Health Insurance Underwriting and the Heterogeneous Effects of the Affordable Care Act *

Pietro Tebaldi§ Andrea Scalenghe¶ Joe Orsini‡

October 2026

Abstract. The Affordable Care Act, the largest US welfare reform since 1965, banned health insurance underwriting and rejections based on pre-existing conditions and introduced large premium subsidies for low- to middle-income individuals. Yet how many people the ban made better off, and how its effects varied by health risk, remain unknown—because data on who was rejected before the reform have been lacking. Using a proprietary dataset of pre-ACA insurance applications and underwriting decisions, we document, model, and estimate the reform's heterogeneous effects by age, income, and health risk. We find its benefits were concentrated among low-income, high-risk, or over-50 individuals.

Keywords. ACA, Insurance Reform, Insurance Rejections, Underwriting
JEL Codes. I11, I13

*We thank Martin Hackmann, Igal Hendel, Liran Einav, and Kosali Simon for the helpful comments. Miguel Borrero provided excellent research assistance. All errors are our own.
§Columbia University and NBER. Email: p.tebaldi@columbia.edu.
¶MIT. Email: ascale@mit.edu.
‡Upstart Network, Inc. Email: joe.orsini@upstart.com.

1 Introduction

The Patient Protection and Affordable Care Act (ACA) was enacted in 2010 and has since reshaped the structure of the US health insurance market for individuals without access to either public or employer-sponsored coverage. The backbone of the reform consisted of prohibiting medical underwriting and risk-based pricing—forcing insurers to sell coverage to all individuals and imposing strict rating regulations—and the introduction of a large, means-tested subsidy program to safeguard the affordability of insurance premiums. As the largest US welfare reform since the introduction of Medicare and Medicaid in 1965, the ACA continues to play a central role in policy discussions, motivating a rich research agenda at the intersection of health economics and industrial organization (see, e.g., Handel and Kolstad, 2022, for a recent review). While this body of work is broad, retrospective assessments of the ACA's impact have been limited by the scarcity and scattered nature of information on the pre-reform market landscape.

In this article, we address this gap by complementing publicly available health insurance data with a proprietary dataset of pre-ACA health insurance applications and underwriting decisions from 2008 to 2012, previously used by Orsini (2015) and Wong (2015). These data contains two critical observables: the premium quoted for the plan selected by each applicant, and the insurer's decision to approve or reject the application after medical underwriting. We use these data in a simple empirical model of insurance demand and (pre-reform) insurance rejections. The model delivers the joint distribution of demographics, health risk, pre- and post-ACA demand, and pre-ACA rejection probability. We use these estimates to measure the ACA's effect across different groups, and quantify the extent to which the reform's two main components—the underwriting ban and premium subsidies—played complementary roles in shaping coverage and welfare.

In Section 2, we begin by illustrating how the data speak directly about the underwriting ban, rating restrictions, and subsidies introduced by the ACA. We document that almost 50 percent of applications were rejected,[1] and that average annual premiums increased from \$2,000 to \$5,700 following the ACA, consistent with an increase in adverse

[1]This share is application-weighted, and reflects the fact that some individuals applied multiple times.

selection brought about by rating and underwriting restrictions (Handel, Hendel and Whinston, 2015). Subsidies absorbed the vast majority of this increase for the lowest-income and oldest individuals, and, for many, net-of-subsidy, effective premium became lower than pre-ACA levels. Adding further evidence of adverse selection before and after the reform, survey data on health risk show that there is a positive relationship between insurance and risk, and that this relationship strengthened after the ACA.

In Section 3 we introduce a simple econometric model of insurance demand and underwriting that allows us to link insurance shares and rejections to the distribution of risk and demographics. We model individuals as making a binary decision of whether to seek insurance coverage or remain uninsured, with preferences varying by income, age, health risk, year, and geographic location. We further assume that, in the pre-ACA period, insurers used a cutoff rule to accept or reject applications based on the applicant's predicted risk, and that rejected individuals may re-apply a random number of times. The model is identified by observed insurance shares, rejection rates, the average number of applications per applicant,[2] and "micro-moments" that describe risk conditional on insurance status and demographics, built from the Medical Expenditure Panel Survey.

In Section 4 we use our estimates to show how insurance demand varies with risk, how this gradient changed between the pre- and post-ACA periods, and how premium subsidies affected it. Prior to the ACA, individuals in the top 20 percent of the risk distribution were more than four times as likely to seek insurance coverage than those in the bottom 20 percent. However, following underwriting and rejections—which, we find, affected approximately 30 percent of the risk distribution, consistent with Claxton, Cox, Damico, Levitt and Pollitz (2016)—the probability of insurance was increasing in risk only for those in the first six deciles of the risk distribution, and decreasing in risk for those in the top four deciles. Those with the highest risk were denied coverage with probability one. Following the ACA, the average probability of insurance is greater for nearly all risk levels. If, however, we consider a partial equilibrium counterfactual in which premium subsidies are removed, the probability of insurance would be, on average, lower for individuals with risk between the second and eighth deciles.

[2]Our data provider informed us that, on average, each individual submitted 1.5 applications.

We find that the ACA increased average consumer surplus by $100 per person per year, but this effect is highly heterogeneous. For those who would have been rejected prior to the ACA, consumer surplus increased much more, and we do not detect any meaningful negative effect of the reform. Instead, the ACA reduced consumer surplus for one fourth of the healthier individuals whose coverage would have been approved. The shares of "winners" varied substantially by age and income, with the largest gains concentrated among subsidized individuals under 40, or over 50. We estimate that less than half of those middle-aged, higher-income, and low-risk benefited from the reform.

Importantly, we measure the effects of the ACA conditional on demographics and realized health risk. As discussed thoroughly by Handel *et al.* (2015), what we measure is an ex-post realized consumer-surplus effect, which reflects changes in selection that are natural when underwriting and risk-based pricing are banned. The premium subsidies introduced by the ACA mitigated some of these effects for the lowest-income. Although we find that for a nonnegligible share of individuals the ACA reduced consumer surplus, our measure does not account for the reduction in reclassification risk. Behind the veil of ignorance, individuals do not know their future health or income, so an ex-ante evaluation could differ from our ex-post realized consumer-surplus estimates.

We make two contributions to the literature.

First, we provide direct empirical evidence on medical underwriting before the ACA. Hendren (2013) and Handel, Hendel and Whinston (2015) emphasized the trade-off between risk-based exclusion and the welfare losses generated by adverse selection (Einav, Finkelstein and Cullen, 2010). Prior empirical work examines pre-ACA individual markets through state community-rating and guaranteed-issue regulations (Pauly and Herring, 2007; Herring and Pauly, 2006; Sasso and Lurie, 2009), the Massachusetts reform (Hackmann, Kolstad and Kowalski, 2012, 2015; Finkelstein, Hendren and Shepard, 2019; Shepard and Forsgren, 2023), and the prevalence of health conditions that could trigger rejection (Claxton *et al.*, 2016). We extend this work by combining observed underwriting decisions and quoted premiums with data on coverage and health risk. This allows us to distinguish willingness to purchase insurance from the ability to obtain it and to connect pre-reform rejection patterns to post-reform outcomes.

Second, we measure how the ACA's effects on coverage and realized consumer surplus vary with health risk, age, and income. Our analysis complements studies documenting the ACA's aggregate coverage effects and the central role of subsidies (Frean, Gruber and Sommers, 2017; Aboulafia, Gruber and Sommers, 2025), as well as work on post-reform insurance demand and adverse selection (Panhans, 2019; Saltzman, 2019; Tebaldi, 2025; Ghili, Hendel, Sullivan and Whinston, 2026). By jointly quantifying the roles of the underwriting ban and premium subsidies, we provide a retrospective assessment of the reform that links pre-ACA exclusion to post-ACA coverage and welfare, informing the broader policy-design agenda discussed by Baicker, Chandra and Shepard (2023).

2 Data and Institutional Background

Pre-ACA Applications and Underwriting The main data source that enables our analysis is a dataset of insurance applications and underwriting decisions (AU, henceforth), obtained under a data-use agreement with a national online health insurance marketplace. This dataset was previously used by Orsini (2015) and Wong (2015).

The data span the 2008-2012 period, during which the marketplace worked as follows. First, after providing age, gender, and zip code, individuals were presented with a menu of contracts. Upon selecting a plan, the underwriting phase began. This required the disclosure of extensive medical history and authorization for the insurer to verify self-reported data against external records. The insurer could then approve or reject the application, or communicate with the applicant about contract adjustments or alternative coverage options. We observe nearly 1.5 million applications, each paired with the insurer's decision on whether to approve coverage, the applicant's age, county of residence, and the premium and deductible for the selected contract.

While we observe whether coverage was approved or denied, underwriting was the result of a highly heterogeneous process. Insurers followed different procedures to collect information on pre-existing medical conditions and prescription drug use, and states imposed different restrictions on these processes (see, e.g., Claxton *et al.*, 2016). Within the online marketplace, premium adjustments following underwriting were limited, or prohibited (Pauly and Herring, 2007). Asymmetric information provides a justification

for rejections to be the preferred response to high-risk applicants (Hendren, 2013).

We use the AU data to measure three market outcomes: premiums, contract characteristics, and the probability of rejection following underwriting. For our analysis, we aggregate the data to the county, year, and age-bin level (18–39, 40–49, and 50–64); we denote this combination with the subscript *cta* henceforth. Moreover, we exclude the states that even prior to the ACA adopted community rating and guaranteed-issue regulation (MA, ME, NH, NJ, NY, VT; cf. Pauly and Herring, 2007). We compute the (application-weighted) average premium P_{cta}, the probability that the application is denied after underwriting R_{cta}, and a measure of actuarial value AV_{cta}.[3]

The AU data lack individual identifiers, making it impossible to observe a sequence of applications by the same individual. Nevertheless, we have learned from the data provider that, on average, the data contain 1.5 applications per individual. This provides us with an important empirical moment to avoid the risk of double-counting individuals who were unable to obtain insurance prior to the ACA.

Figure 1 and Table 1 summarize P_{cta} and R_{cta} in the pre-ACA period. Average annual premium was \$1,970, varying from around \$1,200 among individuals aged 18–39 to over \$3,000 among those aged 50–64. Nearly five out of ten applications were rejected. This share was lower among younger individuals and increased to nearly 7 out of 10 among those aged 50-64. Average actuarial value was 56.3 percent, highlighting that several coverage options were below the minimum mandated by the ACA since 2014.

Insurance Coverage and Post-ACA Data We complement the AU data with public administrative data on insurance coverage for the entire 2008-2019 period and on 2014-2019 post-ACA plan offerings.[4] Although the coverage data begin in 2008, SAHIE reports ages 40–64 as a single group in 2008 and 2009. We therefore construct the three-age-group insurance shares used in estimation beginning in 2010; income-specific shares are

[3]We compute the actuarial value of the plans in the AU data following a procedure similar to the one used to classify plans in "metal tiers" under the ACA. C.f. Department of Health and Human Services (2013), see e.g. `www.cms.gov/cciio/resources/files/downloads/av-calculator-methodology.pdf`; last accessed on 02/18/2026. See Supplementary Appendix for more details.

[4]We focus on a number of years after the reform that is similar to the length of our pre-ACA data, limiting confounding effects from changes enacted by subsequent legislation. Following 2019, the repeal of the individual mandate constituted a major change in the ACA, which we do not consider in our analysis.

available throughout 2008–2019. For a given county, year, age-bin, and income-bin y (138–249, 250–399, 400+, measured as a percentage of the FPL), we measure the insurance share in the individual market as

$$Q_{ctay} = \frac{I^{\text{SAHIE}}_{ctay} \times \iota^{\text{ACS}}_{ctay}}{I^{\text{SAHIE}}_{ctay} \times \iota^{\text{ACS}}_{ctay} + U^{\text{SAHIE}}_{ctay}} \,. \tag{1}$$

In this expression, I^{SAHIE}_{ctay} and U^{SAHIE}_{ctay} represent, respectively, the number of individuals with and without health insurance as measured by the Census Bureau Small Area Health Insurance Estimates Program.[5] The term $\iota^{\text{ACS}}_{ctay}$ is the share of individual-market insurance among the insured individuals in the 1-year sample of the American Community Survey; for more details, see Supplementary Appendix A.

Figure 1 and Table 1 illustrate how the ACA affected insurance coverage in the segment of the individual market not impacted by the Medicaid expansion. The share of covered individuals increased from 0.38 to 0.5, a 20 percent reduction in uninsurance as measured in Frean *et al.* (2017); Aboulafia *et al.* (2025). Changes in coverage were different across income groups. Among individuals who benefited from the most generous premium subsidies and cost-sharing reductions (y in [138–249]), coverage increased by 50 percent from a pre-ACA average of 0.2. Among those who benefited from (less generous) premium subsidies, but did not receive cost-sharing reductions (y in [250–399]), coverage increased by 29 percent from a pre-ACA average of 0.35. Among the unsubsidized, coverage increased by 13 percent from 0.6.

Since our analysis focuses exclusively on the extensive margin decision of whether to purchase insurance, we measure premium for every county-year-age-income group, P_{ctay}. This is built as the enrollment-weighted average of premiums after accounting for age-rating and subsidies, as described in detail in Supplementary Appendix A.

Figure 1 and Table 1 highlight the effect of the ACA on insurance premiums across different groups. Without subsidies, every age group experienced significant premium increases, consistent with the additional adverse selection predicted under rating restrictions and underwriting bans (Handel *et al.*, 2015). Average (pre-subsidy) premium increased

[5] See: `www.census.gov/programs-surveys/sahie.html`; last accessed on 02/18/2026.

from $2,000 to $5,700, nearly tripling overall. After accounting for subsidies, premiums were reduced. For the youngest group, this reduction was relatively small, while for older groups—recipients of much more generous subsidies—annual premiums dropped from $3,000 to $1,000, on average.

These changes in premiums partly reflect differences in coverage quality. The changes brought about by the ACA along this dimension were significant. Focusing on actuarial value, this is now highly standardized within metal tiers (bronze, silver, gold, and platinum), and within silver it varies with income according to the cost-sharing reductions (CSR, see Lavetti, DeLeire and Ziebarth, 2023).[6,7] Table 1 shows that, even without CSR, the average actuarial value increased from 56.3 percent to 68.8 percent. Accounting for CSR, it increased further to 72.4 percent, a 16 percentage-point gain from the pre-ACA.

Risk Scores and Risk Selection. We build a simple risk score measure using the MEPS. For this, we combine the CMS ACA risk adjustment model with an approach similar to that in Ghili *et al.* (2026).[8] The resulting distribution of risk score conditional on age and income, $F(r|a, y)$, is an important input for our analysis. The bottom panels of Figure 1 plot $F(r|a, y)$ across the nine age-income groups, and the distribution of expenditure conditional on risk, $F(e|r)$, for three ranges of r (bottom quartile, with mean expected annual spending lower than $500; second and third quartiles, $3,000; and top quartile, $11,000).

This risk model allows us to take a first look at whether the ACA's ban on underwriting

[6]In 2017, the federal government ended CSR reimbursements to insurers, while insurers remained required to provide the statutory cost-sharing reductions to eligible enrollees. This led to "Silver loading"—a practice by which insurers increased the base rate of silver plans to compensate for the loss of CSR payments.

[7]We set AV^j_{ct} for a given plan j equal to 60 for bronze plans, 70 for silver, 80 for gold, and 90 for platinum plans. For Silver plans and incomes between 138 and 249% FPL, we set the CSR-inclusive value to the equal-weight mean of the 74%, 88%, and 95% variants, or 85.67%, and obtain AV^j_{cty}. We then take the enrollment-weighted average over j's within each county-year, and then compute AV_{cty} within our nine age-income cells.

[8]We begin by computing r^{ACA}_{it} as the most precise version of the ACA risk adjustment factor for each individual in the MEPS at the end of year t. This is a function of age and gender in year t and diagnoses measured in year $t-1$. Then, we fit a lasso regression of log annual health expenditure in year t on a high-order polynomial of r^{ACA}_{it}, age, and gender, together with additional health events measured in year $t-1$. For each individual, we set the risk score r_i equal to the corresponding percentile of the distribution of the fitted values of this model. The Supplementary Appendix provides more details, and shows that our results are robust to this construction.

and rating restrictions led to stronger risk selection relative to the pre-reform period. The bottom panel of Table 1 shows the average risk score conditional on insurance status, age, and income group in the MEPS. This highlights that, while in both periods the insured population was riskier than the uninsured, the difference increased after the ACA for 8 out of 9 age-income groups, and nearly doubled among the subsidized. This is consistent with an increase in adverse selection following the reform.

3 Econometric Model

An individual i is characterized by their county, age, and income, collected in $x_i = (c_i, a_i, y_i)$, and by a risk score r_i, drawn independently from the conditional distribution $F(r \mid x_i) = F(r \mid a_i, y_i)$. In every year t, let $N_t(x)$ denote the number of individuals with characteristics x and let $f(r|x)$ denote the conditional probability mass function of risk. Then $N_t(x, r) = N_t(x)f(r|x)$ denotes the population mass with characteristics and risk (x, r), while $F(r|x)$ denotes the corresponding conditional cumulative distribution.

3.1 Demand

We model the indirect utility from insurance as a function of net-of-subsidy premium, risk, and an unobservable term varying across counties, years, and demographics. Formally, individual i is willing to purchase insurance in year t if $V_{it} \geq 0$, where

$$V_{it} \equiv \delta_t(x_i) + \rho_t \times \Delta r_i + \epsilon_{it}^I - \epsilon_{it}^U \,, \tag{2}$$

$$\delta_t(x_i) \equiv \alpha(y_i) \times P_t(x_i) + \xi_t(x_i) \,. \tag{3}$$

The terms ϵ_{it}^I and ϵ_{it}^U in equation (2) are shocks to the utility of being insured and uninsured, respectively. These are independent draws from a type-1 extreme value distribution, implying that the probability that individual i is willing to purchase insurance is

$$\sigma_{it} \equiv \frac{\exp\left[\delta_t(x_i) + \rho_t \times \Delta r_i\right]}{1 + \exp\left[\delta_t(x_i) + \rho_t \times \Delta r_i\right]}. \tag{4}$$

In the above expressions, $\Delta r_i \equiv r_i - \mathbb{E}_F\left[r \middle| a_i, y_i\right]$ denotes the deviation in i's risk from the mean of same-age and same-income individuals.

The parameter ρ_t captures the effect of risk on insurance demand, and is allowed to vary across periods: $\rho_t = \rho^{\text{pre}}$ for $t < 2014$ and $\rho_t = \rho^{\text{post}}$ for $t \geq 2014$. The ACA impacted the generosity of insurance, market structure, marketing strategies, and the overall shape of the individual insurance market, and these changes are likely to affect the attractiveness of insurance differentially for individuals with different risk.

The term $\delta_t(x_i)$ in equations (2) and (4) measures mean utility from insurance conditional on year, county, age, and income. In equation (3) this term is decomposed into a linear function of net-of-subsidy premium, with slope $\alpha(y_i)$ that varies with income, and a county-year-age-income-specific term $\xi_t(x_i)$ which we will hold constant throughout. This decomposition allows us to study how demand would have changed under different premiums, and to evaluate the consumer surplus effect of the ACA.

3.2 Rejections and Repeated Applications

While the above demand model is standard, it is inadequate for instances where insurance applications could be rejected, and individuals could apply multiple times.

To this end, our key assumption is that insurers used a cutoff strategy to approve or reject applications based on the applicant's risk score. Formally, we assume that:

$$\text{an applicant } i \text{ is approved if } r_i \leq \psi_t(s_i, a_i), \text{ and rejected otherwise.}$$

The cutoff $\psi_t(s, a)$ is a state-year-age-specific object that we estimate for the purpose of measuring the ACA's effects accounting for pre-reform rejections. We model these cutoffs as constant within a state-year, conditional on age, due to data limitations. In the AU dataset, we do not observe income, and, after conditioning on age and year, county-level applications are often too few to estimate rejection rates precisely.

A second assumption we make is that, once an applicant i is rejected, they reapply a

random number of times. We assume that

a rejected applicant re-applies, on average, λ times after the initial rejection.

Combined with the cutoff assumption and our model of demand, this allows us to build a model prediction of the observed rejection rates and average number of applications per applicant as a function of the parameters that we collect in $\boldsymbol{\theta} = (\boldsymbol{\rho}, \boldsymbol{\delta}, \boldsymbol{\psi}, \lambda)$.

3.3 Estimation

We estimate $\boldsymbol{\theta}$ by matching observed insurance shares by age-income-county-year, rejection rates by age-state-year, the average number of applications per applicant in the AU dataset, and minimizing the distance between model-predicted and observed average risk of insured and uninsured by age, income, and period. Formally, we solve

$$\min_{\boldsymbol{\theta}} \sum_{a,y} \sum_{\tau} \sum_{\text{INS}} \left(\mathbb{E}\left[r \,\middle|\, \text{INS}, t \in \tau, a, y; \boldsymbol{\theta} \right] - \widehat{\mathbb{E}}\left[r \,\middle|\, \text{INS}, t \in \tau, a, y \right] \right)^2 \tag{5}$$

s.t.

$$Q_{ctay} = \int \sigma_{it} \mathbf{1}\left[r \leq \psi_t(s,a) \right] dF(r|x_i = (c,a,y)) \qquad \text{for } t < 2014, \text{ all } c, a, y$$

$$Q_{ctay} = \int \sigma_{it}\, dF(r|x_i = (c,a,y)) \qquad \text{for } t \geq 2014, \text{ all } c, a, y$$

$$\frac{\text{Reject} \mid \text{state}, a, t}{\text{Apply} \mid \text{state}, a, t} = \frac{(\lambda+1) \int \sigma_{it} \mathbf{1}\left[r > \psi_t(s,a) \right] dN_t(x, r|\text{state}, a)}{\int \sigma_{it} \left(1 + \lambda \mathbf{1}\left[r > \psi_t(s,a) \right] \right) dN_t(x, r|\text{state}, a)} \qquad \text{for } t < 2014, \text{ all state}, a$$

$$1.5 = \frac{\text{Pre-ACA applications}}{\text{Pre-ACA applicants}} = \frac{\sum_{t<2014} \int \sigma_{it} \left(1 + \lambda \mathbf{1}\left[r > \psi_t(s,a) \right] \right) dN_t(x,r)}{\sum_{t<2014} \int \sigma_{it}\, dN_t(x,r)},$$

where $\tau \in \{\text{pre,post}\}$ indexes the pre- and post-ACA periods, and INS $\in$ {I,U} indexes insured and uninsured individuals. This minimizes the mean-squared error between observed risk conditional on insurance, age, income, and period as shown in Table 1B, which we denote with $\widehat{\mathbb{E}}\left[r \,\middle|\, \text{INS}, t \in \tau, a, y \right]$, and model-predicted risk conditional on the

same characteristics, $\mathbb{E}\left[r \,\middle|\, \text{INS}, t \in \tau, a, y; \boldsymbol{\theta}\right]$. The constraints require insurance shares to be matched for every year, county, age, and income group, rejection rates to be matched for every pre-ACA year, state, and age group, and the average number of pre-ACA applications per applicant to be equal to 1.5 as reported by the AU data provider.

In practice, the post-ACA estimator is a standard "micro-BLP" (Petrin, 2002; Berry, Levinsohn and Pakes, 2004) using survey moments to target the risk composition of insured and uninsured individuals observed in the post-ACA MEPS, and estimating the mean utility terms $\delta_t(x_i)$ by matching insurance shares conditional on year and x_i. For the pre-ACA, we loop over values of ρ and, for each value, we find δ_t, ψ_t, and λ that match the observed insurance shares, rejection rates, and the average number of applications per applicant. These values are unique, as we discuss below.[9] We then select the ρ^{pre} that minimizes the objective in (5).

3.4 Identification

For the pre-ACA period, identification is constructive and closely tracks estimation. Fix candidate values of λ and ρ^{pre}. For any vector of rejection cutoffs, the contraction mapping of (Berry, Levinsohn and Pakes, 1995) yields a unique vector of mean utility parameters that matches the observed insurance shares. Substituting these mean utilities into the rejection-rate equations yields a system in the rejection cutoffs; because predicted rejection rates are strictly decreasing in the cutoffs, the observed rejection rates uniquely determine $\boldsymbol{\psi}$ and hence $\boldsymbol{\delta}$. Thus, each candidate pair $(\lambda, \rho^{\text{pre}})$ uniquely determines $(\boldsymbol{\delta}, \boldsymbol{\psi})$. For each candidate ρ^{pre}, the target of 1.5 applications per applicant uniquely determines $\lambda(\rho^{\text{pre}})$. Finally, substituting the resulting $\lambda(\rho^{\text{pre}})$, mean utilities, and rejection cutoffs into the risk moments yields an objective with a unique minimum in ρ^{pre}, since expected risk conditional on insurance is strictly increasing in ρ^{pre}. Hence, ρ^{pre}, λ, $\boldsymbol{\delta}$, and $\boldsymbol{\psi}$ are identified.

For the post-ACA period, identification of δ_t and ρ^{post} is standard, and it results from the combination of micromoments in Table 1B with observed shares (see Petrin, 2002;

[9] During this process, for a given λ it is possible that some age-county-year-income observations cannot be rationalized by the model and are excluded from the inversion step. These are cases in which, given λ, even setting the application probability to one, the observed rejection rate implies an insured share lower than the observed one.

Berry, Levinsohn and Pakes, 2004). The parameters in $\boldsymbol{\theta}$ are therefore identified. These are sufficient to decompose the impact of the ACA on insurance across groups.

To consider consumer surplus, we need the income-specific price coefficients $\alpha(y_i)$ in equation (3), which we obtain via 2SLS. We regress $\delta_t(x)$ on net-of-subsidy premium interacted with income group, and include county, year, age group, and income group fixed effects. We instrument for premium using the share of individuals younger than 50 among potential buyers in the county-year, interacted with income group indicators. This strategy follows Waldfogel (1999) and mimics Tebaldi (2025). Given the ACA's age-rating restrictions, insurers are incentivized to set lower base rates in markets with a larger share of relatively young buyers, and vice versa. This provides first-stage variation in premiums; the Supplementary Appendix reports the first-stage results. Identification additionally requires that, conditional on county, year, age, and income effects, the age composition of potential buyers affects insurance demand only through premiums. Because the regulatory source of first-stage variation is available only under the reform, we estimate $\alpha(y)$ using only post-ACA data.

4 Results

Figure 2 illustrates our findings on how the ACA impacted insurance demand and coverage across different risk groups, after accounting for changes in premiums, underwriting, and other characteristics of the market. Panel a shows that, prior to the reform, the share of individuals willing to pay the market premium to obtain insurance was steeply increasing in risk (red long-dashed line), from near zero for the healthiest individuals, to 0.5 for those in the middle of the risk distribution, to more than 0.8 for those in the top two risk deciles.

For low- and medium-risk individuals, the pre-ACA demand and coverage curves are nearly identical in panels a and b because their risk generally falls below the rejection cutoffs shown in panel c. This fact could be explained by high deductibles (thus the low actuarial value in Table 1), lack of individual mandate, and limited salience of individual insurance prior to the ACA relative to after the reform.

For individuals in the top four deciles of risk, Figure 2 tells a different story. Although their willingness to pay for coverage was the highest (panel a), their probability of being

insured dropped monotonically to zero (panel b). The probability of rejection, captured by the vertical distance between the red long-dashed curves in panels a and b, follows from the estimated cutoff distribution overlaid on population risk in panel c. Importantly, the model-implied probability of being rejected differs from the raw rejection rates we observed in the AU data, which suffer from the presence of repeated applications. We estimate that, after the initial rejection, a denied applicant applied, on average, two more times ($\lambda = 2$, see Supplementary Appendix).

Moving to the post-ACA period, the solid dark lines in panels (a) and (b) of Figure 2—which are identical because of the underwriting ban—reflect several simultaneous changes relative to the pre-ACA period: higher gross premiums, subsidies, the individual mandate (Frean *et al.*, 2017; Fiedler, 2020), and changes in plan characteristics, market access, and publicity (Aizawa and Kim, 2025). We find that, on average, coverage increased for nearly all individuals across the risk distribution, with the largest gains—of up to 80 percentage points—for those in the top two deciles, for whom (highly valued) coverage is no longer denied. Also for those in the bottom two deciles, coverage increased substantially.

The pre-ACA and ACA results in Figure 2 rely on our estimates of the rejection cutoffs (panel c), and on the parameters ρ^{pre} (9.44, se =1.02) and ρ^{post} (2.81, se =0.07) that discipline the risk-demand gradient, paired with the data presented in Section 2.

To separate the effect of premium subsidies from that of the underwriting ban, we go one step further and use our estimates of the parameters $\alpha(y)$ that discipline demand responses to premium.[10] The short-dash gray line in panels a and b of Figure 2 shows the model-implied average coverage by risk under the ACA if subsidies were removed, all else equal, by replacing net premium $P_t(x_i)$ with gross premium $P_t(x_i)+S_t(x_i)$, where $S_t(x_i)$ is the premium tax credit defined in Supplementary Appendix equation (A.5). This illustrates the importance of subsidies in increasing coverage in the face of the large premium increases that the ACA brought about (see also Aboulafia *et al.*, 2025). Panel a shows that, without subsidies, demand would be lower for 80 percent of the risk distribution.

[10]These estimates are illustrated in Table 2, with the bottom panel showing the implied semielasticities of demand for insurance with respect to premiums, which are decreasing in age, income, and risk. These align well with estimates in the literature, including Panhans (2019); Saltzman (2019); Finkelstein *et al.* (2019); Tebaldi, Torgovitsky and Yang (2023); Tebaldi (2025).

Even taking rejections into account (panel b), if individuals were required to pay the full premium (held fixed at observed 2014–2019 levels; a fully fleshed out equilibrium analysis is beyond our scope), average coverage would be lower under the ACA between the second and eighth deciles of risk, and only slightly higher for the bottom two deciles.

The large role played by subsidies in delivering coverage gains follows from the large premium increases from the pre- to post-ACA period (see also Figure 1 and Table 1). These increases are consistent with the adverse selection predicted in Handel *et al.* (2015). We illustrate this further in panel d of Figure 2, plotting the cumulative distribution of the premium index $100P^{\text{post}}/P^{\text{pre}}$, where an observation is a county-age-income cell. For the unsubsidized, the index lies between 200 and 400, implying that post-ACA premiums are two to four times their pre-ACA levels, or 100–300 percent higher, and the ratio is larger for market-demographic groups in the top decile of rejection probabilities. Instead, for the vast majority of subsidized individuals, premiums decreased, with small differences between markets with high or low pre-reform rejections.[11]

In Table 3 we illustrate richer heterogeneity and summarize our findings in terms of the ACA impact on coverage and consumer surplus by risk, age, and income. The curves in Figure 2 aggregate the period-specific levels across all estimation cells observed in each period. For the pre-post comparisons in Table 3, instead, we first restrict the sample to county-age-income cells $x = (c, a, y)$ observed in both periods. Using the potential-buyer-weighted period probability $q_\tau(x, r)$ defined in the note to Figure 2, we compute the change in coverage within each matched cell and risk percentile as

$$\begin{aligned} \Delta q(x,r) &= q_{\text{post}}(x,r) - q_{\text{pre}}(x,r), \qquad (6)\\ cs_t(x,r) &= \mathbf{1}\{r \le \psi_t(s(c),a)\} \ln\Big(1 + \exp\Big[\delta_t(x) + \rho_t \Delta r\Big]\Big) / |\alpha(y)|,\\ CS_\tau(x,r) &= \left(\sum_{t\in\tau} N_t(x)\right)^{-1} \sum_{t\in\tau} N_t(x)\, cs_t(x,r), \quad \Delta CS(x,r) = CS_{\text{post}}(x,r) - CS_{\text{pre}}(x,r). \end{aligned}$$

[11] Markets with higher pre-ACA rejections (and thus higher risk and post-ACA gross premiums) generally record lower average net-of-subsidy premiums. This is a mechanical consequence of the fact that subsidies target affordability of (the second-lowest silver) plan, and subsidized bronze plans become cheaper when silver plans' gross premiums are higher, as is the case in markets with higher risk.

Thus, $CS_\tau(x, r)$ is annual expected consumer surplus in the individual market, evaluated at realized risk and averaged across years within each period using potential-buyer weights.

In panel A of Table 3, we summarize the overall distribution of the increase in coverage, and also split this between individuals who would have been rejected pre-ACA and those who would have been approved. Overall, the probability of coverage increased for the top 75 percent of the individuals included in our data and for whom we can compute both pre- and post-ACA coverage probabilities. The mean increase of 15 percentage points masks a lot of heterogeneity, with the top 10 percent corresponding to increases of more than 60 percentage points, while the bottom 10 percent experienced a decrease in coverage of more than 20 percentage points. Conditional on being rejected prior to the reform, coverage increased throughout the entire population, with an average increase of nearly 70 percentage points. Instead, conditional on being approved, the average increase was much lower, and estimated at 8 percentage points.

Panel B of Table 3 translates these findings into annual per-person consumer surplus. Here we find that, on average, the ACA increased consumer surplus by $100 per person per year, with a more modest median increase of $50. For some, however, the increase was much higher, and equal to $250 at the 75th percentile, and $1000 at the 90th percentile. At the other end of the distribution, we find that the bottom 10 percent experienced a decrease in consumer surplus of more than $750.

Pre-ACA underwriting explains a portion of this heterogeneity, with positive gains in consumer surplus throughout the distribution for those who would have been rejected. These are as high as $5000 or more for the top 10 percent of the distribution. For the larger number of individuals who, if willing, could have purchased coverage prior to the reform, we find that 25 percent of them experienced a decrease in annual consumer surplus of $120 or more. Income and age also explain a portion of the heterogeneity in the effects of the ACA. In panel C of Table 3, we compute the share of individuals with a positive change in annual consumer surplus broken down across the nine age-income groups of our analysis, and distinguishing between those who could have purchased coverage prior to the reform, and others. Across all age groups, we find a share of "winners" ranging from 0.47 (for the highest income, middle age group who would have been approved) to 1

(for those rejected prior to the reform). Within each age group, higher income corresponds to a lower share of winners, as a natural consequence of premium subsidies. Among the unsubsidized, we find that those aged 50 or older benefited the most from the reform.

5 Conclusions

The ACA replaced a market that addressed adverse selection partly through medical underwriting with one that guaranteed access while regulating prices. Economic theory has long emphasized the resulting trade-off between risk-based exclusion and the costs of adverse selection. Yet empirical assessments of the ACA have rarely quantified how these competing forces affected different individuals, partly because pre-reform underwriting decisions are rarely observed. By linking such decisions to premiums, coverage, and health risk before and after the reform, we quantify the distributional consequences associated with both sides of this trade-off.

Our estimates reveal large gains for individuals who would previously have been rejected, alongside losses for some healthier individuals, especially those receiving little or no subsidy. Aggregate coverage gains therefore conceal meaningful redistribution across health-risk, age, and income groups. The results also highlight the role of premium subsidies in supporting a market with guaranteed issue. At observed post-ACA gross premiums, our counterfactual indicates that eliminating subsidies would substantially reduce coverage across much of the risk distribution. Affordability policies were therefore central to sustaining enrollment in the broader risk pool created by the underwriting ban.

Our consumer-surplus estimates are ex post and conditional on realized health risk and income. They do not capture the value of protection against reclassification risk, public financing costs, or equilibrium responses in premiums, plan design, and insurer participation. A more comprehensive welfare evaluation would need to incorporate these forces and individuals' uncertainty about future health and income. Our estimates provide an empirical benchmark that can inform such analyses.

References

Aboulafia, G., Gruber, J. and Sommers, B. (2025). The rise and fall (and rise) of the affordable care act: Varying impacts on coverage over time and place. *NBER Working Paper*, (w33615).

Aizawa, N. and Kim, Y. S. (2025). Public and private provision of information in market-based public programs: Evidence from advertising in health insurance marketplaces. *American Economic Journal: Economic Policy*, **17** (2), 84–121.

Baicker, K., Chandra, A. and Shepard, M. (2023). Achieving universal health insurance coverage in the united states: Addressing market failures or providing a social floor? *Journal of Economic Perspectives*, **37** (2), 99–122.

Berry, S., Levinsohn, J. and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, **63** (4), 841–890.

—, — and — (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of political Economy*, **112** (1), 68–105.

Claxton, G., Cox, C., Damico, A., Levitt, L. and Pollitz, K. (2016). Pre-existing conditions and medical underwriting in the individual insurance market prior to the aca. *KFF Issue Brief*, **2016**, 1–11.

Department of Health and Human Services (2013). Patient protection and affordable care act; standards related to essential health benefits, actuarial value, and accreditation. *Federal Register*, **78** (37), 12834–12872, 78 FR 12834, 45 CFR Parts 147, 155, and 156.

Einav, L., Finkelstein, A. and Cullen, M. R. (2010). Estimating Welfare in Insurance Markets using Variation in Prices. *The quarterly journal of economics*, **125** (3), 877–921.

Fiedler, M. (2020). The aca's individual mandate in retrospect: What did it do, and where do we go from here? a review of recent research on the insurance coverage effects of the affordable care act's individual mandate. *Health Affairs*, **39** (3), 429–435.

Finkelstein, A., Hendren, N. and Shepard, M. (2019). Subsidizing Health Insurance for Low-income Adults: Evidence from Massachusetts. *American Economic Review*, **109** (4), 1530–1567.

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

Ghili, S., Hendel, I., Sullivan, M. and Whinston, M. (2026). *What Prevents Health Insurance Take-up? Quantifying Frictions and Policy Responses*. Mimeo, MIT, Department of Economics, Boston, MA.

Hackmann, M. B., Kolstad, J. T. and Kowalski, A. E. (2012). Health reform, health insurance, and selection: Estimating selection into health insurance using the massachusetts health reform. *American Economic Review*, **102** (3), 498–501.

—, — and — (2015). Adverse selection and an individual mandate: When theory meets practice. *American Economic Review*, **105** (3), 1030–1066.

Handel, B., Hendel, I. and Whinston, M. D. (2015). Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica*, **83** (4), 1261–1313.

— and Kolstad, J. (2022). The affordable care act after a decade: industrial organization of the insurance exchanges. *Annual Review of Economics*, **14** (1), 287–312.

Hendren, N. (2013). Private information and insurance rejections. *Econometrica*, **81** (5), 1713–1762.

Herring, B. and Pauly, M. (2006). *The Effect of State Community Rating Regulations on Premiums and Coverage in the Individual Health Insurance Market*. Working Paper 12504, National Bureau of Economic Research.

Lavetti, K., DeLeire, T. and Ziebarth, N. R. (2023). How do Low-income Enrollees in the Affordable Care Act Marketplaces Respond to Cost-sharing? *Journal of Risk and Insurance*, **90** (1), 155–183.

Orsini, J. (2015). *Three Essays on the Economics of Health Insurance*. Stanford University.

Panhans, M. (2019). Adverse selection in aca exchange markets: Evidence from colorado. *American Economic Journal: Applied Economics*, **11** (2), 1–36.

Pauly, M. V. and Herring, B. (2007). Risk Pooling and Regulation: Policy and Reality in Today's Individual Health Insurance Market. *Health Affairs*, **26** (3), 770–779.

Petrin, A. (2002). Quantifying the Benefits of New Products: The Case of the Minivan. *Journal of Political Economy*, **110** (4), 705–729.

Saltzman, E. (2019). Demand for health insurance: Evidence from the california and washington aca exchanges. *Journal of Health Economics*, **63**, 197–222.

Sasso, A. T. L. and Lurie, I. Z. (2009). Community Rating and the Market for Private Non-Group Health Insurance. *Journal of public Economics*, **93** (1-2), 264–279.

Shepard, M. and Forsgren, E. (2023). Do insurers respond to active purchasing? evidence from the massachusetts health insurance exchange. *Journal of Risk and Insurance*, **90** (1), 9–31.

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

—, Torgovitsky, A. and Yang, H. (2023). Nonparametric Estimates of Demand in the California Health Insurance Exchange. *Econometrica*, **91** (1), 107–146.

Waldfogel, J. (1999). Preference externalities: An empirical study of who benefits whom in differentiated product markets.

Wong, P. E. (2015). *Essays in Health Economics and Industrial Organization*. Stanford University.

Figure 1: Main Analysis Variables Over Time and Across Groups

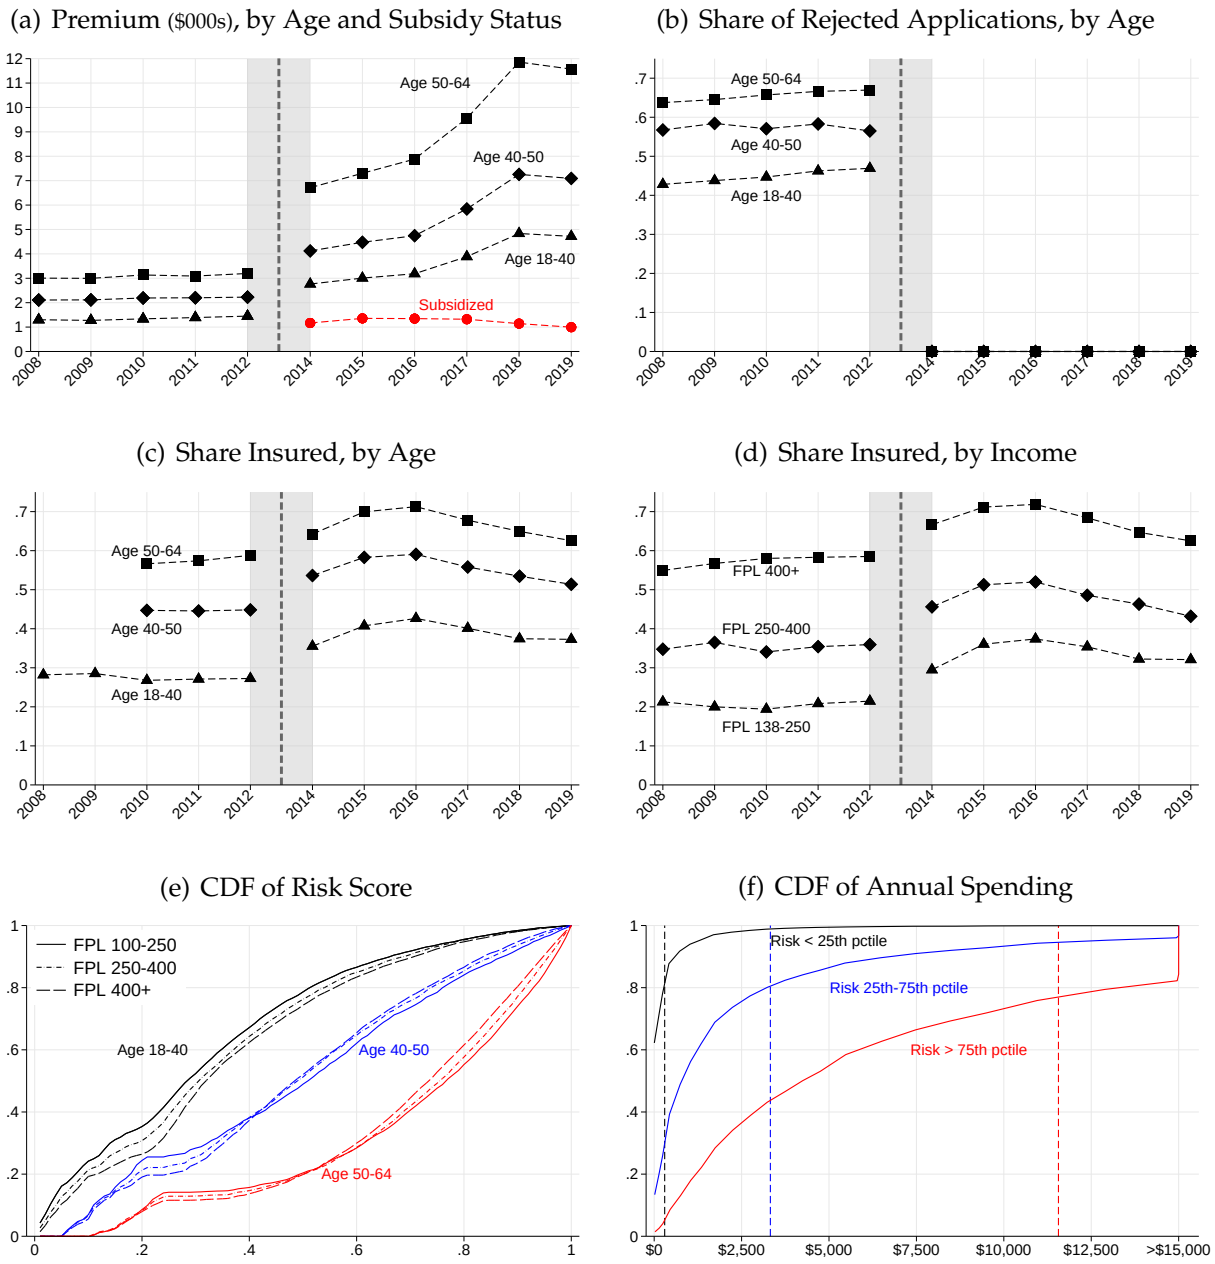

Note: Panel (a) plots the average annual premium by age group and subsidy status, weighted by the number of potential buyers in each county-year-age-income cell. Panel (b) plots the share of rejected applications by age group, weighted by the number of potential buyers in each county-year-age-income cell; these include repeated applications by the same applicant. Panel (c) and (d) plot the share insured by age and income group, respectively, weighted by the number of potential buyers in each county-year-age-income cell. Panel (e) and (f) plot the CDF of risk scores by age and income group and annual spending by risk quartile, respectively. Panel (a), (b), (c), and (d) only show counties that appear in both periods. The three-age-group series in panel (c) begins in 2010 because SAHIE combines ages 40–64 in 2008 and 2009. Panel (d) covers the full period because income groups are available beginning in 2008.

Table 1: Summary Statistics and Risk Score Micro Moments

Panel A: Summary statistics of county-year-age-income level data

	Pre-ACA (N=23383)		Post-ACA (N=71000)	
	Mean	SD	Mean	SD
Share Insured	0.378	0.262	0.503	0.251
Share Approved | Year, County, Age	0.47	0.174	1	0
Annual Premium (post-subsidy; $000s)	1.97	0.829	3.039	3.042
Annual Subsidy (APTC; $000s)	–	–	2.7	3.117
Actuarial Value (approx. % covered)	56.3	8.4	68.8	1.5
Actuarial Value, including CSR	–	–	72.4	5.5
Number of potential buyers	1790	5554	1371	4535
Avg. age of potential buyers	38	14	40	14
Avg. income of pot. buyers (%FPL)	303	88	309	87
Share pot. buyers under 50	0.761	0.153	0.712	0.143
Share pot. buyers under 40	0.567	0.159	0.507	0.139
Share pot. buyers FPL 138–249	0.363	0.157	0.337	0.137
Share pot. buyers FPL 250–399	0.294	0.127	0.294	0.116

Panel B: MEPS Risk score micro moments

Age	Income	Pre-ACA Average risk score when uninsured	insured	Δ^{pre}	Post-ACA Average risk score when uninsured	insured	Δ^{post}	Change in selection $\Delta^{post} - \Delta^{pre}$
	138-249	0.32	0.46	0.14	0.31	0.48	0.17	0.03
18-39	250-399	0.36	0.44	0.09	0.3	0.48	0.18	0.09
	400 or more	0.34	0.44	0.1	0.35	0.47	0.12	0.02
	138-249	0.5	0.62	0.11	0.45	0.64	0.19	0.08
40-49	250-399	0.48	0.6	0.12	0.44	0.63	0.19	0.07
	400 or more	0.48	0.61	0.12	0.51	0.59	0.08	−0.04
	138-249	0.68	0.78	0.1	0.64	0.8	0.17	0.07
50-64	250-399	0.67	0.77	0.1	0.62	0.77	0.15	0.05
	400 or more	0.67	0.76	0.09	0.65	0.77	0.12	0.04

Note: Panel A summarizes 2010–2012 and 2014–2019 county-year-age-income cells observed in the two datasets; *N* denotes the number of cells. Statistics are potential-buyer weighted except the potential-buyer count itself; potential buyers are the denominator of equation (1). Pre-ACA premiums and approvals are application weighted, these are application-level and not individual-level. Repeated applications must be accounted for. Post-ACA net premiums and APTC are computed using CMS data and year-specific IRS subsidy schedules, and aggregated using plan-enrollment weights. Pre-ACA actuarial value is calculated using plan-level deductibles and the distribution of health expenditure from the MEPS. Post-ACA actuarial value uses metal-tier values, varying by income. Panel B reports MEPS-weighted risk scores. Δ^{τ} is insured minus uninsured mean risk (discrepancies reflect rounding); the last column is its post-pre change. The Supplementary Appendix provides additional details on the data and construction of these variables.

Table 2: Post-ACA Demand Model Estimates and Implied Semielasticities

	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Summary of Post-ACA Demand Parameters						
Premium coefficient α	-0.03	-0.02		-1.41	-0.77	
	(0.003)	(0.003)		(0.32)	(0.09)	
$\alpha \times$ 138–249% FPL			0.01			-3.53
			(0.01)			(0.97)
$\times$ 250–399% FPL			-0.03			-2.51
			(0.01)			(0.74)
$\times$ 400+% FPL			-0.03			-0.45
			(0.003)			(0.13)
Mean utility $\times$ age 40–49	0.99	0.99	0.99	2.07	1.58	1.5
	(0.01)	(0.01)	(0.01)	(0.25)	(0.08)	(0.15)
$\times$ age 50–64	1.45	1.45	1.45	4.9	3.29	3.17
	(0.01)	(0.01)	(0.01)	(0.79)	(0.23)	(0.48)
$\times$ 250–399% FPL	0.53	0.54	0.58	1.17	0.87	0.78
	(0.01)	(0.01)	(0.02)	(0.15)	(0.04)	(0.32)
$\times$ 400+% FPL	1.23	1.21	1.26	8.3	4.98	0.41
	(0.02)	(0.02)	(0.02)	(1.62)	(0.47)	(0.57)
$\times$ risk	2.81	2.81	2.81	2.81	2.81	2.81
	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)
Fixed Effects	State	County + Year	County + Year	State	County + Year	County + Year
Panel B: % Change in Share Insured if Price Rises by $120/year						
Age 18–39	-0.18	-0.15	-0.09	-9.22	-5.06	-13.9
	(0.02)	(0.02)	(0.04)	(1.99)	(0.61)	(3.4)
Age 40–49	-0.12	-0.11	-0.06	-6.45	-3.51	-9.79
	(0.01)	(0.01)	(0.03)	(1.44)	(0.43)	(2.56)
Age 50–64	-0.1	-0.09	-0.05	-5.37	-2.91	-8.19
	(0.01)	(0.01)	(0.02)	(1.21)	(0.36)	(2.22)
138–249% FPL	-0.18	-0.15	0.08	-8.93	-4.91	-21.33
	(0.02)	(0.02)	(0.07)	(1.93)	(0.59)	(5.16)
250–399% FPL	-0.15	-0.12	-0.17	-7.56	-4.13	-13.23
	(0.01)	(0.01)	(0.05)	(1.66)	(0.5)	(3.7)
400+% FPL	-0.12	-0.1	-0.11	-5.97	-3.24	-1.92
	(0.01)	(0.01)	(0.01)	(1.33)	(0.4)	(0.56)
I risk quartile	-0.18	-0.15	-0.11	-9.06	-4.96	-12.25
	(0.02)	(0.02)	(0.03)	(1.99)	(0.6)	(3.08)
II-III risk quartile	-0.12	-0.1	-0.07	-6.04	-3.27	-8.11
	(0.01)	(0.01)	(0.02)	(1.36)	(0.4)	(2.18)
IV risk quartile	-0.07	-0.06	-0.04	-3.6	-1.93	-5.6
	(0.01)	(0.01)	(0.01)	(0.83)	(0.24)	(1.62)

Note: Panel A reports the estimates of the demand parameters for the post-ACA period for different specifications. The first three columns report OLS estimates, while the last three columns report IV estimates using the share of potential buyers younger than 50 in the year-county as an instrument for the price. Panel B reports the implied semielasticities of insurance demand with respect to price for different age, income, and risk groups. We compute semilasticities at the individual level as $\frac{\sigma_{it}(p_i+\$120)-\sigma_{it}(p_i)}{\sigma_{it}(p_i)} \times 100$, and then average across individuals in each cell. Standard errors in parentheses are computed from 100 county-cluster bootstrap replications.

Figure 2: Insurance Demand, Enrollment, Rejections, and Premium Changes

(a) Probability of Demand: $V_{it} \geq 0$

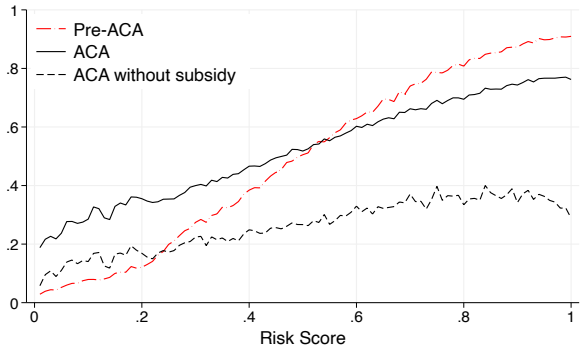

(b) Share Insured: $V_{it} \geq 0$ and $r_i \leq \psi_t(s_i, a_i)$

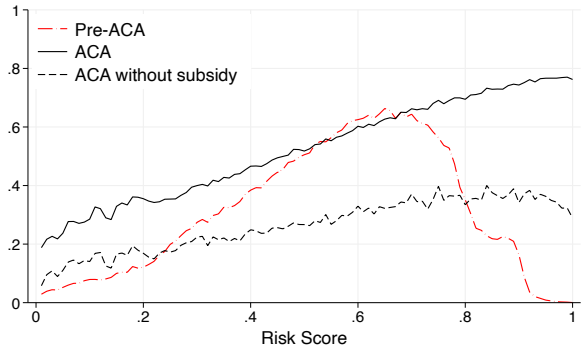

(c) Rejection Cutoffs vs. Risk Scores

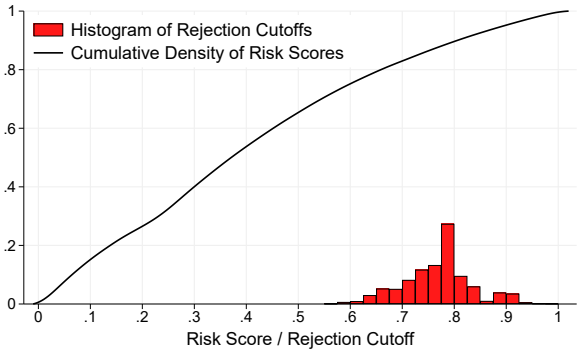

(d) CDF of Net-Premium Index $100\frac{P^{\text{post}}}{P^{\text{pre}}}$

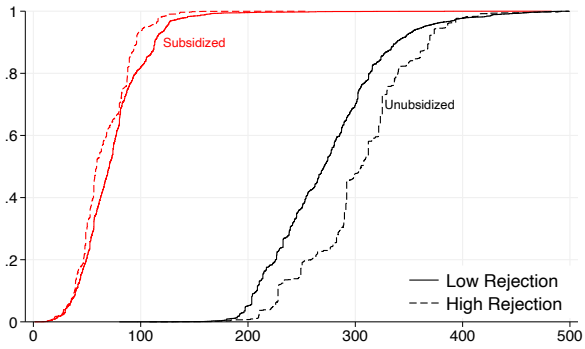

Note: Let $x = (c, a, y)$ denote a county-age-income cell, let $N_t(x)$ be its number of potential buyers, and let $\tau \in \{\text{pre}, \text{post}\}$ collect pre-ACA and post-ACA years in our data. For risk percentile r, define the cell-level period probabilities of demanding insurance and being insured as

$$d_\tau(x,r) = \frac{\sum_{t\in\tau} N_t(x)\sigma_t(x,r)}{\sum_{t\in\tau} N_t(x)}, \qquad q_\tau(x,r) = \frac{\sum_{t\in\tau} N_t(x)\sigma_t(x,r)\mathbf{1}\{r \leq \psi_t(s(c),a)\}}{\sum_{t\in\tau} N_t(x)},$$

where $\sigma_t(x,r)$ is the model-implied demand probability in equation (4) and $\psi_t(s,a) = \infty$ for $t \in$ post. Panels (a) and (b) plot potential-buyer-weighted averages of $d_\tau(x,r)$ and $q_\tau(x,r)$ across cells at each risk percentile. The short-dashed "ACA without subsidy" curves recompute post-ACA demand after replacing net premium $P_t(x)$ with gross premium $P_t(x) + S_t(x)$, holding all other post-ACA objects fixed. Panel (c) overlays the potential-buyer-weighted distribution of estimated pre-ACA state-year-age rejection cutoffs with the potential-buyer-weighted CDF of population risk scores; an applicant is rejected when $r > \psi_t(s,a)$. Panel (d) uses county-age-income cells observed in both periods and defines $P^\tau(x)$ as the potential-buyer-weighted mean net premium within period τ. It plots CDFs of $100P^{\text{post}}(x)/P^{\text{pre}}(x)$ by subsidy status and whether the mean pre-ACA rejection probability is in its top decile. An index of 100 denotes no change, and 200 denotes a doubling of net premium.

Table 3: Consumer Surplus and Pre-Post Comparison

Panel A: Distribution of Changes in Share Insured

Δ prob. insured conditional on age, income, county, risk	p5	p10	p25	p50	p75	p90	p95	Mean
Rejected	0.36	0.42	0.57	0.72	0.85	0.91	0.93	0.69
	(0.004)	(0.004)	(0.005)	(0.004)	(0.003)	(0.002)	(0.002)	(0.004)
Approved	-0.28	-0.2	-0.06	0.1	0.2	0.32	0.41	0.08
	(0.0154)	(0.0143)	(0.0081)	(0.0054)	(0.01)	(0.015)	(0.0148)	(0.002)
All	-0.27	-0.19	-0.03	0.12	0.28	0.59	0.78	0.15
	(0.0159)	(0.0143)	(0.0066)	(0.0061)	(0.0107)	(0.003)	(0.0051)	(0.0001)

Panel B: Distribution of Changes in Consumer Surplus

ΔCS(\$000s) conditional on age, income, county, risk	p5	p10	p25	p50	p75	p90	p95	Mean
Rejected	0.12	0.16	0.28	0.53	3.15	5.09	5.94	1.65
	(0.04)	(0.05)	(0.09)	(0.18)	(0.9)	(1.45)	(1.7)	(0.47)
Approved	-2.4	-0.95	-0.12	0.04	0.13	0.72	1.2	-0.12
	(0.84)	(0.35)	(0.05)	(0.01)	(0.05)	(0.21)	(0.35)	(0.07)
All	-2.11	-0.75	-0.06	0.05	0.26	0.99	2.24	0.1
	(0.75)	(0.29)	(0.03)	(0.02)	(0.08)	(0.29)	(0.64)	(0.06)

Panel C: Share with Positive Consumer Surplus Change by Age and Income

Age	18–39			40–49			50–64		
Income	138–249	250–399	400+	138–249	250–399	400+	138–249	250–399	400+
Rejected	0.98	0.98	0.99	0.99	0.99	0.99	1	1	1
	(0.005)	(0.006)	(0.004)	(0.003)	(0.003)	(0.002)	(0.001)	(0.002)	(0.001)
Approved	0.8	0.7	0.57	0.62	0.55	0.47	0.69	0.62	0.59
	(0.0034)	(0.005)	(0.009)	(0.003)	(0.003)	(0.0094)	(0.0016)	(0.0035)	(0.01)
All	0.81	0.72	0.61	0.68	0.62	0.54	0.77	0.71	0.67
	(0.004)	(0.0056)	(0.0093)	(0.0011)	(0.0042)	(0.0096)	(0.0014)	(0.0048)	(0.0097)

Note: The sample comprises county-age-income cells observed in both periods, each crossed with 100 risk percentiles. Within-period outcomes and their post-ACA minus pre-ACA changes are defined in equation (6). Reported means and quantiles weight cell-risk observations by average potential buyers and the estimated age-income risk distribution. Panel A reports changes in the model-implied probability of insurance. Panel B reports changes in annual consumer surplus in thousands of dollars; post-ACA eligibility is one and $\alpha(y)$ is the income-specific premium coefficient. Panel C reports the share with $\Delta CS \geq 0$. Mean pre-ACA rejection probabilities within cell-risk observations are nearly binary, with rare exceptions. We classify an observation as Rejected when this probability is at least 0.5, and as Approved otherwise. Consequently, the winner share for the Rejected group in Panel C can be less than one. Parentheses contain standard deviations across 100 draws based on the bootstrap variance-covariance matrix of the model parameters.