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

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

October 31, 2025

Preliminary and incomplete — do not cite or distribute

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 various factors. We find that uncompensated care is the main friction impeding take-up. Uncompensated care depresses consumers' willingness to pay for insurance to the extent that the exchange market fully unravels without premium subsidies. Subsidies are necessary because of uncompensated care: eliminating both subsidies and uncompensated care raises participation from 39% to 84%.

*We thank Amy Finkelstein and seminar participants at Princeton, Emory, and Stanford 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 increase participation. 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 incomes too high to qualify for Medicaid. The exchanges are coordinated markets on which private insurers list plans that adhere to regulated cost-sharing features. The ACA offers generous premium subsidies, especially for consumers with lower incomes, and initially featured a penalty for failure to purchase insurance. This fine, commonly called the individual mandate, was as high as 2.5% of yearly 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 in the exchanges 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 insurance on the ACA exchanges suggests the presence of major impediments to individuals' 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 likely future medical expenses fails to rationalize such low participation.

The health insurance literature has documented consumers' low willingness to pay for medical 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. Uninsured consumers' bills often go unpaid, and consumers even 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, but could instead (or in addition) reflect financial constraints that make consumers more concerned about immediate payments than about future ones. Other possible frictions that may reduce uptake of insurance include 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.

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

To explain observed participation patterns, we incorporate several frictions—including uncompensated care and myopia—into the standard model of insurance choice. Combining the augmented model of consumer choice with a model of supply side conduct on the exchanges, we evaluate the impact of the different frictions on exchange participation, welfare, and the effectiveness of alternative government policies. By estimating a model of medical expenses and participation in the exchanges that incorporates these frictions we essentially run a horserace among them to see assess which are the most important.

Our primary data source is the Medical Expenditure Panel Survey (MEPS), which provides information on panelists' source of insurance (if any), medical conditions, medical events (e.g., emergency room visits or prescription drug purchases), and medical expenses by payer (e.g., insurance, out-of-pocket, or state/local programs). 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.

The main empirical innovation of our paper is using direct measures of uncompensated care, to assess its impact on willingness to pay for exchange insurance. The key data in MEPS, information on payments for medical services by *type of payer* (e.g., out-of-pocket, charity or state/local programs). Enabling us directly measure OOPs the uninsured are responsible for, separately from the rest of the expenses. These OOP distributions represent the actual risk faced by uninsured individuals given their health risk, income and demographics. Incorporating the implicit coverage enjoyed by the uninsured enables us to assess the extent to which uncompensated care reduces the willingness to pay for coverage.

We find that while insured individuals pay about 20% of their medical expenses (80% is covered by insurance), uninsured individuals pay about 50%. They pay more than insured individuals, but far less than footing the entire bill. In addition, we find that the share that goes uncompensated increases in total expenses. This suggests that uncompensated care as reported in MEPS works like a de facto catastrophic coverage or high deductible insurance plan, offered at zero premium to the uninsured.

Uncompensated care takes two forms: first, payments for medical services made by entities other than the patient or an insurer (e.g., charities or government programs) and, second, non-recouped expenses that are implicitly paid by providers. Although the former sort of uncompensated care is recorded in our data, the latter is not. This is because uninsured consumers may decline to pay their bills, providers' collections operations may fail, and consumers may declare bankruptcy to avoid legal obligations to pay for treatment received (Mahoney (2015)). These non-recouped expenses can be estimated as the extent to which medical providers receive greater payments from insured consumers than from uninsured ones. The gap indicate implicit subsidization of uninsured consumers by medical providers. We call this implicit subsidization *unobserved uncompensated care*. We find that the mean amount of unobserved uncompensated care in the entire uninsured population is $1377, almost as high as the mean observed amount of $1520. We also find that the unobserved components of uncompensated care is higher for

higher deciles of the total expense distribution. As such, uncompensated care bears resemblance to catastrophic coverage or high deductible insurance plans when viewed as a form of implicit insurance. Once unobserved uncompensated care is accounted for, OOPs paid by the uninsured represent about 26% of total expenses.

The ingredients of our model are the estimated frictions, preferences toward risk, transitions across health types (costs) over time as well as the distribution of health costs condition on health type. Consumers with constant absolute risk aversion (CARA) utility make a binary choice of coverage to maximize their expected utility taking expectations over the estimated OOP distributions. In addition to the uncompensated care embedded in the estimation (through its effect on the observed OOP distributions), we allow for myopia à la Abaluck and Gruber (2011) that varies across income groups, as well as heterogeneity in tastes for insurance unrelated 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.

We combine this model of insurance demand incorporating frictions with a competitive supply side as in Handel et al. (2015). The estimated model permits simulation of equilibrium participation and welfare under counterfactual regimes featuring varying frictions and regulations.

Using this estimated model we find that uncompensated care is the main friction limiting insurance uptake. Under the baseline regime with status quo regulations (premium subsidies but no individual mandate), the model predicts a 38.4% participation rate among the exchange-eligible population. We estimate that participation would increase to 94.3% upon the elimination of uncompensated care. Based on our estimates, myopia plays a relatively modest role: if consumers were to weight OOPs equally to premiums when they buy insurance, participation would be 10 percentage points higher than under the baseline regime but still under 50%.

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 raises the participation rate to 84%. Thus, heavy subsidization of premiums is necessary to prevent market unravelling due to the presence of uncompensated care. We estimate that the mandate penalty had a smaller 12 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 enact-

ment 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 consumers' pre-existing medical conditions. Healthier individuals pay lower premiums under risk-based pricing, which consequently raises participation to 52.5% from 38.4% under the baseline community rating regime. Yet, while eliminating adverse selection by risk-based pricing boosts participation by more than eliminating myopia, it does not nearly achieve the participation increase of eliminating uncompensated care.

Having characterized the participation effects of frictions and inducement policies, we conduct welfare analysis to understand their normative impact. First, the welfare analysis highlights the importance of capturing behavioral frictions. While an ACA-type exchange that forces full participation is the preferred outcome of fully rational consumers, full participation delivers the lowest welfare in the eyes of myopic customers who discount OOPs relative to premiums.

Our second welfare finding is that the performance of different regimes in terms of social efficiency is tightly connected to their participation impacts. Participation increases, however achieved, limit the extent that myopic consumers make costly mistakes. We additionally find that the reclassification risk induced by risk pricing is not as harmful to consumer welfare as previously found using the fully rational model Handel et al. (2015). Both subsidies and uncompensated care attenuate the welfare losses associated with risk pricing. Subsidies absorb premium hikes, while uncompensated care sharply limits the harms from going uninsured when premiums are the highest under risk pricing. Absent uncompensated care, participation is high across the various regimes, delivering similar aggregate welfare levels. One exception is risk pricing: absent uncompensated care, reclassification risk is especially painful as individual can no longer escape high premiums by relying on the implicit, free insurance of uncompensated care. For that reason aggregate welfare is low under risk pricing when uncompensated care is eliminated.

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 takeup. 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 studies the effects of a randomized governmental outreach to indi-

viduals paying the individual mandate penalty on insurance uptake participation 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, suggesting the relevance of informational frictions. 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. Our study joins Abaluck and Gruber (2011) and Abaluck and Gruber (2016) in using a revealed preference approach exploiting data describing 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.

There is additionally a body of literature studying hassle costs impeding insurance adoption. Articles in this literature include Shepard and Wagner (2022) and McIntyre et al. (2021), who study the impact of auto-enrolment, finding 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 provides an explanation for consumers' 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 them insurance, 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). A Kaiser Family Foundation report from 2013 (link) found that 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 is present: Camilleri (2018) finds that uncompensated care fell at hospitals that treat a

disproportionately high number of low-income patients.

Our article also adds to a literature studying policies that encourage participation in health insurance exchanges. 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. Saltzman (2019) estimates demand in the California and Washington exchanges, reporting a modest response of enrollment to penalties. Frean et al. (2017) using the American Community Survey (ACS) finds 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

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 noninstitutionalized 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.

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 state of residence; insurance status; medical conditions applying to each panelist in the form of three-digit International Classification of Diseases (ICD) codes; and medical events. Medical events include outpatient, inpatient, emergency room, office-based, and home health visits in addition to drug prescriptions and various miscellaneous expenses.

The key information to quantify uncompensated care is that the medical expenses are reported by payer (insured, insurance company, charities, etc.), allowing us to observe the share of total expenses paid by the consumer. Throughout the article, we use the term expenses to refer to payments received by providers for administering medical services. In no part of the analysis do we use the "charges" variables provided by MEPS, which may refer to list prices or billed amounts rather than actually exchanged dollars. The distribution of OOP expenses (net of expenses paid by other sources) that we observe conditional on health type and insurance status represent the risk faced by the individual under each insurance status given her health type. We use the distribution of OOP expenses of insured and insured individuals to compute expected utility.

The ACA exchanges cater to individuals who are ineligible for Medicaid, a public program that provides health insurance for individuals with low incomes. In particular, Medicaid covers those with household incomes below 138% of the Federal Poverty Line (FPL) in states that expanded

[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.

the program under the ACA ("expansion states") and below 100% of the FPL in non-expansion states. The FPL depends on household size; in each year t, it takes the form

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

where n is the household size and α_t and β_t are parameters. The values of α_t and β_t are reported in Appendix A. In 2019, the most recent year that we study, the FPL for a single-member household was $12,490 and increased $4420 for each additional household member.

We construct an individualized household income measure that can be compared to $FPL_t(1)$, the single-member household FPL, to determine eligibility for Medicaid. When Y_{jt} is panelist j's household income in year t, their individualized income in that year is

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

When y_{jt} is greater than 100% (138%) of the single-member-household FPL, then panelist j is ineligible for Medicaid in non-expansion (expansion) states.

As we will describe in Section 4.2, eligibility for premium subsidies on the ACA exchanges also depends on income relative to the FPL. Thus, we can also compare y_{jt} to the single-member-household FPL to determine the extent of premium subsidies available to the panelist.

More generally, we use the individualized household income y_{jt} as our measure of income throughout the article. This measure reflects the resources available to panelist j in a manner that reflects cohabitation economies and is consistent with the US government's determination of how living costs scale with household size.

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 Federal Poverty Level (FPL, which varies by household size). 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 determining the sample of exchange-eligible consumers, we eliminate individuals who do not meet age- and income-related criteria: we drop panelists who are younger than 25 years or older than 64 years, and who have household incomes below 138% of the Federal Poverty Level (FPL, which varies by household size). This is because individuals under 25 years of age are typically covered by their parents' plans and because those 65 years and older are eligible for Medicare. Also, individuals under 138% of the FPL are eligible for Medicaid in states that expanded this program, which covers low-income adults. Although some states failed to expand Medicaid, meaning that only those under 100% of the FPL were eligible for the program in these states,

we use the 138% threshold to ensure that all Medicaid-eligible consumers are excluded from the sample.

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. MEPS includes panelist-level weights that are recommended for use in cross-panelist analyses. We use these weights in estimating distributions of OOP expenses and in estimating the insurance choice model.

The publicly accessible data do not contain state of residence of the panelist, for confidentiality reasons. Because health exchange premiums vary by state, we gained access to the MEPS restricted data at the AHRQ's data center in Rockville, Maryland to be able to use this premium variation in estimating our insurance choice model.

2.2 Medical utilization and expenditures data

We use medical expenditure, medical utilization, and demographic information data in MEPS to build an ACG-type risk prediction. The measure predicts future medical utilization for each individual in the data. We use h_j to denote an individual j's expected log medical expenses, and often refer to h_j as the consumer's *health status*. After characterizing the expected medical expenses of each individual, we estimate the distribution of OOP medical expenses faced by each consumer conditional on insurance status, income and λ.

2.3 Premiums data

We use data from the Kaiser Family Foundation on premium data by state for insurance plans purchased on ACA exchanges. These data describe premiums for Silver plans, a tier of insurance coverage in the exchanges that is especially attractive to consumers as the tier to which income-based cost-sharing restrictions apply (see Section 4.3). In particular, Kaiser provides premiums for the Silver plan with the second-lowest premium in one major city within each state for a 40-year-old non-smoker on an annual basis. We use these premiums as our measures of state/year-level premiums for 40-year-olds. To obtain premiums for other ages, we apply the ACA age schedule that we will later describe later in Section 4.1.

Since premium subsidies depend on income, information on premium and income are used to compute the effective, post-subsidy premium paid by each individual in the sample. Whereas some states operate their own insurance exchanges following the enactment of the ACA, others use the federal exchange Healthcare.gov. The plans, however, listed on Healthcare.gov are state-specific, and the KFF premium data describe state-level premiums for both state-based and federally-operated exchanges.

2.4 Risk aversion

Handel et al. (2015) use data on medical claims and insurance plan choice from a large employer

to estimate a model of risk preferences. Consumers in the model have CARA preferences and heterogeneous risk aversion coefficients. Handel et al. (2015) estimate a distribution of consumers' CARA risk aversion coefficients conditional on age. We use these estimates of risk aversion, which appear in the fourth column ("Robustness 3") of Table 2 of Handel et al. (2015), in our study.[3] Under these estimates, the intercept of the risk aversion coefficient's conditional mean is 2.54×10^{-4}, the slope of the conditional mean in age is 3.99×10^{-6}, and the standard deviation of the risk aversion coefficient is 1.20×10^{-4}. These estimates imply that a consumer with the average risk aversion coefficient among 40 year olds would be indifferent about accepting a 50-50 gamble with a winning pay-off of $1,000 and a losing penalty of $704.

3 Uncompensated care

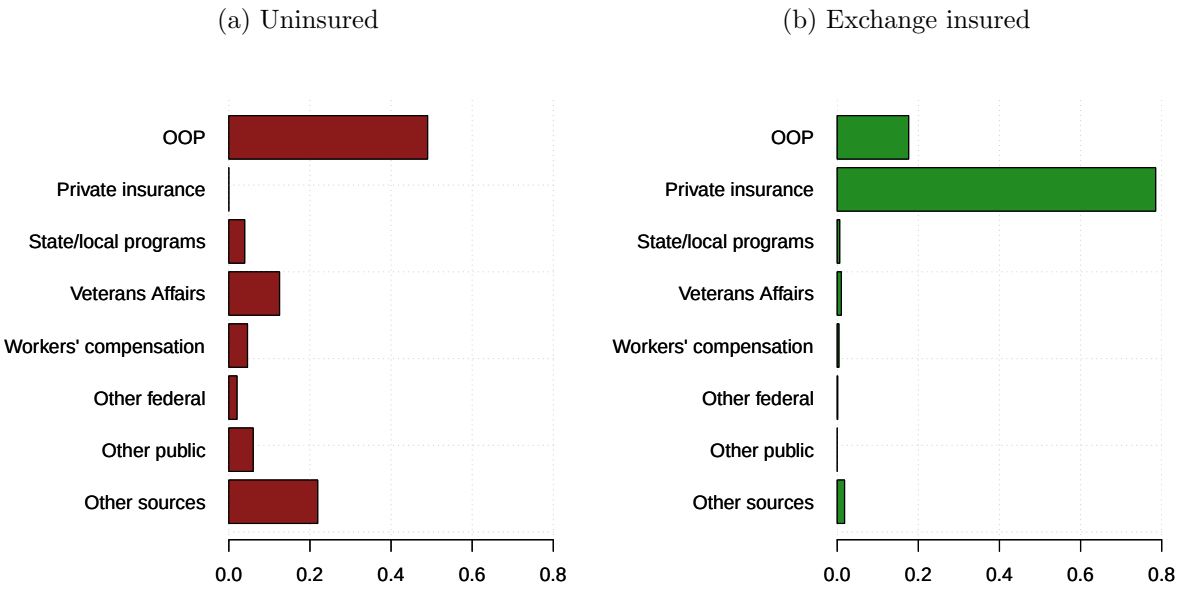
Uncompensated care provides one explanation for the low willingness of consumers to pay for exchange insurance. One of our paper's principal contributions is to directly measure uncompensated care and OOPs net of uncompensated care using MEPS. MEPS allows us to quantify two components of uncompensated care. First, MEPS includes data on sources of payment for medical services; these data allow us to measure the share of consumers' total expenses that are covered by insurance, OOPs, and other payers. We call payments made by these other payers *observed uncompensated 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 *unobserved uncompensated care*.

First, we assess the extent of observed uncompensated care by decomposing medical payments by payer. MEPS provides payments made by the following sorts of payers: out-of-pocket, private 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 medical providers across various payment sources, separately for the uninsured and the privately insured. The figure shows that about half of uninsured consumers' payments are covered by OOPs with uncompensated 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.

Next, we assess the relationship between observed uncompensated care and total medical expenses. This relationship is critical for the insurance value of uncompensated care to the unin-

[3]Handel et al. (2015) allow risk preferences to depend on the individual's job position at the employer—specifically, whether or not the individual was a manager at the organization—in addition to age. We use the estimated distribution of risk preferences for non-managers, who comprise 77% of employees at the large employer studied by of Handel et al. (2015)

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



sured and thus consumers' willingness to pay for formal insurance. Uncompensated care, for example, would provide greater insurance value and significantly lower willingness to pay for formal insurance if it disproportionately covered especially painful larger expenses.

To measure the relationship between uncompensated care and total expenses, we compute mean observed uncompensated care for panelist/year observations in MEPS within each decile of total expenses in this population. Throughout this analysis, we consider only observations with positive total expenses and we limit the sample to MEPS panelists who were uninsured for an entire calendar year. Figure 2 shows the results. Figure 2a plots the median and interquartile range of observed uncompensated care share within each total expenses decile. The figure shows that observed uncompensated 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 heterogeneity in use of uncompensated care.

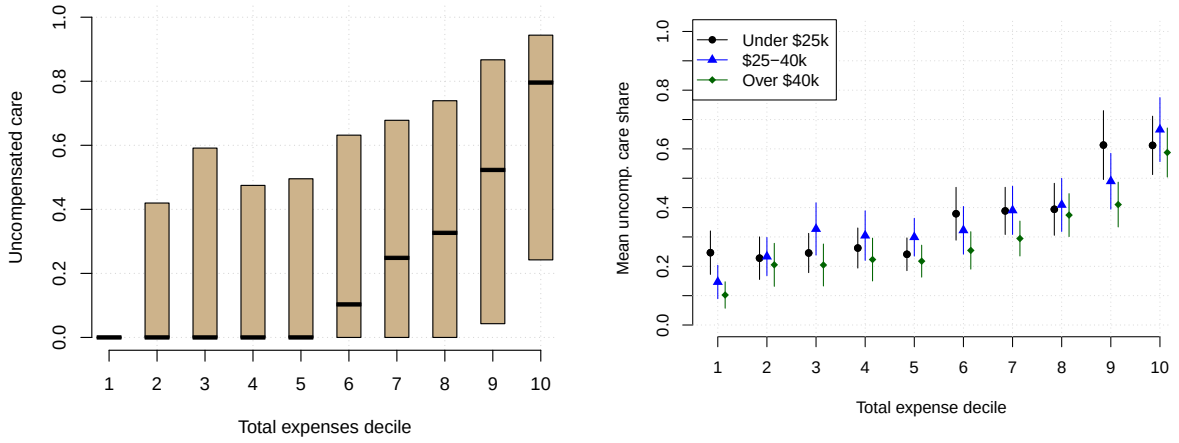
The second panel, Figure 2b, plots the mean observed uncompensated care share separately for three income groups. Within each group the mean share is increasing in total expenses, but income does little to explain utilization of uncompensated care.

As we noted above, *observed* payments by non-insurance sources for treatment of the uninsured is not the only source of uncompensated care. In addition, service providers may be able to recoup less from all sources for treating an uninsured patient relative to an insured patient. This is because uninsured consumers may decline to pay their bills, providers' collections operations may fail, and consumers may declare bankruptcy to avoid legal obligations to pay for treatment received. A desirable feature of MEPS is that it allows us to determine the extent to which medical service providers recover different payment amounts from insured and uninsured individuals

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

(a) Interquartile ranges by total expense decile

(b) Mean uncompensated care share by income



Notes: Figure 2a plots the interquartile range and median of uninsured panelists' shares of total expenses (payments received by providers) constituted by payments other than OOPs (i.e., observed uncompensated care). We conduct the analysis at the panel/year level and include only panelists who are uninsured for the entire year. We also include only panelists with positive total expenses. Figure 2b plots the mean share of total expenses constituted by observed uncompensated care by income group. The whiskers running through each point indicate 95% confidence intervals for the estimated means. Here, income refers to the individualized household income measure.

for similar medical services. To do so, we regress the log total payments providers received for a medical event on an insured indicator and a rich set of controls for the characteristics of the medical event:

$$\log(\text{TOT}_j) = \beta a_j + x_j'\gamma + \varepsilon_j.$$

Here, j indicates a medical event, TOT_j is the sum of payments received by the provider for the event, a_j is an insured indicator (1 if insured, zero otherwise), and the controls x_j include an intercept term as well as indicators for the type of procedures performed (e.g., vaccinations, diagnostic tests, x-rays, or surgeries) and the ICD diagnostic codes applying to consumer j.[4] Note that the variable a_j measures whether the panelist was covered by insurance plan, not whether insurance contributed to payment for medical event j (a divergence may arise if, e.g., the patient has not yet paid their insurance deductible). In this linear model, the coefficient β provides the approximate percentage effect of the consumer being insured on the sum of payments received for the consumer's treatment holding fixed the type of consumer

[4]MEPS provides three-digit ICD codes applying to panelists in the sample. The ICD is an international system for classifying diseases. The US officially implemented the 10th version of the ICD, ICD-10, in 2015, replacing the ICD-9 system. In 2016, MEPS switched from reporting ICD-9 codes to ICD-10 codes. We include indicators for both sets of codes to enable pooling together events from before and after the switch to ICD-10. Examples of common ICD-10 codes are: I10 — Essential (primary) hypertension; E78 — Disorders of lipoprotein metabolism and other lipidaemias (this includes conditions related to cholesterol levels); F32 — Depressive episode; and J45 — Asthma. There were 432 unique ICD-9 codes and 474 unique ICD-10 codes recorded for MEPS panelists in the years 2012–2019.

attended to and the characteristics of the consumer's visit. In Table 1 we show separate ordinary least squares estimates for an event-level data set for four types of medical event— inpatient, outpatient, emergency room, and office-based visits.

Table 1 shows that, across these four event categories, there exist significant gaps between the payments that providers receive from the insured and the uninsured. For inpatient visits, which are especially costly, hospitals recover 1.34 log points (282%) more in payments from treated insured patients than from treating uninsured patients. Outpatient and emergency room (ER) visits also feature large gaps in payments received for treating insured and uninsured panelists. The gap is smaller at about 4% for office-based visits, which are often less costly, less acute, and more elective in nature than hospital visits. The table provides a sense of the extent and type of expenses that providers do not get paid for. In our welfare analysis, we account for the gap in payments received by medical service providers from treating insured and uninsured patients to fully capture the welfare implications of insurance take-up.

Table 1: Difference in log payments received between insured and uninsured

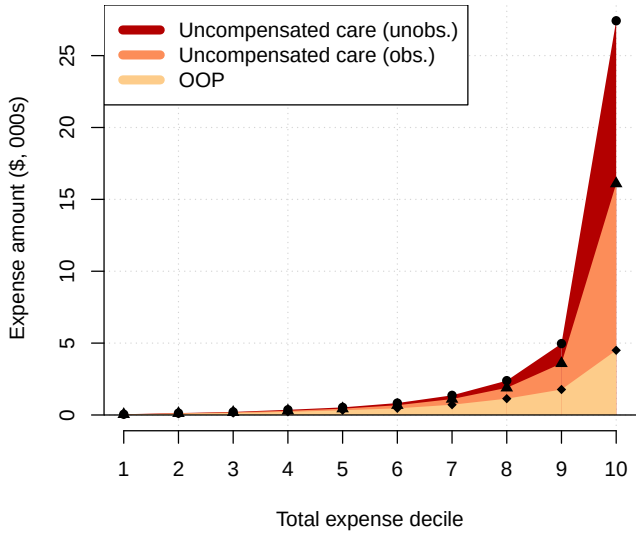
Event type	Estimate	SE
Inpatient	1.34	(0.303)
Outpatient	0.38	(0.092)
ER	0.45	(0.100)
Office-based	0.04	(0.012)
All	0.42	(0.015)

The fact that medical providers receive lower payments from insured panelists motivates our definition of *unobserved uncompensated care* as the additional amount that medical providers would have received from an uninsured patient if that patient was instead insured. Unobserved uncompensated care includes unpaid amounts of uninsured consumers' medical bills—perhaps owing to declarations of bankruptcy—and uninsured consumers' use of publicly funded health clinics that extend care to the uninsured at a lower cost than conventional providers. Our method for computing unobserved uncompensated care is similar to that used in the regression analysis underlying Table 1: we regress log payments received by providers on visit characteristics and an indicator for the patient being insured. The coefficient on this indicator measures the extent to which providers receive more payments from insured patients. In constructing the unobserved uncompensated care measure, however, we account for sorts of expenses not accounted for by the Table 1 regressions, including those for home health visits and prescription drugs. Additionally, we account for variation in unobserved uncompensated care at a finer level than "Event type" (inpatient, outpatient, etc.) by interacting the insurance indicator with many medical event characteristics. Appendix B details the procedure.

Figure 3 displays the results of the exercise. Note that the deciles are based on total expenses before unobserved uncompensated care and that only uninsured consumers with positive total expenses are included in the analysis. We draw two conclusions from the figure. First, uncompensated care primarily protects consumers who experience large expenses: both the observed

and unobserved components of uncompensated care are largest relative to OOPs for consumers in the higher deciles of the total expense distribution. In this respect, uncompensated care bears resemblance to catastrophic coverage or high deductible insurance plans when viewed as a form of implicit insurance. Second, the unobserved and observed components of uncompensated care are similar in magnitude. In the entire uninsured population, the mean amount of unobserved uncompensated care is $1377 whereas the mean amount of observed uncompensated care is $1520.

Figure 3: Uncompensated care and OOPs among the uninsured



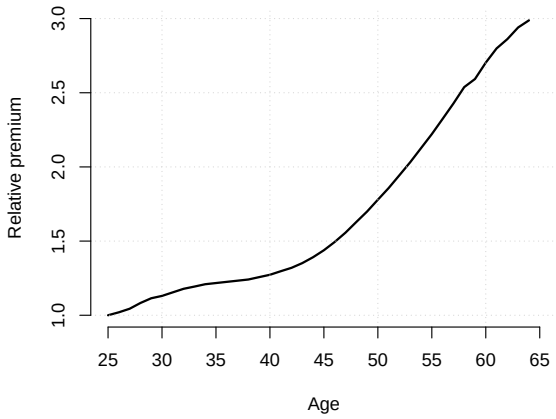
4 Features of the ACA exchanges

This section describes institutional features of the ACA exchanges with the goals of clarifying the environment in which we estimate the model and highlighting sources of variation in our data.

4.1 Age-adjusted community rating

As we noted above, insurers are not permitted to price discriminate on the basis of consumers' pre-existing conditions, which we collectively refer to as consumer health risk. The practice of insurers charging all consumers in a market the same premium is known as community rating. The ACA enacted a modified version of community rating with two distinctions from a pure community rating regime. First, an insurer may charge tobacco smokers premiums up to 50% more than the premiums charged to non-smokers. Second, premiums for an exchange insurance plan vary by consumer age according to a regulated schedule. Figure 4 displays the ratio of the premiums charged to consumers of various ages relative to those charged to a 25-year-old

Figure 4: Schedule of relative premiums by age



consumer. This schedule, which applies across all states, is mandated by the ACA. Consumers who are 64 years old pay premiums about three times those paid by a 25-year-old consumer. We often refer to this age-adjusted community-rating system simply as a community-rating regime for brevity's sake. As we noted in Handel et al. (2015), community rating creates adverse selection because insurers are not allowed to price based ob health conditions that a consumer knows. We compare community-rating regimes with risk-based pricing regimes in which insurers price based on consumer health risk, eliminating (or, at least, ameliorating) adverse selection.

4.2 Premium subsidies

Under the ACA, low income consumers are eligible for subsidies that cap premiums for exchange insurance as a share of income. Figure 5a displays the level of these caps for each income level under the 2019 Federal Poverty Line Premium subsidies generate price variation that is helpful in estimation. Roughly 60% of exchange-eligible consumers qualified for premium subsidies, and the average (pre-subsidy) premium in 2019 was about $5,800 per year. At an income level of 133% of the federal poverty line (FPL), the maximum premium exhibits a discontinuous jump from 2.08% to 3.11% of income. The cap then increases continuously with income until income reaches 400% of FPL, at which point the cap will be approximately $4,920 per year. Figure 5b shows the distribution of maximum premiums within the subsidy-eligible population.

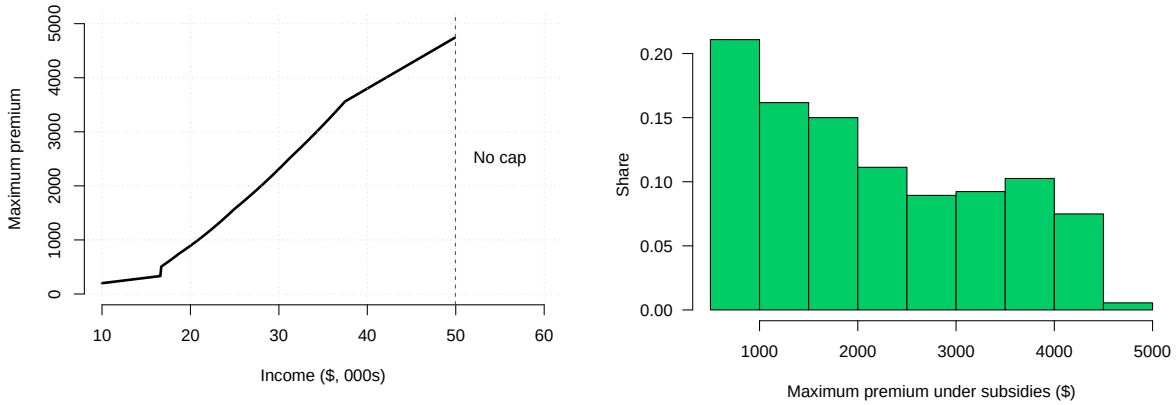
4.3 Cost-sharing restrictions

Plans marketed under the ACA Silver tier also feature cost-sharing restrictions (CSRs) that reduce coinsurance rates for consumers with low incomes. Table 2 reports the cost-sharing features of the Silver tier for income groups that are defined by cutoffs. Group (IV) is not subsidized; hence, its effecive actuarial value remains at 70%, the designated level for Silver plans before CSRs. Each level of cost-sharing reductions has a specified actuarial value, deductible,

Figure 5: Premium subsidies

(a) Maximum premium by income (2019)

(b) Maximum premium after subsidies among subsidy-eligible population



and cap on OOP expenses ("Maximum OOP"). Silver plans also specify coinsurance schedules for each of inpatient care and physician visits. These changes in cost-sharing rates at income thresholds creates useful variation for estimating the extent of myopia (undervaluing of OOPs).

Table 2: Cost-sharing levels under the ACA Silver plan

	(I)	(II)	(III)	(IV)
Lower income cutoff	\$0	\$17,235	\$22,980	\$28,725
Actuarial value	94%	87%	73%	70%
Deductible	\$0	\$250	\$1,750	\$2,000
Maximum OOP	\$1,000	\$2,000	\$4,000	\$5,500

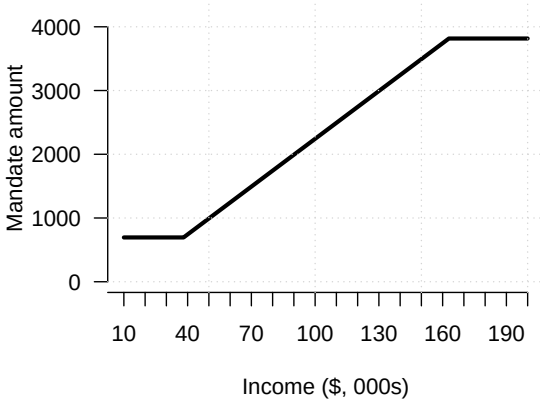
4.4 Individual mandate

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 6 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.

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 US individual health insurance

Figure 6: Individual mandate penalty by income (2018)



market. In addition to adverse selection (due to community rating), the model features (i) public premium subsidies, (ii) various forms of government penalties designed to deter consumers from non-participation, (iii) uncompensated care paid by other parties such as public programs, charitable organizations or providers, (iv) distaste for insurance, and (vi) 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 purchasing health insurance in the individual market. Each consumer j is between 25 and 64 years old; consumers under the age of 25 often remain on their parents' health insurance plans under the ACA whereas consumers older than 64 are eligible for Medicare. In each year t, consumer j has a concave utility index $U_{jt}(\cdot)$ over total consumption. We specify U_{jt} as a constant absolute risk aversion (CARA) utility function

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

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

Consumption is a function of the consumer's action $a_{jt} \in \{0, 1\}$ in the beginning of the year, where $a_{jt} = 0$ and $a_{jt} = 1$ respectively represent being uninsured and insured during year t. Consumption c_{jt} is given by

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

where y_{jt} is j's income during year t and p^c_{jt} is the total premium net of government subsidies paid by insured consumers. (Consumers pay $p^c_{jt} \leq p_{jt}$ where the government covers the gap between p_{jt} and p^c_{jt} as a subsidy.) $OOP^c_{jt}(a_{jt}, m_{jt})$ captures the individual's out-of-pocket medical expenses; these may fall below the individual's total medical expenses m_{jt} due to insurance coverage, government programs (including cost-sharing regulations), and/or uncompensated care.

As in (Abaluck and Gruber, 2011) we allow for the possibility that consumers may discount a dollar in out-of-pocket expenses relative to a dollar in premiums, generating a gap between expenses as perceived by the consumer OOP^c_{jt} and OOP_{jt}. We formulate this difference by letting

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

where $\mu_{jt} \in [0, 1]$ depicts what fraction of OOP_{jt} consumer jt takes into account when making insurance choices.

The term z_{jt} governs consumer tastes for insurance that are unrelated to the pecuniary elements of premiums and coverage of medical expenses. We assume this taste is a linear function of observed consumer characteristics:

$$z_{jt} = \alpha' x_{jt}, \tag{4}$$

where x_{jt} includes income, census region, and year in addition to an intercept term.

Finally, l_{jt} denotes a penalty charged to the uninsured such as the ACA individual mandate penalty.

At the beginning of a year, when the consumer chooses whether to purchase insurance, each component of consumption in equation (2) is known to the consumer except $OOP^c_{jt}(a_{jt}, m_{jt})$. As in Handel et al. (2015), at the time of purchasing insurance both the consumer and insurers know the consumer's *health risk* $_{jt}$, which determines the joint distribution over end-of-year expenses m_{jt} and the following year's health status $h_{j,t+1}$ according to $f(\cdot, \cdot|h_{jt})$. The consumer has no additional information on their health risk than the commonly known h_{jt}. In the counterfactual analysis, we will discretize the continuous measure of health risk h_j into discrete groups $\lambda \in \{1, ..., 7\}$, with a health score of $\lambda = 1$ representing lowest possible health risk and $\lambda = 7$ representing the highest possible health risk.

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

$$a^*_{jt} = \arg \max_{a \in \{0,1\}} \mathbb{E}U_{jt}\Big(c_{jt}(a_{jt})\Big). \tag{5}$$

where utility function $U_{jt}(\cdot)$ and consumption $c_{jt}(a_{jt})$ are described by equation 1 and 2 respec-

tively.

5.2 Insurers

Insurers partially protect their enrollees against uncertainty in realized medical costs m_{jt}. We assume that insurers are risk neutral and thus only consider the expected cost $c(h_{jt})$ of insuring a consumer of health risk h_{jt}. Given that the financial characteristics of ACA insurance plans are regulated, we assume that these expected costs are exogenous primitives of the model and do not vary across insurers. Furthermore, we assume all insurers in the market are homogeneous in terms of their quality and network of providers. Insurers charge premiums p_{jt} to each consumer j in year t. Note that p_{jt} may differ from p^c_{jt} in equation (2) due to government subsidies. As we will detail later, p_{jt} is determined by the insurer's profit maximization constrained by the regulatory restrictions from the government. We do not index premiums by insurer as the model equilibrium does not feature across-insurer heterogeneity in premiums. Last, the relationship between premiums and consumer characteristics—such as health status and age—depends on the regulatory regime. (In our estimation, it is governed by the ACA community-rating regulations; our counterfactuals consider other possibilities.)

5.3 Equilibrium among insurers

We assume that the insurance market is imperfectly competitive but do not explicitly model the nature of imperfect competition among insurers. Rather, for simplicity assume 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 $\mathbf{p}^* = \{p^*_{jt}\}_{jt}$ that they charge yields a common and fixed markup of $\tau \geq 0$. In our setting, the markup τ reflects 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.

Our notion of equilibrium, in which both demand and pricing depend on market primitives and regulation, allows 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. Government subsidies on premiums and cost sharing, the age-adjusted community standard, and uninsurance penalty regulations all impact participation. Finally, market frictions such as a distaste for insurance by consumers, myopic discounting of out-of-pocket expenses (Abaluck and Gruber, 2011), and uncompensated care (Finkelstein et al., 2019) also impact participation. The model with all these estimated ingredients will allow us to assess the relative importance of these frictions and their interactions.

5.4 Welfare measures

We want to measure not only the participation induced by different regimes, but also consumer and social welfare. The presence of myopia requires distinguishing paternalistic from myopic welfare.

5.4.1 Consumer welfare

There are two ways to measure consumer welfare from the perspective of a social planner, depending on the planner's stance on whether the "myopic" discounting of out-of-pocket expenses documented by Abaluck and Gruber (2011) requires "correction." We start with a "myopia neutral" consumer-welfare measure that takes the stance that if a consumer acts as if they underestimate the magnitude of their out-of-pocket expenses, then the welfare analysis should too. Thus, this welfare measure does not seek to "correct" this irrational behavior.[5] We then introduce a "myopia correcting" welfare measure which takes a paternalistic view aiming to correct irrationality.

The first step in constructing the myopia-neutral welfare measure is to compute a dollar-value welfare for each person in the population. This measure, W^N_{jt}, is defined by

$$\Sigma_{k=0,\ldots,64-age_{jt}}\left(\beta^k \times U_{j,t+k}(W^N_{jt})\right) = \Sigma_{k=0,\ldots,64-age_{jt}}\left(\beta^k \times \mathbb{E}U_{j,t+k}\left(c_{j,t+k}\left(a^*_{j,t+k}\right)\right)\right) \quad (6)$$

In words, consuming the sure amount W^N_{jt} starting in period t for k periods until she leaves the market at age 64, makes the consumer indifferent to the stochastic consumption stream under the equilibrium $(g^*, \mathbf{a}^*, \mathbf{p}^*)$, where g^* represents the regulatory regime, where $\mathbf{a}^*$ represents the consumer's insurance purchase decisions and $\mathbf{p}^*$ represents insurance prices (both over time).

Given the individual welfare measures W^N_{jt} we construct a population-level myopia-neutral welfare metric. Our metric, denoted W^N is constructed by taking the average of W^N_{jt} among all individuals at a fixed snapshot t:

$$W^N = \mathbb{E}_j(W^N_{jt}) \quad (7)$$

Note that the left-hand-side is not indexed by t even though the right-hand-side is. This is because, in steady state, the welfare measured is the same at any time snapshot.

We now turn to constructing a myopia-correcting welfare measure W^C. The steps of constructing this measure are the same as those for W^N with one modification. In equation (6), the utility function uses a "corrected" measure of consumption levels. Instead of the formula for c_{jt} from equation (2), it uses corrected consumption c'_{jt} defined as follows:

[5] This welfare measure would be appropriate if the source of "myopia" is actually financial frictions the consumer faces.

$$c'_{jt}(a_{jt}) = y_{jt} - p^c_{jt}(a_{jt}) - \frac{1}{\mu} \times OOP^c_{jt}(a_{jt}, m_{jt}) - l(a_{jt}) \tag{8}$$

Recall that μ is the factor by which consumers discount the magnitude of OOP expenses when making insurance decision. This formula replaces the perceived out-of-pocket expenses OOP^c with the true expenses $\frac{OOP^c}{\mu}$. As a result, if a regulator attempts to maximize this welfare measure, it means that she, among other things, attempts to protect consumers against their "behavioral mistakes."

5.4.2 Social welfare

Myopia neutral social welfare $W^{s,N}$ is defined as the sum of consumer welfare and annual monetary surpluses by the government and insurers:

$$W^{s,N} = W^N + S^i + S^g \tag{9}$$

where S^i is insurers' annual per-capita profit and S^g is the net annual per-capita surplus to government and charitable organizations. In a similar fashion, one can define the myopia-correcting social welfare measure $W^{s,C}$ as follows:

$$W^{s,C} = W^C + S^i + S^g$$

5.4.3 Interpretation

Notice that equations (6) and (7) first translate utils of each individual consumer into a dollar measure and then take the expectation across consumers. We intentionally depart from what is customary in the literature (see Ghili et al. (2021) for instance), which involves reversing the two steps: in **?**, consumer-welfare measures are calculated by first taking the expectation over ex-post utilities and then translating to a dollar amount.

The rationale behind our construct is to separate the allocative and redistribution impact of a regime. The proposed welfare measure above effectively takes the stance that regulation of this market is a separate matter from income redistribution. If we computed welfare by first taking the expected utility across consumers and then dollarizing the value, our analysis would favor any policy that redistributes income to lower-income consumers. In other words, the insurance aspect of the problem would be dwarfed by the income redistribution aspect were we to adopt the more commonplace approach to welfare. Our approach evaluates market regulation of the insurance market separately from income redistribution.

Finally, notice that our baseline welfare measure reflects society at the time of transition. At the time of a reform, society is composed of individuals of all ages, whose welfare is differentially impacted. We could instead measure welfare from the perspective of the 25 year old, namely, individuals welfare for their whole working life. These measures might differ. Consider a reform

that favors younger at the expense of older consumers. From a 25-year-old's perspective the reform might be welfare improving as she discounts the far future, while the reform might be unpopular in society as a whole. Comparing both welfare perspectives highlights who benefits from different regimes.

6 Estimation

6.1 Overview

The estimation proceeds as follows. We construct a measure h_j of MEPS panelists' health risk on the basis of their medical conditions. In estimation, we use a continuous measure equal to the panelist's log predicted total medical expenses when insured, although we discretize this measure into a seven-point sccore λ_j for the purpose of solving our dynamic model. To use the discretized health scores in analyzing our equilibrium model, we also estimate health state transitions using MEPS panelists' computed scores. With measures of MEPS panelists' medical risk in hand, we estimate OOP distributions conditional on medical risk, income and insurance status. These estimated distribution reflect the uncertain health expenses consumers of a specific health and income, face under the event they purchase or do not purchase coverage. Finally, we estimate a binary choice model of insurance take-up, using the estimated conditional OOP distributions as an input.

6.2 Determining medical risk

We summarize MEPS panelists' medical risk using their predicted log total medical expenses based on their medical conditions. We predict expenses the panelist incurs in round 5, based on the medical conditions recorded in the panelists' first four rounds in the panel. Predicted expenses, reflect the panelists' risk of future medical need. This is because future medical needs relevant to consumers and insurers at the time of buying coverage. We predict expenses incurred in round 5 so that we can use as long of a time frame as possible—the duration of a panelist's first four rounds—to obtain as complete of a record of a panelist's medical conditions as we can.

Given that panelists spend differing lengths of time in round 5 (typically 4–6 months), we annualize their round 5 expenses by dividing these expenses by the number of months spent in round 5 and multiplying them by 12. The equation that we use to predict total medical expenses is

$$\log(\mathrm{TOT}_j) = x_j'\beta + \varepsilon_j, \tag{10}$$

where TOT_j is annualized total medical expenses (plus one, to ensure positivity) in round 5, and the right-hand side variables x_j include all ICD diagnostic codes, sex, age, and BMI as well as self-reported health, abilities, and mood. By contrast, x_j does not contain characteristics of panelist's j's medical events. We use the Lasso to circumvent problems related to overfitting that affect ordinary least squares in settings with many regressors. The estimation sample

Table 3: Mean OOPs and costs by health score

(a) OOPs

λ	Uninsured	Insured
1	142	117
2	287	198
3	527	306
4	890	455
5	1381	653
6	1992	909
7	3039	1271

(b) Cost to insurer

λ	$c(\lambda)$
1	1001
2	2833
3	4596
4	6273
5	7823
6	9384
7	11129

Note: tables display mean OOPs and mean costs to the insurer among MEPS panelists classified with health status λ and indicated insurance status. Costs to insurer are mean insurer payments for medical services among exchange-insured panelists within each health status.

includes only insured panelists. [6] The R^2 of the regression is 0.41. Based on the regression, we compute predicted log expenses

$$h_j = x_j' \hat{\beta} \tag{11}$$

for each MEPS panelist. Importantly, we compute predicted log expenses, h_j, for both insured and uninsured MEPS panelists using the lasso regression coefficients.

Finally, we make an adjustment to our h_j measures for uninsured individuals to reflect potential adverse selection on unobservable health risk. More specifically, we find that conditional on our x_j variables, the average uninsured consumer in our data has lower realized medical expenses TOT_j relative to the average insured one. We interpret this to mean that conditional on observables, an uninsured individual is in expectation healthier than an insured one, due to reasons unobservable to us. For that reason we allow for an additional dimension of consumer heterogeneity, observed by the consumer and used to make her insurance choice, but unobserved by the econometrician. As a result, we "adjust down" the h_j measures for the uninsured in a way that rationalizes their lower spend given observables. Appendix C details this adjustment.

Our (adjusted) predictions h_j provide us with a continuous ex-ante health score for each individual. For our equilibrium analysis, we discretize them into 7 groups. We choose thresholds that evenly divide insured panelists into seven groups by score h_j. These groups constitute our 7 health status groups, denoted λ from 1 to 7. Table 3 reports mean OOPs and costs to the insurer by health and insurance status.

[6] For the uninsured unpaid expenses may reflect list prices, of little relevance, while the amount paid may reflect write downs, also not reflective of the cost of treatment.

6.3 Conditional OOP distributions

Key to our exercise are OOP distributions that determine the risk associated with going uninsured upon taking into account uncompensated care. We estimate a parametric model of OOP expense distributions conditional on health risk, insurance status, and income using the non-discretized measure of health risk described in the preceding subsection. This model captures the extensive margin of whether a consumer incurs positive OOPs using a logit component and the intensive margin governing the magnitude of a consumer's positive expenses using a log normal component. The model also features persistence of expenses across months on both the intensive and extensive margins as well as the dependence of expenses on health risk and income. We estimate this model separately for uninsured and exchange insured consumers. The data used in estimation are a consumer/month-level panel covering the months for which consumers belonged to round 5 of MEPS. We include only these months in order to condition on consumer health risk based on medical conditions recorded in the first four rounds of MEPS. Appendix D details the model and presents maximum likelihood estimates of its parameters.

Figure 7 plots quantiles of annual OOPs for a consumer with the median income across the seven discretized health scores for both insured and uninsured consumers; for each score, we compute mean continuous health risk h_j among the consumers sharing that score, and compute the quantiles of OOPs conditional on that score. The figure shows that OOPs increase with health risk, and that the uninsured have generally higher levels of OOPs than the insured. Figure 8 displays mean OOPs across various income levels for a consumer with the mean health risk among consumers in the highest of the seven discrete health groups. We see that income plays a relatively minor role in shaping OOPs for high risk consumers.

6.4 Health status transitions

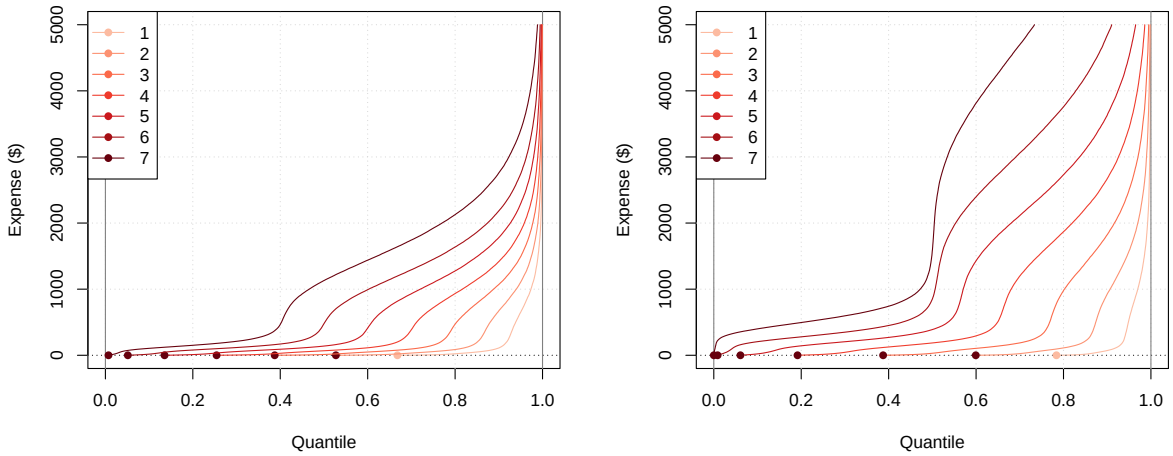
Welfare calculations requires probabilities $f_t(\lambda_{jt}|\lambda_{j,t-1})$ of health status transition. To estimate these probabilities, we first compute each MEPS panelist's discrete health score $\lambda_{j,t-1}$ based on medical conditions reported in their first year in the sample and another score λ_{jt} based on medical conditions reported in their second year in the sample.[7] We then estimate the probability of transitioning from health score λ to λ' as the empirical frequency with which MEPS panelists with health score λ in their first year in the panel switch to health score λ' in their second year. We estimate these transition probabilities separately for each of the five year age groups 25–29, 30–34, and so on.

Given that we computed predicted log expenses based on one year of medical conditions rather than four rounds of medical conditions here, the predicted log expenses that we compute on the annual data are not directly comparable to those that we compute in determining panelists'

[7]Note that to compute a health status for consumer i in year t, we no longer concern ourselves with the "rounds" in which the consumer's expenses were measured. The key function of the analysis that used those rounds was to help estimate $\hat{\beta}$ in equation 11 and construct predicted expenses. Once equipped with this $\hat{\beta}$ and the thresholds for h_j that determine the 7 health statuses, one can directly measure health statuses on an annual basis for any it by plugging the right x_j into equation 11.

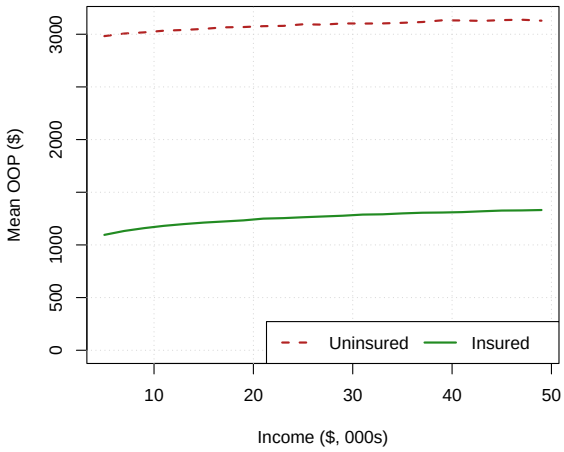
Figure 7: Estimated conditional distributions of annual OOPs

(a) Insured consumers (b) Uninsured consumers



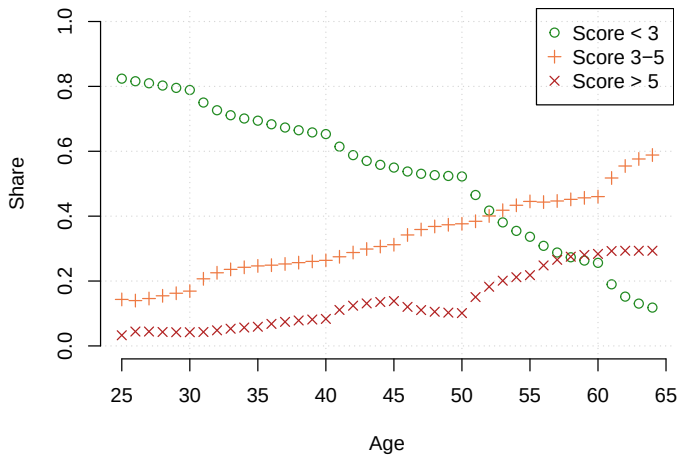
Note: the darker curves correspond to higher levels of health risk. The OOP quantiles are computed for a consumer with the median income in the population eligible for exchange insurance. The shaded circles on the horizontal axis indicate the share of consumers with zero OOPs in the course of a year among those consumers with a particular health risk level and insurance status.

Figure 8: Mean OOP by income for high risk consumer



Notes: this figure plots the mean OOP across varying levels of income for an uninsured and an insured consumer with the mean health risk of a consumer in the highest risk of the seven discrete risk groups.

Figure 9: Health status shares by age implied by estimated transition probabilities



discrete health scores as described in Section 6.2. In particular, the thresholds that partition the exchange-insured population into seven equally sized groups based expenses predicted from four rounds of conditions do not partition observations of expenses predicted from annual medical conditions into seven equally sized groups. Thus, we use different thresholds in estimating transitions; we choose thresholds such that the distribution of discretized health scores in a population simulated from the estimated transition probabilities and the initial distribution of health scores in the age 25 population is the same as the the distribution of discretized health scores in the observed population.

Figure 9 plots the share of consumers in various health score groups as implied by the estimated transition probabilities and the initial distribution of age 25 consumers across discrete health scores.

6.5 Choice model

To recover preferences for insurance, we estimate a binary choice model of insurance coverage. The dataset used in estimation is a repeated cross-section of MEPS panelists who were exchange eligible: they were between 25 and 64 years of age and had incomes above 138% of the Federal Poverty Line. Additionally, we subset to consumers who were either insured via an exchange for their entire second year in MEPS or uninsured for their entire second year. The outcome variable used in estimation is insurance status in the second year. We compute each panelist's health status group using the procedure described in Section 6.2. Last, we include MEPS panelists whose second year in the panel was one of 2015, 2016, 2018, or 2019; we do not include 2017 in the estimation sample because MEPS switched from using ICD-9 to ICD-10 diagnostic codes in this year, which means that panelists who began in MEPS in 2016 and ended in 2017 are assigned codes from one code system in their first year in MEPS and another in their second

year. This complicates computation of health risk among these panelists. Also, we exclude 2014, which was the first year in which the ACA exchanges operated, because low uptake that year (under 10%) may have reflected idiosyncratic factors associated with the launch of the exchanges (low consumer information about the exchanges and administrative problems in the roll-out of exchanges) rather than persistent aspects of consumer preferences. This sample includes 4173 panelists, 3154 who were uninsured and 1019 who were insured.

Our data is generated under ACA rules. Because the decision to get coverage under ACA rules is static, a consumer with health h_j purchases insurance if her expected utility under insurance is higher than under no insurance. That is, consumer j purchases insurance when

$$-\mathbb{E}[\exp\{-\gamma_j(y_j - p_j^c - \mu_j OOP_j^1(1) + z_j)\} \mid h_j, y_j] \geq -\mathbb{E}[\exp\{-\gamma_j(y_j - l_j - \mu_j OOP_j(0)\} \mid h_j, y_j], \tag{12}$$

where the left hand side reflects expected utility if insured and the right hand equals expected utility when uninsured. The expectations are taken over distributions $\text{OOP}(a)$ for $a \in \{0, 1\}$ conditional on health risk h_j and income y_j. Given that we estimate the model on a repeated cross section rather than a panel, we suppress the year subscript t. Recall that we parameterize tastes for insurance z_j that are not related to plan financial characteristics as

$$z_j = \alpha' x_j,$$

where x_j includes indicators for consumer j's membership of various income groups (under \$20k, \$20k to \$30k, \$30 to \$40k, ..., \$70k to \$80k, and above \$80k), indicators for x_j residing in various US census regions (Northeast, Midwest, South, and West), and indicators for years of observation (2015, 2016, and 2018; 2019 is the omitted category). Participation in the ACA exchanges increased over time, which may be attributable to increases in consumer knowledge of the exchanges and by improvement of exchange infrastructure (e.g., exchange websites). The dependence of z_j on observation year allows the model to capture such factors that boosted exchange participation over time. Last, p_j^c is the after-subsidy premium facing consumer j whereas l_j is the individual mandate penalty faced by the consumer.

We allow for consumer heterogeneity in myopia as governed by μ_j by letting it differ with income:

$$\mu_j = \begin{cases} \mu_{\text{low}}, & \text{income}_j \leq \$30{,}000 \\ \mu_{\text{high}}, & \text{income}_j > \$30{,}000. \end{cases}$$

There are several reasons to think that low-income consumers may underweight OOPs relative to premiums. These consumers may have reduced access to information about their medical risk. They may also lack the liquid assets required to make upfront premium payments. The model can capture consumer inability to pay for premiums by assigning a larger weight to upfront premium payments. Last, it is possible that a lower μ_j among low-income consumers represents the fact that such consumers are more likely to rely on financial aid from family to make their OOPs; such help differs from uncompensated care given that MEPS includes

payments made by family in their measurement of OOPs.

Equation (12) permits the simulation of the consumer's choice for values of the unobservable risk aversion γ_j. We simulate as such to compute our simulated method-of-moments estimator. We choose the model parameters $\theta = (\mu, \alpha)$ to match the following moments in the observed data

1. The covariance between predicted log medical expenses (continuous health score) and participation among consumers with incomes under \$30,000 and among consumers with incomes exceeding \$30,000;
2. The participation rate among consumers in each of the following income brackets: under \$20k, \$20k to \$30k, \$30 to \$40k, ..., \$70k to \$80k, and above \$80k.
3. The participation rate among consumers in each of the following census regions: the South, the West, and the Midwest.
4. The participation rates in the years 2015, 2016, and 2018.

We approximate the model-implied values of these statistics by simulation, simulating each consumer's choice for five draws of γ_j, conditioning on observed incomes, premiums, and census regions throughout. Note that risk aversion is the sole unobservable that generates variance in insurance choice conditional on observables in the model.

6.6 Identification

The objects estimated in Sections 6.2 through 6.4 are measured directly from our data: medical risks and conditional OOP distributions, uncompensated care, and health status transitions. Such direct measurement makes the identification of these objects straightforward. As a result, we focus our discussion here on objects identified through the choice model in section 6.5.

Our goal in estimating the choice model is to quantify frictions that explain limited uptake of exchange health insurance plans. One such friction is myopia, i.e., incomplete consumer internalization of out-of-pocket payments ($\mu_j < 1$). When μ_j is low, the consumer's perceived payoffs from buying insurance have a low degree of dependence on the consumer's health risk; when μ_j is high, health risk exercises a high degree of influence on perceived payoffs from buying insurance. Thus, the relationship between health risk and particpation is informative about μ_j. We pin down μ_j for low- and high-income consumers by matching observed covariances between health risk and participation.

The z_j term in consumer utility represents preferences for insurance that are unrelated to insurance plans' financial characteristics. We specify that z_j depends on income, census region, and year, and we pin down the coefficients on these attributes in the linear specification of z_j by matching the observewd participation rates within groups defined by income, census region, and observation year.

We finish this subsection by noting two points. First, there is no explicit "price/premium sensitivity" parameter in our model. We treat health insurance as a purely financial product, which implies that any individual i shall make a purchase decision by comparing distributions of financial outcomes between scenarios in which she does and does not take up insurance. Premiums are part of what shapes this comparison between distributions over financial outcomes. As a result, the objects we estimate/calibrate (e.g., CARA coefficient, uncompensated care, myopia), collectively enable the model to predict how an individual's insurance decision would react to the price/premium change.

Second, the effect of uncompensated care on participation is also not separately estimated as a primitive parameter, but is instead determined by our direct measurement of OOP distributions for the insured and uninsured. Once μ is identified from how participation varies with health status, the impact of changes in uninsured OOPs on participation follows mechanically from these distributions together with risk aversion. In this sense, uncompensated care's effect on take-up is mediated entirely through other model parameters (risk aversion, myopia) rather than through a separate parameter of its own.

6.7 Results

Table 4 presents estimates of the choice model parameters. Our estimates of the μ_j parameters imply that consumers underweight OOPs relative to premiums, and that the extent of underweighting varies by income. Consumers with incomes under \$30,000 internalize 59% of OOPs in their insurance uptake decisions whereas consumers with incomes over \$30,000 internalize 85% of OOPs. Lower internalization of OOPs among consumers with lower incomes has several explanations. These consumers face greater informational frictions. Additionally, low income consumers may face greater financial constraints that make the upfront payment of premiums challenging and thus lead them to overweight premiums relative to OOPs when choosing whether to adopt insurance.

The estimates imply that tastes for insurance were lower in the earlier years of the ACA exchanges. This finding may reflect that expanding consumer information about the exchanges as time passed since their establishment, operational improvements in the exchanges over time, or an attenuation of ideological opposition to the exchanges in the years following the ACA's implementation.

Additionally, we find that tastes for insurance increase in income. The element z_j of tastes for insurance that is not related to plan financial characteristics includes a saturated set of indicators for income ranges, and consequently z_j does not include a distinct constant term. Thus, the results imply that z_j is positive for all consumers earning above \$30,000 across years and census region. One reason that consumers may generally prefer to be insured is that it spares the effort involved in obtaining uncompensated care.

Table 4: Parameter estimates

Parameter	Estimate	SE
μ_{low}	0.587	(0.1599)
μ_{high}	0.853	(0.1502)
$\alpha^{\text{inc}}_{20-}$	0.452	(0.3371)
$\alpha^{\text{inc}}_{20-30}$	1.302	(0.3211)
$\alpha^{\text{inc}}_{30-40}$	2.033	(0.4392)
$\alpha^{\text{inc}}_{40-50}$	2.456	(0.6118)
$\alpha^{\text{inc}}_{50-60}$	2.907	(0.6522)
$\alpha^{\text{inc}}_{60-70}$	3.126	(0.6487)
$\alpha^{\text{inc}}_{70-80}$	3.240	(0.9375)
$\alpha^{\text{inc}}_{80+}$	2.458	(0.8725)
α_S	-0.181	(0.3731)
α_W	-0.130	(0.4951)
α_{MW}	-0.036	(0.3368)
α_{2015}	-1.792	(0.6771)
α_{2016}	-1.873	(0.4115)
α_{2018}	-1.157	(0.2838)

Notes: this table provides estimates of the choice model parameters. The α^{inc} parameters are the effects of income group membership on z_j; α_S, α_W, and α_{MW} are the effects of residing in the South, West, and Midwest on z_j; and $\alpha_{2015}, \alpha_{2016}$, and α_{2018} govern differences in mean tastes for insurance betweeen the years indicated by the parameter subscripts and 2019.

6.8 Fit of the Model

We evaluate the fit of of the model by comparing the observed MEPS data to simulated outcomes from the estimated model conditional on observed covariates. Figure 10 shows participation rates by age and by health risk h_j in the observed and simulated data. The model predicts the relation between participation and age, as well as income quite well.

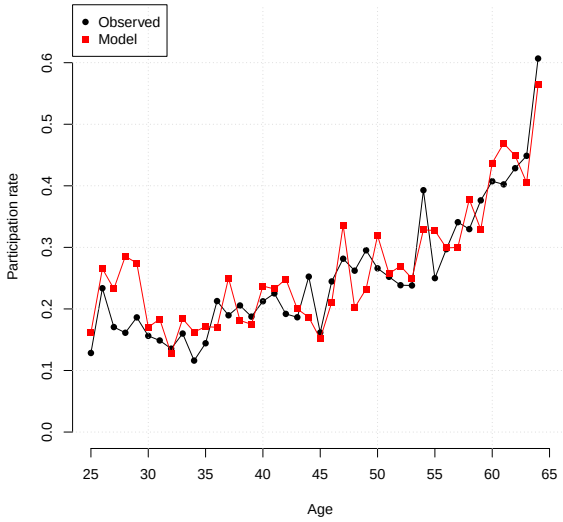
To evaluate the fit, keep in mind that one of the moments we match in the estimation is the correlation between participation and continuous health score. However, we do not match any moment linking participation with age, yet the model replicates the observed profile with a high degree of accuracy. Two points are worth emphasizing. The first is that while we match the correlation between participation and health type, panel b of figure 10 shows are remarkable fit, beyond just correlation. That is, the non-linear relation is picked quite well despite not being rewarded by the estimator. To understand what is behind the models prescribed participation, let us consider the ingredients that determine participation. Participation is determined by willingness to pay for insurance surpassing premiums. Willingness to pay for insurance depends on the OOP distribution given health state, and the risk aversion coefficient. It is quite reassuring that the estimated OOP distribution, passed through the risk aversion distribution (which we do not estimate, but take from elsewhere) end up delivering an apparently tight convex relation between health and participation.

Second, and more remarkable, is the relation between participation and age, which our estima-

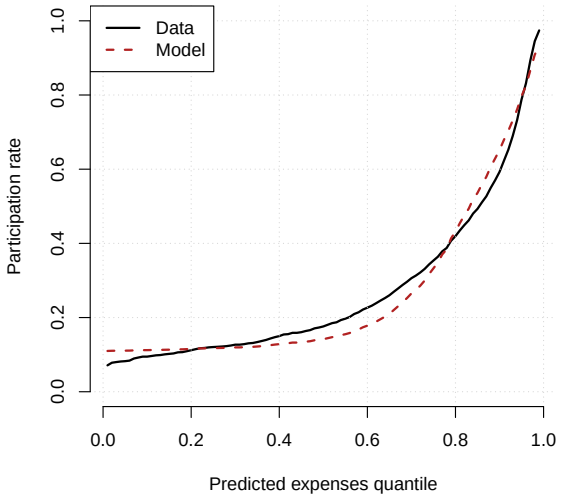
tion does not impose. None of the moments used in the estimation relate participation and age. Naturally, health risk and income increase in age, and the risk aversion estimated in Handel et al. (2015) increases in age, making older consumer willing to pay more insurance. The ingredients are bound to deliver a positive relation. However, the accurate relation between participation and age means that the sources of willingness to pay we infuse into the model, risk and risk aversion types, capture remarkably well the observed willingness to pay for insurance. Recall that both risk aversion and risk type are data, in the sense, that we do not estimate them to fit any moment of the choice model. These different ingredients could have combined to deliver a poor fit. It is reassuring that they deliver predicted behavior so similar to the data.

Figure 10: Participation by age and health risk

(a) Age



(b) Health risk



7 Analysis of participation and welfare

7.1 Construction of population used in counterfactual analysis

In this section, we describe the construction of the simulated population on which we compute exchange equilibria throughout our counterfactual analysis. The starting point is the subsample of MEPS panelists who meet the age- and income-related exchange eligibility criteria described in Section 2.1 and who are either exchange-insured or uninsured for the entirety of their second year in MEPS. Additionally, we limit the subsample to panelists whose second year of MEPS participation was 2019 and use the 2019 FPL in determining subsidy eligibility.

To introduce heterogeneity in risk preferences, we replicate each panelist 20 times, assigning to each replicate a distinct risk aversion coefficient drawn from the age-specific conditional distribution estimated by Handel et al. (2015) as described in Section 2.4. Next, we discretize panelists' health scores h_j into health status groups λ_j. The purpose of this discretization is to

allow the modelling of health risk evolution with health state transition matrices. To discretize health risk, we compute the (weighted) mean health score h_j within each health status group $\lambda \in \{1, 2, \ldots, 7\}$. We then overwrite each panelist's health risk h_j with the mean health risk within their health status group, and use this mean health risk in determining the panelist's conditional distribution of OOPs when insured and when uninsured.

We construct the final simulated population by simulating forward 25-year-old MEPS panelists to the age of 64. For each 25-year-old panelist, we compute their income percentile within their five-year age group (25–29) and assume that this income rank relative to their five-year age group (e.g., 25–29, 30–34, 35–39, etc.) remains fixed throughout their life cycle. Thus, the income of a consumer in age group A who was at the qth income percentile among the 25–29 age group at age 25 would be the qth income percentile among members of age group A. We compute these quantiles among the MEPS subsample described in the first paragraph of this subsection.

Risk aversion evolves with age according to the specification in Handel et al. (2015): each year, a panelist's risk aversion coefficient increases by 3.99×10^{-6}. Forward simulation of the population then proceeds under these assumptions, with health state transitions governed by the estimated process described in Section 6.4.

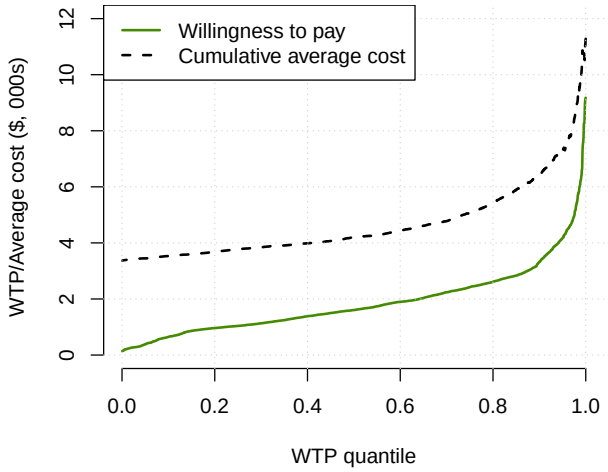
Our reason for simulating forward a cohort of 25-year-olds is to ensure a stationary environment; it guarantees that the cross-sectional distribution of age, health, and medical costs in the simulated population is constant over time. Under stationarity, the current equilibrium premiums continue to be valid equilibrium outcomes in the future. Knowing the path of premiums is necessary to evaluate consumers' lifetime welfare, which is a central object of our analysis.

Our counterfactual analysis evaluates a regime without uncompensated care. We implement this analysis by reallocating uncompensated care to consumers' OOPs. For each bin k defined by insurance status a and health group λ, we compute the mean OOP and mean amount of uncompensated care, form the ratio $R_{a,\lambda}$ of the latter over the former, and then inflate all OOP realizations for consumers in the bin by $1 + R_{a,\lambda}$. This procedure shifts uncompensated care uniformly to OOPs within each health/insurance bin in a multiplicative manner. Note that the uncompensated care amount for a panelist is defined as the sum of observed and unobserved uncompensated care. Observed uncompensated care is the sum of payments received for medical services from sources other than insurance. Unobserved uncompensated care is equal to the panelist's total expenses times $e^{\hat{\beta}} = e^{0.42} = 1.52$, where $\hat{\beta}$ is the estimate of the extent to which providers receive more payments from insured patients than from uninsured patients as described in Section 3 (see the "All" row of Table 1).

7.2 Willingness to pay for insurance

Before analyzing the sources of impediments to participation it is useful to zoom into the key output of our estimation: the distribution of willingness to pay for exchange insurance

Figure 11: Willingness to pay and average costs absent subsidies



Notes: this figure provides, for each value of $\tau \in [0,1]$, (i) the τ quantile of consumers' willingess to pay for insurance, not accounting for premium subsidies, and (ii) the average cost (adjusted for the 20% minimum loss ratio afforded to insurers for overhead and profits) of insuring all consumers with a willingness to pay above this τ quantile.

conditional on health type (as well as age and income).

Figure 11 plots consumers' willingness to pay for insurance (not accounting for premium subsidies) by quantile $\tau \in [0,1]$. The dashed line shows the average cost of insuring those with a willingness to pay above τth quantile. The estimated frictions make consumers' willingness to pay uniformly lower than actuarial costs. Thus, a community rating regime without subsidies would fully unravel: for any premium level, the cost of insuring those willing to pay for insurance exceeds the premium. Thus, we find a similar result to that of Finkelstein et al. (2019) in the context of the Massachusetts Connector, who report that for 70% of the population the average cost to insure them is twice their willingness to pay.

These two lines will play a key role in finding equilibrium under different regimes and rules. It is worth reminding the reader how the lines were derived. The willingness to pay distribution congeals our friction estimates, with our estimates of the OOP distributions, with data on the distribution of health status, and the distribution of risk aversion. The average cost line summarizes our ACG-type model prediction for individuals with different willingness to pay (derived from the inputs just described).

7.3 Impediments to participation

We now unpack the sources of this low willingness to pay. Specifically, 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,

Table 5: Participation rates (%) upon friction elimination

Regime	Current regime (no penalty)	Risk-based pricing
Baseline	38.4	52.5
No myopia	48.8	59.8
No uncompensated care	94.3	97.4
Neither	97.1	99.5

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.

we simulate what equilibrium would be in their absence.

The column labeled "Current regime (no penalty)" in Table 5 describes how the removal 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, 38.4% of eligible consumers purchase exchange insurance. Removing myopia significantly raises participation to 48.8%. However, this increase pales in comparison to the effect of eliminating uncompensated care (making an uninsured consumer incur 100% of their medical expenses), which raises participation to 94.3%. Removing myopia once there is no uncompensated care leads to a small further increase in participation, to 97.1%. This analysis indicates that uncompensated care is the primary friction depressing exchange participation.

The second column in Table 5 describes a similar exercise conducted under the alternative assumption of risk-based pricing, which removes the adverse selection caused by community rating. Were there no frictions at all (as in the classic textbook insurance setting), participation in the presence of premium subsidies and cost sharing would be above 100%. Table 5 shows that with risk-based pricing we estimate that participation in the absence of both myopia and uncompensated care would be 99.5% (note that the model still has a distaste for insurance for some consumers). This 99.5% participation rate compares to a rate of 52.5% when both frictions are present, 97.4% if only myopia is present, and 59.8% if only uncompensated care is present. Two conclusions emerge from this. First, as in the first column of the table where there was community rating, these frictions have a large impact and the vast majority of this is due uncompensated care, whose removal on its own raises the participation rate to 97.4%. Second, absent uncompensated care and myopia, adverse selection due to community rating would not be a significant problem for participation rates: we estimate that adverse selection lowers participation only from 99.5% to 97.1%.[8]

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

Table 6: Participation rates (%) by policy regime

Policy regime	Baseline	No uncomp. care	Risk-based pricing
Mandate and subsidies	51.8	96.4	66.8
No mandate	38.4	94.3	52.5
No subsidies	0.0	86.8	28.3
Neither	0.0	83.9	21.1

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 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.

7.4 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 frictions identified above.[9] In particular, we compare the effects of a mandate penalty (of the magnitude enforced just before its 2019 elimination) against the effects of premium subsidies.

Table 6 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 38.4% (matching Table 5). Adding the mandate penalty raises participation by 13.4 percentage points. However, subsidies have a much larger impact on participation than the mandate penalty. Indeed, the market fully unravels without the subsidies even if there is a mandate penalty. At the same time, at their current magnitude, the premium subsidies are not sufficient to fully overcome the friction from uncompensated care and generate full participation.

The significance of uncompensated care can also be seen in another way in Table 6: if there were no uncompensated care (see the second column), participation would be 83.9% even absent both subsidies and a mandate penalty.

7.5 Welfare analysis

This section conducts welfare analysis. We start by analyzing the baseline scenario that features both premium subsidies and uncompensated care. The subsequent analyses compare policy regimes upon the elimination of subsidies and, next, upon the elimination of uncompensated care. Last, we examine the interactions between subsidies and uncompensated care.

Recall that our welfare measure differs in two ways from measures commonly used in the literature (e.g., Ghili et al. (2021)). First, it aggregates welfare among all individuals rather than

[9] 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 billionCongressional Budget Office (2023, 2017).

Table 7: Welfare by policy regime

Penalty regime	No exchange	Risk pricing	Mandate penalties		
			Zero	2016 ACA	Extreme ACA
Part. (%)	0	52.5	38.4	51.8	100
Social deficit	**1384**	**2056**	**2104**	**1521**	**1296**
Myopic welfare	-2039	-215	0	-437	-683
Patern. welfare	**-5406**	**-551**	**0**	**-110**	**101**
Agg. welfare	**-6790**	**-2607**	**-2104**	**-1631**	**-1195**

Notes: this table reports components of total welfare by policy regime on a per capita basis. Note the following abbreviations: "Uncomp." indicates uncompensated care, and "Patern welfare" indicates paternalistic welfare. "Myopic welfare" is the welfare measure computed using consumers' estimated utility functions inclusive of myopia. "Patern. welfare" is the welfare measure computed under the assumption that consumers ultimately pay for unanticipated medical expenses. Note that the equilibria are always computed under the assumption of myopic behaviour; the paternalistic welfare measure computes consumer welfare based on their decisions made according to myopic utility functions using non-myopic utility functions.

from the perspective of an "unborn" consumer before the age of 25. Second, our welfare measure first translates each individual's welfare from utils into dollars and then aggregates, rather than the other way around. See section 5.4.3 for more details on our choice of welfare metrics.

7.5.1 Welfare analysis under premium subsidies

Table 7 summarizes our welfare results under the status quo regime in which premium subsidies are offered. As points of comparison, we provide welfare results for the regimes discussed in Section 7.3 in addition to a "no exchange" regime in which exchange-eligible consumers have no option to buy insurance and an "extreme ACA" regime with an inescapable mandate to buy insurance (and which thus features full participation). There are two main lessons. First, regimes with higher participation generally deliver greater aggregate welfare. This is because of myopia: higher participation (even if by force) helps correct the mistake that consumers would otherwise make by not enrolling. To verify that this is indeed the mechanism, note that the extreme ACA penalty, which forces full participation, is preferred under the paternalistic consumer welfare measure, and the most unpopular regime under the myopic welfare measure.

The discrepancy between paternalistic and myopic criteria highlights the importance of capturing beahvioral frictions. While the classical model without frictions naturally prescribes full insurance as optimal, with the estimated frictions full participation delivers the lowest consumer welfare. Actually, consumer welfare is maximized under the zero penalty regime, namely, when consumers are left alone.

Second, the reclassification risk associated with risk pricing has a modest impact on welfare. We found in prior work that reclassification risk takes a substantial toll on welfare (Handel

Table 8: Welfare without premium subsidies

Penalty regime	No exchange	Risk pricing	Mandate penalties		
			Zero	2016 ACA	Extreme ACA
Part. (%)	0	21.1	0	0	100
Social deficit	**1384**	**1285**	**1384**	**408**	**-1105**
Myopic welfare	0	51	0	-972	-2030
Patern. welfare	**0**	**24**	**0**	**-983**	**1927**
Agg. welfare	**-1384**	**-1261**	**-1384**	**-1391**	**3032**

et al. 2015), of the order of $3,000/year/person. In the current setting, reclassification risk is less onerous. Myopic consumers even prefer risk pricing to the ACA. The consumer-welfare advantage of a (no-penalty) regime without risk-pricing relative to one with risk pricing is reduced to $551/person/year under the paternalistic criterion and $215/year/person under the myopic one. The chief reasons for the modest gap are premium subsidies and uncompensated care. Premium subsidies provide implicit insurance against reclassification risk to lower income consumers whose premiums rise when transitioning to a high-expense health state. The latter provides implicit within-year insurance to those who, under risk pricing and without penalties, would opt to not insure themselves. Note that uncompensated care is accounted for in the social deficit: the risk pricing regime runs a higher deficit than a full-force ACA penalty.

7.5.2 Role of premium subsidies

Table 8 shows welfare for a regime without subsidies. Similar to the previous table, here we normalize the welfare measures for the no penalty regime to zero (note that this precludes meaningful comparisons between welfare numbers in Tables 7 and 8). Without subsidies, aggregate welfare is higher under risk pricing than under the non-extreme ACA regimes. The reason is that the non-extreme ACA community rating regimes now feature zero participation and thus can be improved upon by any non-degenerate insurance market. Once again, full insurance is the preferred paternalistic outcome but delivers the lowest welfare to myopic customers.

7.5.3 Role of uncompensated care

Table 9 shows participation and welfare results when there is no uncompensated care. We draw several conclusions form this table. First, there is low variability in welfare across the community rating ("Mandate penalties") regimes. This owes to the fact that there is little variation in participation rates across these regimes; even under the lowest-participation regime, 94.3% of consumers purchase exchange insurance. Although the risk pricing regime features a relatively high rate of participation (97.4%), it subjects consumers to reclassification risk given the dependence of premiums on consumers' health states. This means that it delivers

Table 9: Welfare without uncompensated care

Penalty regime	No exchange	Risk pricing	Mandate penalties		
			Zero	2016 ACA	Extreme ACA
Part. (%)	0	97.4	94.3	96.4	100
Social deficit	**0**	**1615**	**1310**	**1255**	**1258**
Myopic welfare	-40387	-919	0	-14	0
Patern. welfare	**-60440**	**-1055**	**0**	**116**	**281**
Agg. welfare	**-60440**	**-2670**	**-1310**	**-1139**	**-977**

Notes: see the notes for Table 7.

a lower level of consumer welfare than the community rating regimes that insure consumers against reclassification risk. Reclassification risk is especially harmful to consumers absent uncompensated care because it deprives them of the option of switching out of the market when facing high premiums to obtain the free, implicit insurance afforded by uncompensated care.

8 Conclusions

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 about 10 percentage points in the baseline specification, and 7 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, participation is low even under risk pricing ($\approx 20\%$) given that adverse selection alone depresses participation less than the combination of uncompensated care and myopia. Moreover, subsidies are necessary to sustain participation because of uncompensated care. Participation would be high absent uncompensated care, even if without subsidies.

The welfare analysis highlights the importance of capturing behavioral frictions. While an ACA-type exchange that forces full participation is the preferred outcome of fully rational consumers, full participation delivers the lowest welfare in the eyes of myopic customers who discount OOPs relative to premiums.

We find that welfare is tightly connected to participation. This owes to the fact that participation increases, however achieved, limit the extent that myopic consumers make costly mistakes. We additionally find that the reclassification risk induced by risk pricing is not as harmful to consumer welfare as previously found using the fully rational model Handel et al. (2015). Both subsidies and uncompensated care attenuate the welfare losses associated with risk pricing. Subsidies absorb premium hikes, while uncompensated care sharply limits the harms from going uninsured in high-risk health states when premiums are the highest under risk pricing.

Absent uncompensated care, participation is high across the various regimes, delivering similar aggregate welfare levels. One exception is risk pricing: absent uncompensated care, reclassification risk is especially painful as consumers can no longer escape high premiums by relying on the implicit, free insurance of uncompensated care. Thus, aggregate welfare is especially low under risk pricing when uncompensated care is eliminated, consistent with previous findings in the literature.

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.

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.

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.

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.

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.

Ghili, Soheil, Benjamin Handel, Igal Hendel, and Michael D Whinston (2021) "Optimal long-term health insurance contracts: characterization, computation, and welfare effects."

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.

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

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.

A Federal Poverty Line

Table 10 describes the parameters of the Federal Poverty Line (FPL) in each year from 2012 to 2019. Recall that, for a household with n members, the FPL in year t is

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

Table 10: Federal Poverty Line, 2012–2019 ($)

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

B Calculation of unobserved uncompensated care

In calculating unobserved uncompensated care, we first determine the extent to which medical providers recover higher payments from the insured than the uninsured as a function of medical event characteristics. To do so, we compare expenses among the insured and uninsured. We make this comparison on datasets of medical events than span the years 2013 to 2019. For the inpatient visit, outpatient visit, emergency room visit, home health visit, and office-based visit categories, we estimate the following regression:

$$\log(y_i) = \sum_j \delta_j \left(\text{INS}_i \times x_{ij}\right) + \beta' x_i + \varepsilon_i,$$

where y_i is the sum of the payments received by the provider for event i, x_{ij} is a characteristic of event i, and INS_i is an indicator for the patient of event i being insured. The x_{ij} characteristics differ across event type and are numerous. We enumerate them below. The δ_j parameters govern how the insured/uninsured gap in payments received varies with event characteristics. The β parameters govern systematic differences in total fees across events with different characteristics. With the regression results in hand, we compute the fitted values of log expenses $\log(y_i)$ for all events involving uninsured panelists. We also compute the predicted log expenses for these events when their underlying panelists are counterfactually insured. Let LE_i^1 denote these expenses under counterfactual insurance and let LE_i^0 denote the fitted values. Next, define $r_i = \exp(LE_i^1 - LE_i^0)$. Unobserved uncompensated care at the visit level is then defined as

$UC_i^u = (r_i - 1) \times y_i$. To obtain a panelist's total unobserved uncompensated care, we sum UC_i^u across all events i involving the panelist in question.

We treat prescription drugs slightly differently: for each unique drug i, we divide the mean expense across the insured for drug i by the mean expenses across the uninsured for drug i. Call this quantity r_i. Unobserved uncompensated care among the uninsured for drug i is $UC_i^u = (r_i - 1) * y_i$, where y_i is the panelist's expenses (from all sources) for drug i.

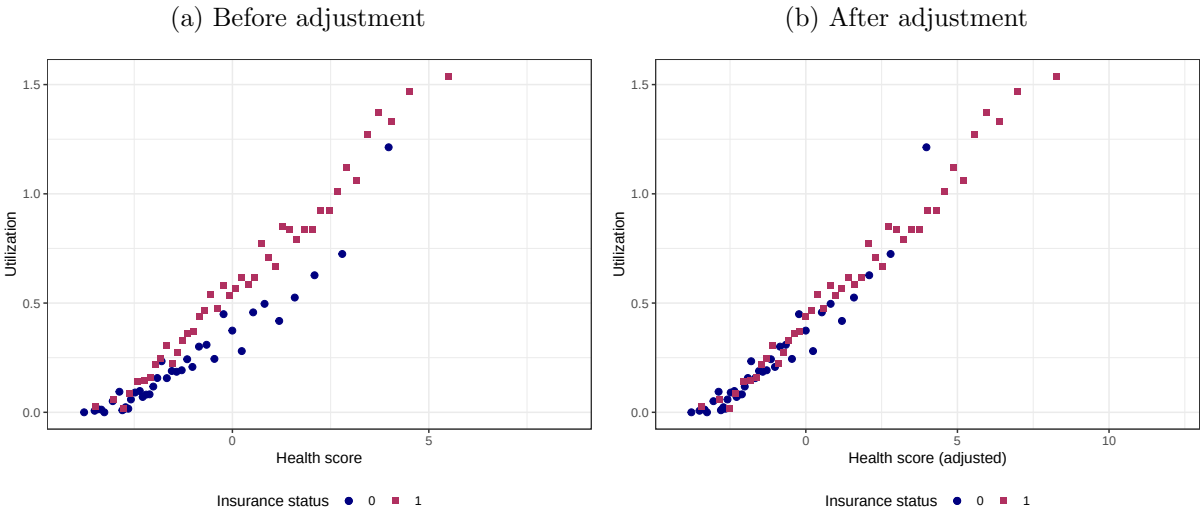
The x_{ij} characteristics are listed below.

- For inpatient visits: indicators for an operation/surgery, for visit originating from an emergency room visit, for the visit being related to a medical condition, and for the reason for admission (see here). Another characteristic is the log number of nights in hospital.
- For outpatient visits: indicators for mammogram, MRI, vaccination, the patient having seen a doctor in the visit, sonogram, surgery or operation, swab, X-ray, visit being related to a medical condition, and for the prescription of a medicine during the visit. Also included are indicators for visit categories (these are "general checkup," "diagnosis or treatment," "emergency," "psychotherapy/mental health counselling," "follow-up or post-operative visit," "immunizations or shots," "vision exam," "pregnancy-related," "well child exam," "laser eye surgery," and "other") and the speciality of the doctor seen (if any).
- For office-based visits: indicators for lab tests being conducted, EKG or EEG, mammogram, medical prescription, MRI, vaccination, patient seeing a doctor, patient speaking to provider over the phone rather than in person, surgery or operation, swab, X-ray, visit being related to a medical condition, doctor's specialty, type of medical professional seen, and category of visit (these are the same categories as specified for outpatient visits).
- For emergency room visits: indicators for lab tests, EKG or EEG, mammogram, sonogram, medicine being prescribed, MRI, vaccination, other service, surgery or operation, swab, X-ray, visit being related to medical condition, and visit categories.
- For home health visits: indicators for various type of care professionals being present, for visit being related to a medical condition, and for frequency of visits.

C Adverse selection correction

We find that, conditional on our health score of predicted log total medical expenses when insured, consumers who are in fact insured have greater medical utilization than those who are not insured. Figure 12a displays this difference when utilization is defined as the log of the number of medical visits (inpatient, outpatient, emergencyroom, and office-based) in round 5 plus one. We attribute the higher medical utilization among the insured to an unobservable component of health risk that drives adverse selection: the uninsured have lower health risk,

Figure 12: Medical utilization and health risk



Notes: this figure contains bin-scatter least squares regression plots that characterize the relationship between medical utilization, defined as the log of the number of medical events in round 5 plus one, and panelists' predicted log medical expenses ("health score"). Figure 12a displays the relationship for unadjusted health scores whereas Figure 12b displays the relationship for adjusted health scores.

a part of which is unobserved, and this drives their decision not to purchase insurance. We correct for unobservable health component by adjusting uninsured consumers' health scores h_i as follows: for a given uninsured-consumer health score h^U, the health score h^I of an insured consumer that yields the same expected utilization n_i is

$$\mathbb{E}[n_i \mid h^I, a_i = 1] = \mathbb{E}[n_i \mid h^U, a_i = 0] \tag{13}$$

We specify the conditional expectation function (CEF)

$$\mathbb{E}[n_i \mid h_i, a_i] = \phi + \alpha a_i + \beta h_i + \gamma a_i h_i$$

Then, equating the CEFs (13) establishes that after transforming uninsureds' health scores by

$$h^U \mapsto \frac{h^U - (\alpha/\beta)}{(\beta + \gamma)/\beta},$$

expected utilization of insured equals that of uninsured (conditional on h_i). Figure 12b shows that, after making this adjustment, the relationship between our health risk measure and medical utilization does not noticeably differ between the insured and uninsured.

D Conditional OOP distribution specification

In this appendix, we detail our specification of the conditional OOP distribution model and report estimates of its parameters. The model captures persistence of expenses across months in both (i) the extensive margin of whether a consumer incurs any expenses whatsoever in the course of a month and (ii) the intensive margin of the magnitude of a consumer's expenses conditional on having positive expenses. We take care to model the process by which consumers incur zero expenses in a month given that this occurs at the high rate of 70% among consumers eligible for exchange insurance. Persistence of expenses across months affects the stakes associated with the insurance choice. To see why, consider two data-generating processes that imply mean annual expenses of \$12,000. In the first process, the individual realizes either \$0 or \$2000 expenses in the course of a month, with the probability of each amounts independently determined across months. In the second process, the individual either realizes \$0 in every month or \$2000 in every month, with each of these two outcomes occurring with probability one half. Although the mean annual expense is \$12,000 in each case, the standard deviations of expenses are \$3464 and \$12000 in cases 1 and 2.

OOPs in the model are persistent on account of the presence of distinct types that each consumer randomly draws at the outset of a year. Consumer type has two dimensions: each consumer is either an NZ (for "non-zero") type or a Z type, and either a High type or a Low type. The interaction of these two dimensions implies four distinct types. An NZ type never has non-zero expenses in any month of the year, whereas the Z type may (or may not) have zero expenses in a month. Next, the High type has OOPs that are stochastically greater than those of a similar Low type conditional on positive OOPs.

We assume that consumers do not know their type (NZ or Z, High or Low) at the time of making an insurance decision, and types are not correlated across years. Instead, the health score h_j represents the consumer's knowledge of their health state when choosing insurance, and its transition process captures persistence in expenses across years. The distinct types serve as a modelling device: they allow the OOP distribution to be represented as a flexible mixture of simple distributions, but we do not consider the types specification as a literal description of heterogeneity across consumers.

Under our parametric model of conditional OOP distributions, consumers who realize the Z type incur positive OOPs in the course of a given month with a conditional probability following a logit functional form that depends on the consumer's health risk and insurance status. The monthly expenses of those consumers who do incur positive expenses follow a log normal distribution whose location parameter depends on the consumer's health risk as well as whether the consumer is of the High or Low type. Also, the scale parameter of this log normal distribution depends on the consumer's health risk. We now formally describe the model. Under our parametric specification, consumer j's OOPs in month t under insurance status $a \in \{0, 1\}$, call them OOP_{jt}, are assumed to have the following distribution conditional on health risk h_j,

income y_j, and the consumer's type:

$$OOP_{jt}(a) = \begin{cases} 0, & \text{with probability } 1-\rho_j^a \\ x_{jt} \sim \log N\left(\mu_j^a, (\sigma_j^a)^2\right), & \text{with probability } \rho_j^a. \end{cases} \tag{14}$$

Thus, conditional on the ρ_j^a, μ_j^a, and σ_j^a variables that will soon be explained, $OOP_{jt}(a)$ has a Bernoulli part that determines whether expenses are zero or non-zero in addition to a log-normal (denoted "$\log N$" above) part that determines the magnitude of positive expenses. Each of ρ_j^a, μ_j^a, and σ_j^a depends on the individual's insurance status as well as their health risk and income. First, the probability of non-zero expenses ρ_j^a is determined as

$$\rho_j^a = \mathrm{NZ}_j + (1-\mathrm{NZ}_j) \times \frac{e^{\phi_0^a + \phi_h^a h_j + \phi_y^a \log(y_j)}}{1 + e^{\phi_0^a + \phi_h^a h_j + \phi_y^a \log(y_j)}}.$$

Here, NZ_j is an indicator for the consumer having type NZ in the year in question. Additionally, $(\phi_0^a, \phi_h^a, \phi_y^a)$ are model parameters that govern the probability with which a Z type experiences expenses of zero. The probability with which consumer j is type NZ in a year is

$$\Pr(\mathrm{NZ}_j = 1 \mid h_j, y_j, a) = \frac{e^{\alpha_0^a + \alpha_h^a h_j + \alpha_y^a \log(y_j)}}{1 + e^{\alpha_0^a + \alpha_h^a h_j + \alpha_y^a \log(y_j)}},$$

where $(\alpha_0^a, \alpha_h^a, \alpha_y^a)$ are model parameters.

We next detail the variable μ_j^a governing the location of the log-normal deviate x_{jt}. This variable is defined as

$$\mu_j^a = \gamma_0^a + \gamma_h^a h_j + \mathrm{Hi}_j \times (\zeta_0^a + \zeta_h^a h_j),$$

where $(\gamma_0^a, \gamma_h^a, \zeta_0^a, \zeta_h^a)$ are model parameters and Hi_j is an indicator for whether consumer j is a High type. The conditional probability with which consumer j is a High type in the current year is

$$\Pr(\mathrm{Hi}_j = 1 \mid h_j, y_j, a) = \frac{e^{\beta_0^a + \beta_h^a h_j + \beta_y^a \log(y_j)}}{1 + e^{\beta_0^a + \beta_h^a h_j + \beta_y^a \log(y_j)}},$$

where $(\beta_0^a, \beta_h^a, \beta_y^a)$ are model parameters.

The last unexplained variable appearing in (14) is σ_j^a, which governs the scale of the log-normal part of the OOP distribution. We specify

$$\sigma_j^a = e^{\varsigma_0^a + \varsigma_h^a h_j},$$

where $(\varsigma_0^a, \varsigma_h^a)$ are model parameters. Note that the scale parameter σ_j^a does not depend on income y_j. When specifying the conditional OOP model, we face a trade-off between flexibility and statistical precision. Given the precision costs of a richer model, we do not allow the scale parameter to depend on income.

To complete our specification of the conditional OOP distribution model, we assume that that,

conditional on h_j, y_j, and a, NZ_j and Hi_j are independent of each other. We also assume that, conditional on the consumer's type, health risk, income, and insurance status, the log-normal deviate x_{jt} is independent of whether the consumer has non-zero OOPs in a month t. Also, x_{jt} and whether the consumer has non-zero OOPs in month t are each independent across months t.

We estimate the model by maximum likelihood, weighting each observation by its associated panelist's analysis weight as provided by MEPS. Each observation is a panelist/month pair, and the months included are those for which the panelist in question belonged to round 5 of MEPS. We drop the first of these months, as MEPS rounds often begin in the middle of the month, meaning that OOPs recorded for the first month of round 5 may not fully reflect expenses incurred in that month.

Appendix D.1 describes our method for conducting inference on the estimated parameters. Table 11 reports the parameter estimates and their standard errors.

Table 11: Parameter estimates for OOP model

Parameter	Insured		Uninsured	
	Estimate	SE	Estimate	SE
α_0^a	-1.25	(0.68)	-1.20	(0.17)
α_h^a	0.46	(0.04)	0.66	(0.02)
α_y^a	0.04	(0.06)	0.06	(0.02)
β_0^a	-2.05	(2.34)	-0.24	(0.33)
β_h^a	0.02	(0.06)	-0.04	(0.06)
β_y^a	0.25	(0.22)	0.05	(0.03)
ϕ_0^a	-4.61	(1.24)	-5.15	(0.67)
ϕ_h^a	0.23	(0.04)	0.38	(0.04)
ϕ_y^a	0.15	(0.12)	0.18	(0.07)
γ_0^a	1.09	(0.19)	1.70	(0.10)
γ_h^a	0.13	(0.05)	0.24	(0.06)
ζ_0^a	2.67	(0.16)	2.53	(0.06)
ζ_h^a	-0.04	(0.04)	-0.04	(0.04)
ς_0^a	0.27	(0.05)	0.26	(0.02)
ς_h^a	-0.01	(0.01)	-0.03	(0.01)
$N_{\text{panelists}}$	1225		8759	
$N_{\text{panelist/months}}$	4922		35056	

Notes: this table provides estimates of the parameters of the conditional OOP distribution model. The table provides standard errors in parentheses. The $N_{\text{panelists}}$ row provides the number of unique panelists in the estimation subsamples whereas the $N_{\text{panelist/months}}$ rows provides the number of panelist/month observations in the estimation subsamples. These sample sizes are provided separately for the insured subsample and the uninsured subsample.

To draw from a consumer's annual OOP distribution, we first draw the consumer's types (Z or NZ; Low or High). Next, we take 12 draws the estimated monthly OOP distributions conditional on these types. We then sum across these 12 draws to obtain our draw of the consumer's annual OOPs.

D.1 Inference for OOP distribution model

We now describe our procedure for conducting inference on the parameters of the model of OOP conditional distributions. Let n denote the size of the estimation sample; let w_j denote the MEPS individual weight for observation j; let z_j denote a vector of observation j's insurance status, health status, and OOPs; let $\log \ell$ denote the conditional log likelihood contribution function; let $\nabla \log \ell$ denote the gradient of ℓ with respect to the model parameters θ; and let $\nabla^2 \log \ell$ denote the Hessian of ℓ with respect to θ. Then, the asymptotic variance of $\sqrt{n}(\hat{\theta} - \theta_0)$, where $\hat{\theta}$ is the maximum likelihood estimator and θ_0 is the true parameter value, is

$$\Sigma = H^{-1} \Omega H^{-1},$$

where H is the probability limit of

$$\hat{H} = \frac{1}{n} \sum_{j=1}^{n} w_j \nabla^2 \log \ell(z_j, \hat{\theta})$$

and Ω is the probability limit of

$$\hat{\Omega} = \frac{1}{n} \sum_{j=1}^{n} \Big(w_j \nabla \log \ell(z_j, \hat{\theta}) \Big) \times \Big(w_j \nabla \log \ell(z_j, \hat{\theta}) \Big)'.$$

The estimate of Σ that we use in conducting inference on the model parameters is

$$\hat{\Sigma} = \hat{H}^{-1} \hat{\Omega} \hat{H}^{-1}.$$

E Welfare Analysis of Republican proposals

E.1 Background

Republican legislators have proposed a type of uninsurance penalty that is different from the ACA-style individual mandates. Specifically, they proposed to penalize individuals who *game the system* by remaining uninsured when healthy and seeking insurance when sick. One possible justification for such a proposal would be the idea that an ACA-style mandate restricts consumers' individual freedoms. In contrast to the ACA individual mandate, which penalizes consumers while uninsured, the Republican-proposed policies would penalize consumers upon return to exchanges. Such policies encourage exchange participation without punishing consumers who never purchase insurance.

Policymakers have advocated for two variations of such dynamic penalty schemes. One variation requires any individual who has not purchased insurance in the previous year to pay a surcharge $\kappa_{jt} \geq 0$ upon re-enrolling in exchange insurance. This surcharge may depend on the premium p_{jt} faced by the consumer; in fact, the dynamic penalty proposed by Republican lawmakers in the House of Representatives entailed a surcharge equal to 30% of the premium of the consumer's

insurance plan. The second sort of dynamic penalty prohibits consumers who have lapsed in insurance coverage from purchasing insurance for a fraction $\eta \in [0, 1]$ of the year t in which they seek to re-enroll in exchange insurance. If a consumer chooses to buy insurance in year t, they will have to pay a prorated premium $\eta \times p_{jt}$. The OOPs of a consumer who holds insurance for a fraction η of the year equal

$$(1 - \eta)OOP_{jt}(0, m_{jt}) + \eta OOP_{jt}(1, m_{jt}).$$

This section studies the welfare effects of such proposals, especially as they compare against ACA-style penalties. We first describe how such policies are formally captured, and then turn to the results.

E.2 Modeling of Penalties and Optimal Insurance Decisions by Consumers

The republican-proposed inducements discussed above are both dynamic in nature: they penalize a consumer's uninsurance in year $t-1$ during year t. As a result, these policies render the consumer's decision problem also dynamic. This problem features two sorts of state variables: first, consumer characteristics X_{jt} including health status, age, and, income, and, second, a penalty status L_{jt} that indicates whether the consumer is subject to dynamic penalties. Formally, $L_{jt} = 1 - a_{j,t-1}$. Note that the nature of the penalty depends on the prevailing regulatory regime.

We now turn to the value function $V(X_{jt}, L_{jt})$. At the beginning year t, the value function for a 64 year old consumer is defined as

$$V(X_{jt}, L_{jt}) = \max_{a_{jt}} \mathbb{E}U_{jt}\Big(c_{jt}(a_{jt}|X_{jt}, L_{jt})\Big). \tag{15}$$

For consumers younger than 64 years of age, the value function is defined as

$$V(X_{jt}, L_{jt}) = \max_{a_{jt}} \left[\mathbb{E}U_{jt}\Big(c_{jt}(a_{jt}|X_{jt}, L_{jt})\Big) + \beta \mathbb{E}V(X_{j,t+1}, 1 - a_{jt}|X_{jt})\right] \tag{16}$$

where β is the consumer's discount factor. These last two equations deliver consumer jt optimal action a^*_{jt} given state (X_{jt}, L_{jt}). The vector $\mathbf{a} = \{a_{jt}\}_{jt}$ denotes all of the insurance decisions made by consumers. Note that vector of optimal decisions $\mathbf{a}^* = \{a^*_{jt}\}_{jt}$ depends on the vector of premiums (including premiums paid later in life) charged by insurers.

E.3 Results and Interpretation

Table 12 presents the results, comparing the performances of ACA-style penalties to these dynamic policies. The comparisons are made under the baseline specificaiton, namely, with subsidies and uncomopensated care.

The key observation from the table is that which type of penalty is preferred depends on the

welfare metric used. Under a myopic measure, dynamic penalties are preferred. This comparison flips under the paternalistic welfare metric.

Dynamic penalties outperform ACA style ones under myopic regimes. This is because dynamic penalties "respect" the heterogeneity in individual-level taste for insurance z_{jt}. These taste parameters, described by equation 4, vary significantly across individuals but minimally within individuals over time. As a result, remaining uninsured due to a low z_{jt} is likely to remain a consistent behavior over time and, hence, unpunished by dynamic penalties which only penalize a behavior change. ACA-style penalties, however, punish for uninsurance regardless of circumstances. As a result, under ACA mandates, individuals with low z_{jt} either have to pay a penalty or inefficiently sign up for insurance. This result is consistent with the interpretation that dynamic penalties avoid posing risks on consumers' individual freedoms.

The welfare-comparison between the two penalty types flips under the paternalistic measure: ACA-style penalties now do better. This is because the impact of myopia on insurance decisions, similar to that of distastes, is expected to remain relatively constant within each individual over time. As a result, myopia is less likely to lead to inconsistent insurance decisions over time, rendering it unpunished by dynamic penalties. ACA mandates, however, punish uninsurance unconditionally. Thus, they do penalize uninsurance due to myopia. As a result, a welfare metric that considers myopia "a mistake to correct", would favor ACA penalties over dynamic ones.

Table 12: Comparison of ACA inducements and dynamic penalties

Penalty regime	Mandate penalties			Dynamic penalties	
	Zero	2016 ACA	Extreme ACA	Surcharge (30%)	$\frac{1}{2}$ year OOM
Part. (%)	38.4	51.8	100	43.9	54.4
Social deficit	**2104**	**1521**	**1296**	**1910**	**1811**
Myopic welfare	0	-437	-683	-146	-178
Patern. welfare	**0**	**-110**	**101**	**-331**	**-481**
Agg. welfare	**-2104**	**-1631**	**-1195**	**-2241**	**-2292**

Notes: see the notes for Table 7.