10.1016/j.jhealeco.2018.11.004.

——— (2021) "Managing adverse selection: underinsurance versus underenrollment," *The RAND Journal of Economics*, 52 (2), 359–381.

Shepard, Mark and Myles Wagner (2022) "Do Ordeals Work for Selection Markets? Evidence from Health Insurance Auto-Enrollment," Working Paper 30781, National Bureau of Economic Research, 10.3386/w30781.

Tebaldi, Pietro (2025) "Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca," *Review of Economic Studies*, 92 (1), 586–620.

Tebaldi, Pietro, Alexander Torgovitsky, and Hanbin Yang (2023) "Nonparametric estimates of demand in the california health insurance exchange," *Econometrica*, 91 (1), 107–146.

Treasury Inspector General for Tax Administration (2017) "Affordable Care Act: Implementation of the Notification Requirement for Individual Filers Not Enrolled in Health Insurance," https://www.treasury.gov/tigta/press/press_tigta-2017-19.htm, Report TIGTA-2017-19, August 3, 2017.

A Federal Poverty Line

Table 10 describes the parameters of the Federal Poverty Line (FPL) in each year from 2012 to 2020. Recall that, for a household with n members, the FPL in year t is

$$FPL_t(n) = \alpha_t + \beta_t \times (n-1).$$

B Computation of exchange equilibria

This appendix describes how we solve condition (5) for the equilibrium premium p^*_{25}. One feature of the model that simplifies calculation of equilibria is that a consumer's willingness to pay for insurance has a closed form:

$$WTP_{jt} = \phi \ell_{jt} \; + \; z_{jt} \; + \; \frac{1}{\gamma_j} \log \frac{\mathbb{E}\big[e^{\gamma_j OOP^c_{jt}(0)} \mid \lambda_{jt}, \varepsilon^0_{jt}\big]}{\mathbb{E}\big[e^{\gamma_j OOP^c_{jt}(1)} \mid \lambda_{jt}, \varepsilon^0_{jt}\big]}. \tag{17}$$

Table 10: Federal Poverty Line, 2012–2020 ($)

Year	First person (α_t)	Each additional person (β_t)
2012	11170	3960
2013	11490	4020
2014	11670	4060
2015	11770	4160
2016	11880	4140
2017	12060	4180
2018	12140	4320
2019	12490	4420
2020	12760	4480

Because the regulated schedule is proportional to the age-25 premium, $\bar{p}(p_{25}, a_{jt}) = \bar{p}(1, a_{jt})\, p_{25}$, we refer to $\bar{p}(1, a_{jt})$ as consumer j's age factor. Dividing WTP_{jt} by it expresses the consumer's reservation premium in age-25 units, which we denote p^r_{jt}: consumer j insures if and only if $p_{25} \leq p^r_{jt}$.

The first step in our algorithm involves sorting simulated consumers j in descending order of p^r_j. Here, we use the term "simulated consumer" to reflect that these consumers are simulated clones of MEPS panelists, each with a distinct value of the unobserved heterogeneity $(\gamma_j, \varepsilon^0_j)$. Note that we suppress the time subscript t. Also, without loss of generality, let j indicate the consumer's ordering in the ranking of p^r_j so that $p^r_1 \geq p^r_2 \geq \cdots \geq p^r_J$, where J is the number of simulated consumers.

It is useful for characterizing the algorithm to introduce two functions, average premium revenue per enrollee

$$\text{Revenue}(p) = \mathbb{E}\big[\, \bar{p}(p, a_{j'}) \mid p^r_{j'} \geq p \,\big]$$

and average cost per enrollee, inclusive of the insurer markups that τ captures,

$$\text{Cost}(p) = \tau\, \mathbb{E}\big[\, m_{j'} - OOP_{j'} \mid p^r_{j'} \geq p \,\big]\,.$$

For each p^r_j we compute the break-even function

$$\text{BE}(p^r_j) = \text{Revenue}(p^r_j) - \text{Cost}(p^r_j).$$

The condition $\text{BE}(p) = 0$ is equivalent, whenever participation is positive, to the equilibrium condition (5) of the model section. The pool of enrollees changes discretely as p crosses each p^r_j, however, so BE jumps at every reservation premium and generically has no root among the p^r_j that we evaluate. We therefore identify the equilibrium by the largest pool of consumers that at least breaks even,

$$j^* = \max\big\{\, j \;:\; \text{BE}(p^r_j) \geq 0 \,\big\}\,,$$

and charge that pool the premium at which it exactly breaks even. Because $\bar{p}(\cdot, a)$ is proportional to the age-25 premium, Revenue is linear in p while the pool is held fixed, so that premium is

$$p_{25}^{*} = \frac{\text{Cost}(p_{j^*}^{r})}{\mathbb{E}\Big[\,\bar{p}(1, a_{j'}) \mid p_{j'}^{r} \geq p_{j^*}^{r}\,\Big]},$$

the denominator being the average age factor among those enrollees. The set defining j^* may consist of several disjoint blocks, corresponding to multiple equilibria; taking the largest j selects the one with the highest participation. If $\text{BE}(p_j^r) < 0$ for every j, no premium supports positive participation and the market fully unravels, as in Figure 18.

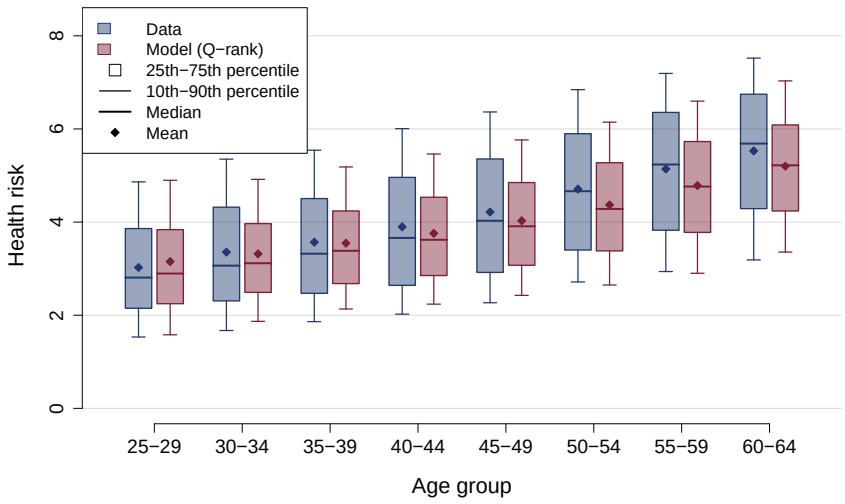
A consumer eligible for premium subsidies pays the smaller of $\bar{p}(p_{25}, a_j)$ and a cap set by the applicable-percentage schedule of Section 4.3, which depends on income. Any such consumer whose willingness to pay exceeds that cap therefore enrolls at every market premium, so we set $p_j^r = \infty$ and place them at the top of the ranking, where $\text{BE}(p_j^r) \geq 0$ holds trivially.

C Health-risk and income transitions

The lifetime welfare calculations of Section 8 require simulating each consumer's health risk and income from age 25 to 64. We model the evolution of health risk as a first-order autoregressive process in quantile ranks. Using the two-year MEPS panel, we convert each panelist's risk score in each year into a quantile rank within that calendar year's cross-sectional distribution, and estimate, separately for each five-year age group, a linear transition of the current rank on the prior-year rank with normally distributed innovations whose standard deviation varies linearly with the prior rank, by maximum likelihood. To simulate a lifetime trajectory, we initialize each simulated life cycle at the observed rank of a consumer aged 25–29, iterate the estimated rank transition forward one age at a time, and convert the simulated rank at each age into a risk-score level by inverting the empirical distribution of the risk score at that age. Figure 20 compares the distribution of simulated risk scores with the observed cross-sectional distribution within each age group. Because the simulation is initialized only from the youngest group, the agreement at older ages is a genuine test of the estimated transition process; the model tracks the rise in both the level and the dispersion of health risk over the life cycle, with a modest understatement of dispersion at the oldest ages.

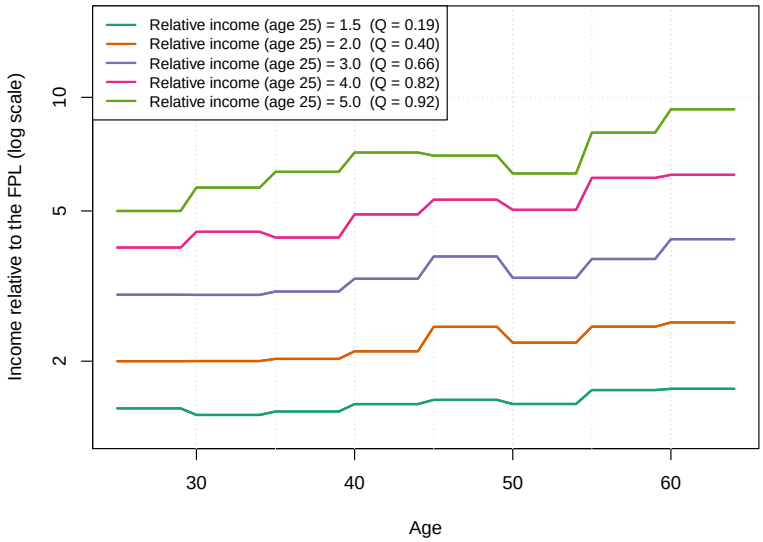
Because our focus is insurance of health risks rather than income risks, we keep the income process deliberately simple: each consumer retains for life their age-25 rank in the income distribution. Upon entering each five-year age group, the consumer's income relative to the FPL is set to the value at that preserved rank within the age group's cross-sectional income distribution in our sample. Income paths therefore exhibit life-cycle progression—the cross-sectional income profile changes with age—but no idiosyncratic income risk. Figure 21 displays the resulting trajectories for consumers whose age-25 incomes range from 150% to 500% of the FPL.

Figure 20: Fit of the health-risk transition process



Notes: This figure compares, within five-year age groups, the survey-weighted distribution of the risk score observed in the MEPS sample with the distribution obtained by simulating the estimated rank-transition process forward from the observed distribution among consumers aged 25–29. Boxes span the interquartile range, whiskers the 10th to 90th percentiles, horizontal lines mark medians, and diamonds mark means.

Figure 21: Rank-preserving income trajectories



Notes: This figure displays the income trajectories implied by the rank-preserving income process for five consumers whose age-25 incomes equal 150%, 200%, 300%, 400%, and 500% of the federal poverty line. Each consumer's rank in the survey-weighted income distribution is fixed at its age-25 value; within each subsequent five-year age group, income relative to the poverty line is the value at that rank in the age group's cross-sectional distribution. The legend reports each trajectory's preserved rank.

D Moment values

Table 11: Choice model moments

Moment	Model	Data	
Chiappori–Salanié coefficient: male; age <45	0.253	0.267	(0.129)
Chiappori–Salanié coefficient: male; age ≥45	0.184	0.154	(0.086)
Chiappori–Salanié coefficient: female; age <45	0.348	0.348	(0.095)
Chiappori–Salanié coefficient: female; age ≥45	0.184	0.173	(0.059)
Participation: income 100–150% FPL; age <45	0.212	0.217	(0.040)
Participation: income 100–150% FPL; age ≥45	0.394	0.356	(0.053)
Participation: income 150–200% FPL; age <45	0.181	0.184	(0.019)
Participation: income 150–200% FPL; age ≥45	0.305	0.304	(0.027)
Participation: income 200–250% FPL; age <45	0.196	0.207	(0.022)
Participation: income 200–250% FPL; age ≥45	0.365	0.374	(0.030)
Participation: income 250–300% FPL; age <45	0.202	0.154	(0.027)
Participation: income 250–300% FPL; age ≥45	0.402	0.469	(0.036)
Participation: income 300–350% FPL; age <45	0.195	0.222	(0.038)
Participation: income 300–350% FPL; age ≥45	0.543	0.412	(0.044)
Participation: income 350–400% FPL; age <45	0.232	0.241	(0.042)
Participation: income 350–400% FPL; age ≥45	0.449	0.367	(0.049)
Participation: income 400–500% FPL; age <45	0.274	0.267	(0.041)
Participation: income 400–500% FPL; age ≥45	0.304	0.394	(0.041)
Participation: income 500%+ FPL; age <45	0.353	0.328	(0.035)
Participation: income 500%+ FPL; age ≥45	0.313	0.448	(0.031)
Participation: year 2016	0.361	0.335	(0.018)
Participation: year 2017	0.301	0.291	(0.019)
Participation: year 2018	0.320	0.313	(0.020)
Participation: year 2019	0.350	0.326	(0.020)
Participation: region Northeast	0.296	0.313	(0.026)
Participation: region Midwest	0.330	0.333	(0.021)
Participation: region South	0.261	0.275	(0.012)
Participation: region West	0.320	0.337	(0.018)
Participation: low income; age <45; λ quartile 1	0.155	0.118	(0.015)
Participation: low income; age <45; λ quartile 2	0.183	0.166	(0.023)
Participation: low income; age <45; λ quartile 3	0.190	0.252	(0.029)
Participation: low income; age <45; λ quartile 4	0.353	0.363	(0.044)
Participation: low income; age ≥45; λ quartile 1	0.199	0.190	(0.044)
Participation: low income; age ≥45; λ quartile 2	0.194	0.224	(0.032)
Participation: low income; age ≥45; λ quartile 3	0.269	0.319	(0.030)
Participation: low income; age ≥45; λ quartile 4	0.525	0.516	(0.027)
Participation: high income; age <45; λ quartile 1	0.205	0.155	(0.028)
Participation: high income; age <45; λ quartile 2	0.239	0.299	(0.035)
Participation: high income; age <45; λ quartile 3	0.282	0.280	(0.042)
Participation: high income; age <45; λ quartile 4	0.491	0.476	(0.054)
Participation: high income; age ≥45; λ quartile 1	0.009	0.244	(0.065)
Participation: high income; age ≥45; λ quartile 2	0.076	0.283	(0.042)
Participation: high income; age ≥45; λ quartile 3	0.169	0.388	(0.036)
Participation: high income; age ≥45; λ quartile 4	0.666	0.519	(0.031)
Mandate penalty coefficient	0.085	0.095	(0.010)

Notes: The table reports the 45 moments used to estimate the choice model. "Model" is the simulated moment at the estimated parameters; "Data" is the sample moment, with its standard error in parentheses. Participation moments are take-up rates within the indicated cell. Coverage coefficients are the take-up slopes from the Poisson regressions described in the main text, estimated separately by gender and age group, with classical standard errors that allow for overdispersion.