The Demand for Individual Insurance: Evidence from a Private Online Marketplace

Conor Ryan, Roger Feldman, Stephen Parente

We use a novel data set from a private online marketplace to estimate the demand for individual health insurance among a set comprising many high-income households across 18 states. Households earning more than 4 times the federal poverty level (FPL) are willing to pay $30 to $135 per month to increase the actuarial value of their insurance by 10 percentage points, much less than households earning less than 2.5 times FPL. Higher-income households are also less likely to forgo insurance due to a premium increase. These results are important for understanding the effect of health reform proposals targeting higher-income populations.

Keywords: Health Insurance, Consumer Demand

JEL Codes: D12, I13

The non-group health insurance market, or individual market, is the only insurance market available to nearly 40 million Americans who do not qualify for any public health insurance program or receive an offer of health insurance through an employer. Prior to the Affordable Care Act (ACA), this market exhibited low take-up rates and poor protection against risk. The reforms introduced to the individual market by the ACA can be sorted into two broad categories. The first category targets the supply of health insurance. The law directly regulates the premiums and characteristics of insurance plans, requiring a threshold of insurance quality and restricting explicit price discrimination based on health status. While restricting price discrimination is desirable from a policy perspective, it worsens the traditional problem of adverse selection that plagues insurance markets.

The second category of reforms targets consumers with a set of subsidies and a penalty for being uninsured. In addition to directly assisting with insurance affordability, these policies intend to mitigate adverse selection by encouraging broad consumer participation. This is similar to the design of Medicare Part D, which regulates the private prescription drug insurance market for Americans over the age of 65. A 75% premium subsidy and steep penalties for late enrollment have been largely successful in solving the severe adverse selection problem in prescription drug coverage (Atherly and Dowd 2009; Heiss, McFadden, and Winter 2009).

Higher-income consumers in households that earn more than 400 percent the federal poverty level (FPL) have previously been ineligible for premium assistance and are an increasingly frequent target of policy making in the individual market. Recent reforms include a temporary extension of health insurance premium subsidies to higher-income households and efforts to make less comprehensive insurance products available which might be attractive to households that do not receive subsidies and are more able to bear the out-of-pocket risk of medical expenditures.

The necessity and efficacy of these reforms depend on the demand for health insurance among the targeted set of consumers. In this paper, we estimate the demand for health insurance using a novel data set that provides evidence on the elasticities of consumer demand in a new segment of the individual market: higher-income consumers. A large and growing literature estimates discrete choice demand models for individual health insurance (K. M. M. Ericson and Starc 2015; Shepard 2016; Jaffe and Shepard 2017; Tebaldi 2020;

Saltzman 2019; Drake 2018), but because of data availability this literature focuses on purchases made through government-run exchanges in a few states—primarily in Massachusetts and California. While these data provide opportunities to study many aspects of market design, they typically do not include many higher-income households that make up roughly a third of eligible consumers and face higher, unsubsidized premiums.

Since government-run exchanges are the primary source of subsidies and the target of enrollment efforts, most exchange consumers earn less than 400 percent of FPL and currently qualify for subsidies—87% of households buying insurance through the exchanges receive premium subsidies (ASPE 2016)—and a substantial amount of previous work has found that income is important in explaining health insurance choices (e.g., Tebaldi 2020; Saltzman 2017). This same income group represents only 69% of the national market.[1]

Our data, which come from a large private online insurance broker, contain roughly equal shares of lower-income individuals who are eligible for subsidies and higher-income individuals who are not. Therefore, it provides new evidence on the demand for individual health insurance among higher-income households who do not typically purchase insurance through the government-run exchanges. An additional advantage of our data is its broad geographic coverage that complements existing insurance demand studies. Our estimation sample has 126 local rating areas in 18 states.

We model consumer demand for individual health insurance using a nested logit model, with

[1] Authors' calculations using the American Community Survey.

different levels of insurance coverage in a single nest and the outside option of being uninsured in a single-component nest. The data provide two sources of identification for the premium elasticity of insurance demand. The first is within-market variation in household-level premiums that results from the premium regulations in the ACA (Tebaldi 2020; Saltzman 2019; Drake 2019). The second source of variation comes from cross-sectional differences in premiums across different geographies. Geographic variation in health care costs and spending has been extensively documented; this provides an additional source of plausibly exogenous variation in the cost and price of insurance (Finkelstein, Gentzkow, and Williams 2016; Dartmouth Atlas).

We find that consumers' willingness to pay for additional insurance falls dramatically with income. Single-person households in the lowest income group, earning less than 250 percent of FPL are willing to pay $164 to $177 per month for an additional 10 percentage points in actuarial value. This is analogous to reducing the gross expected coinsurance rate by 10 percentage points. Single-person households in the highest income group, earning more than 400 percent of the federal poverty level, are willing to pay $30 to $46 per month, a reduction of 74 to 82 percent. This same relationship between income and willingness to pay for additional insurance is present among non-single households as well, but the differences are less drastic. We also find that consumers in our data are highly price-elastic, and the lowest-income consumers are especially sensitive to insurance premiums.

The main contribution of this paper is to the literature on the demand for individual health insurance. A recent literature has investigated many aspects of consumer demand in

the markets newly regulated by the ACA (Frean, Gruber, and Sommers 2017; Dafny, Gruber, and Ody 2015; Abraham et al. 2017; Drake 2018; Tebaldi 2020; Saltzman 2019). Our paper builds on this work by studying the demand of consumers who purchase insurance through a third party, online broker. We find price elasticities at the upper end of estimates in the literature, which suggests that consumers seeking insurance via this third party platform may be more price-conscious than previously reported.

Our demand estimates imply that the semi-elasticity of insurance coverage with respect to a $100 increase in the mandate penalty is -7.0, -1.8, and -4.2 for households who earn less than 250 percent of FPL, 250 to 400 percent of FPL, and more than 400 percent of FPL, respectively. For example, a $100 increase in the annual penalty leads to an increase of 7.0 percent in the share of insured among the lowest-income consumers. These estimates are consistent with the literature, which finds estimates of the semi-elasticity of insurance coverage between -1 and -10, and higher elasticities among low-income populations. Many of the estimates on the high end of this range are based on low-income households (Finkelstein, Hendren, and Shepard 2019; Jaffe and Shepard 2017), and lower estimates among higher-income populations (Tebaldi 2020; Sacks, Lurie, and Heim 2018; Hackmann, Kolstad, and Kowalski 2015; Saltzman 2019).

Section I describes the institutional details of the post-ACA individual health insurance landscape. In Section II, we describe the data. In Section III, we outline the demand model. In Section IV, we present the results of the demand estimation.

I. The Individual Market for Health Insurance

The individual market covers roughly 18 million people and offers coverage to 20 million more individuals who remain uninsured. Local markets are typically characterized by competition among several insurance companies offering dozens of different health insurance plans. The set of available plans varies widely over 501 state-demarcated geographic "rating areas," but the financial characteristics of all plans are somewhat standardized. Every insurance plan is required to cover certain benefit categories and preventive services. Total allowable out-of-pocket expenditures must be below a federal limit of $8,200 for an individual in 2020. The main feature of product differentiation is the network of medical providers that are available within a particular insurance plan.

Companies offer menus of health insurance contracts that fall into four "metal" categories of ascending generosity: Bronze, Silver, Gold, and Platinum. The metal levels correspond to plans with actuarial values (the percentage of expected expenditures covered) of 60%, 70%, 80%, and 90%, respectively. A fifth category of Catastrophic plans covers many of the same benefits but with exceptionally high deductibles, typically equal to the maximum allowable out-of-pocket expenditure. Companies set base premiums that may vary by insurance plan and rating area.

In contrast to group health insurance, premiums in the individual market are household-specific. The premium of any given plan varies by household size, the age of each household member, and whether or not the enrollees smoke, according to a formula specified by regulation. Most states have adopted the federal default age rating curve which, in 2015, required insurance companies to charge the base premium to consumers between 21 and 24

years of age. The premium increases monotonically to 3 times the base rate for a 64-year old and charges a flat 0.635 times the base rate for children under the age of 21.[2] A household purchasing a family plan pays the cumulative premium for all adults and the first three children. Additional children are covered for free. States may require that smokers pay up to 50% higher premiums, and 40 states have chosen the maximum smoker surcharge.

Consumers are eligible for an income-adjusted premium subsidy if they do not have an affordable offer of employer-sponsored insurance and are not eligible for another public insurance program. Subsidy eligibility begins at 100 percent of the Federal Poverty Level (FPL)—an income threshold that accounts for household size—and gradually phases out up to 400 percent of FPL. Households at 100 percent of FPL are offered a fixed subsidy that allows them to purchase the second-lowest price Silver plan for approximately 2 percent of household income. The subsidy declines monotonically with income, until households at 400 percent of FPL pay 9.5 percent of household income for the second-lowest price Silver plan. Consumers who earn more than 400 percent of FPL do not receive any subsidy.[3] Premium contributions for those receiving subsidies are adjusted annually to reflect the excess premium growth over income growth. To date, these increases have been small (1-2 percent every year).

Consumers between 100 percent and 250 percent of FPL receive additional cost sharing

[2] In 2018, the flat rate for children under 20 was replaced by a descending rate for each year down to 0.765 for children under 15.

[3] Beginning in 2020, California became the first state to extend subsidies beyond the federal limit, offering premium assistance up to 600 percent of FPL.

subsidies that increase the generosity of Silver plans through lower deductibles and maximum out-of-pocket expenditure limits. These benefits are very large for households earning between 100 and 150 percent of FPL and smaller for those earning over 200 percent of FPL.

A household that did not purchase health insurance was subject to the individual mandate penalty until 2019. The penalty was the maximum of a fixed per-person fee and a percentage of income. The fixed fee was $95 per adult in 2014, $325 per adult in 2015, and $695 per adult since 2016. The penalty for children was half of the adult value. The percentage of income penalty was 1% of income above the tax-filing threshold for the household's filing status in 2014, 2% in 2015, and 2.5% since 2016. In both cases, the maximum penalty was the national average annual premium for a Bronze plan.

Catastrophic plans were considered sufficient coverage to avoid the penalty only for individuals under the age of 30 or households that qualified for a hardship exemption. Others were required to purchase at least a Bronze plan. The individual mandate penalty was cut to $0 by the Tax Cuts and Jobs Act of 2017 (PL 115-97), effective in 2019. Individuals are nominally required to purchase insurance, but there is no penalty for choosing not to.

II. Data

A. CHOICE DATA

The primary data for the choice model are individual plan choices through a private online marketplace that sells plans both on and off the ACA health insurance exchanges.

Since the implementation of the ACA, households seeking coverage in the individual market can purchase insurance plans through three main channels. First, households can purchase insurance directly from an insurance firm, either through the firm's website or by phone. Second, households can enroll in insurance plans through the government-run ACA health insurance exchanges in each state.[4] The plans offered through the government-run website are often referred to as "on-exchange" and are subject to some additional regulations, such as requiring each participating insurance firm to offer both Silver and Gold plan options. Finally, households can use the assistance of a third-party exchange or broker to view available options and enroll in a plan. We look at choices made through one such broker—a large private online marketplace. In addition to "on-exchange" plans, the online marketplace offers some "off-exchange" plans as well. This distinction is minimal for insurance firms that offer some plan types on and off the exchange.[5] However, some insurance firms do not participate in the government-run exchanges in some states and are therefore unavailable through the government-run platform.

Consumers can visit this marketplace to view the available insurance options in their rating area and enroll in one of those plans. In 2015, this marketplace was authorized to sell subsidized health insurance plans in many states, including states that use the federal HealthCare.gov

[4] In many states, this is Healthcare.gov. However, some states run their own insurance portal, e.g. CoveredCA.com in California.

[5] Analysis of the RWJF HIX 2.0 data show that premium and cost-sharing differences for plans on and off the exchange, among firms that offer both, are negligible.

exchange.[6] The data contain the choices of subsidized and unsubsidized consumers in 48 states.

The data contain information on the age of the head of household, household income (with some missing values), the number of members in the household, and whether or not the household contains a smoker. The data also contain the value of a subsidy received and information about the insurance plan selected including the metal level and the plan name. Because the choice sets used in estimation are aggregated to the metal level (further explained in Section II.C), the only information we use from this data about the selected plan is the selected plan's metal level.

An observation in the data represents a household, but we observe only one age and smoking status. We assume this is the age and smoking status of the head-of-household who purchased the plan. However, to match the household to its relevant choice set, we have to know the ages of every adult (over the age of 14) in the household. If the household contains two people, we assume that it contains two adults of the same age. If the household contains more than two people, we assume that it contains two adults of the same age and that all additional persons in the household are children under the age of 20. We assume that the smoking status refers only to the head of household and apply the maximum smoking penalty allowed by the household's state.

[6] In all states in our estimation sample (and some others), the government-run marketplaces allowed the private online marketplace in our data to interface with the back-end of the government enrollment process in order to administer subsidies to their consumers.

We can test our assumptions on household composition for a subset of observations where where we observe the total price paid for the selected plan. Using the observed total price and the the median premium within the selected metal level and insurance firm, we can impute the household-level age rating. The correlation between our simple age rating rule and the age rating in the imputed sample is 0.90. We experimented with more detailed rating imputations based on the observed sample of prices paid and found it does not significantly affect our results.

We observe or can compute the income of every individual who receives a subsidy. The subsidy formula comes from a direct function of the price of the second-lowest price Silver plan in the rating area, the household-level age rating, and household income expressed as a percentage of FPL. By inverting this equation, we can compute the income of any household that receives a positive subsidy. There is some error in this computation due to the assumption about household age ratings, but it is negligible. Among the 41,000 households for which we observe both household income and the value of a subsidy, the correlation between household income and the imputed income measure is more than 0.99.

We do not observe the incomes of most individuals who do not receive subsidies. We assume these individuals have income levels that make them ineligible for subsidies. This assumption is not terribly restrictive. It requires us to assume that every individual eligible for a subsidy receives a subsidy, or at least selects a plan as if they would receive the subsidy for which they are eligible. There is some evidence that a non-trivial amount of subsidy-eligible consumers do not receive them on a monthly basis. However, all eligible consumers should eventually receive

the full value of the subsidy for which they are eligible when they file their income taxes.

We restrict the analysis to markets in which we observe the entire choice set and can be reasonably confident that the marketplace represents the complete choice set of health insurers. Using Medical Loss Ratio reporting data, we observe aggregate state-level market shares for health insurance firms. We throw out any markets in which there are no purchases from insurance firms that have more than 5% market share in the state. In this way, we hope to ensure that the sample of choices is not limited to only a portion of the market. We discuss the choice sets in Section II.C.

After dropping additional observations because of missing data, the remaining data set includes roughly 92,500 household health insurance choices across 126 rating areas in 18 states.[7]

B. UNINSURED

Our data on insurance product choices do not include the outside option: the choice to be unsinsured. However, we can observe households that choose to be uninsured in the 2015 American Community Survey (ACS). We use the ACS to construct market shares for uninsurance among households with similar observable characteristics as those in our data. We match uninsurance rates in the ACS to households in our choice data conditional on the state in which they live, whether or not the head of household is over the age of 35, whether

[7] We drop 10 rating areas where the ACS sample does not have any responses from a household in one or more of the demographic groups used to calculate uninsurance rates.

or not the household is eligible for a subsidy, and whether the household has one, two, or at least 3 members. One limitation is that this relies on the assumption that households that select insurance through the private online marketplace and appear in our data are similar to households in the ACS, conditional on age, income, and household size.

We consider the population that might purchase individual health insurance to be any legal US resident who is not eligible for Medicaid or Medicare and does not have affordable access to health insurance through their employer. Technically, any individual can switch from these insurance categories to the individual market at any time; however, the insurance plans in the individual market are considerably more expensive and typically require more cost sharing. This type of switching is likely to be infrequent. We consider an individual to have an offer of health insurance through an employer if they are currently enrolled in such a plan, their spouse is enrolled, or their parent is enrolled and they are still a tax dependent. We consider this offer to be affordable for family coverage if the average employer-sponsored premium in 2015, $4,955, is less than 9.5% of household income (Claxton et al. 2015). Dependents who have access to employer-sponsored insurance through the head of the household that exceeds 9.5% of household income are still eligible for premium subsidies in the individual market. Some individuals have an offer of employer-sponsored insurance but do not accept it, and we cannot observe them. We treat these consumers as identical to other participants in the individual market, though by law they cannot receive health insurance subsidies. This population is small (Planalp, Sonier, and Fried 2015).

To identify Medicaid coverage and tax dependents, we adapt a methodology outlined by the

Government Accountability Office (GAO 2012). Medicaid eligibility is determined by state-level eligibility categories defined by income, age, and family status. We assume that everyone who is enrolled in Medicaid in the ACS is eligible. To address under-reporting of Medicaid enrollment, we define any parent who receives public assistance, any child of a parent who receives public assistance or is enrolled in Medicaid, any spouse of an adult who receives public assistance or is enrolled in Medicaid, or any childless or unemployed adult who receives Supplemental Security Income payments as being enrolled in Medicaid. An individual is considered eligible for Medicaid or CHIP if his or her household income falls within state-specific eligibility levels. We assume that individuals who are determined to be eligible for Medicaid, but report enrollment in private individual or group coverage, are enrolled in Medicaid. We believe this corrects for those who confuse Medicaid managed care programs with private coverage, and Medicaid with employer-sponsored insurance.

C. CHOICE SETS

We observe only the ultimate choices made by consumers, not the set of available options (the choice set). To construct choice sets, we use the HIX 2.0 data set compiled by the Robert Wood Johnson Foundation. This data set provides detailed cost-sharing and premium information on plans offered in the individual market between 2014 and 2017. The data are nearly a complete depiction of the individual health insurance market for the entire US, but in some markets, cost-sharing information is missing or insurance firms are missing altogether.

The choice sets in each market are large. The median number of choices per market is 166, and these plans do not necessarily overlap with other markets. Because we observe only

a sample of choices, we do not observe many plans being chosen. This does not necessarily imply that these plans have no market share, but simply that the choice set is large relative to the observed number of choices.

To address the large number of plans relative to the number of choices, we model only five categorical choices, which correspond to the four metal levels and the Catastrophic plan. This requires us to aggregate plans into plan types. We assume the consumer's plan choice is the plan with the 25th percentile premium in each product category in that rating area. If that premium corresponds to plan A, then we set the cost sharing characteristics (deductible and maximum out-of-pocket spending) equal to those of plan A. We prefer this method over separately aggregating the premium, deductible, and maximum out-of-pocket expense, as those variables may be related in endogenous ways. This specification is similar to taking the median plan characteristics, or aggregating separately to the mean or median of each individual characteristic.[8] This approach abstracts from demand differences across insurance firms, which is not the focus of our paper.

D. DESCRIPTIVE FACTS

We present two sets of descriptive facts about the data. First, we describe the demographics of consumers in our data, and we show evidence that income is an important factor driving selection into the private online market relative to other avenues for purchasing insurance. Second, we show descriptive evidence on how consumer choices of whether to buy insurance and which insurance plan to buy depend on the price of insurance, which varies across

[8] We have estimated our model using the median and mean product characteristics among all offered plans and the share-weighted mean of product characteristics. Our results under each of these methods are quantitatively similar.

households and geography.

In Table 1, we summarize the private online market and compare it to two other on the individual insurance market: the 2015 American Community Survey (ACS); and enrollment through the federal health insurance website, HealthCare.gov. The ACS offers the broadest depiction of the health insurance market across all purchasing platforms. The US Department of Health and Human Services (HHS) publishes detailed descriptive statistics for health insurance purchases through HealthCare.gov, which allows us to make plan choice comparisons with our data. All data sources are restricted to the same set of US states that are in our estimation sample.

The maximum age of households in the private online market is similar to the population of eligible households in the ACS. The ages of consumers who purchase insurance through HealthCare.gov are similar, but we only have data on the age distribution of all members rather than the maximum age of each household.

Consumers in the private online market have substantially higher incomes than those in the ACS and those purchasing through Healthcare.gov. In the private online market, 57% of consumers earn over 400 percent of FPL, while only 20% of consumers in the ACS and only 9% of consumers purchasing through Healthcare.gov meet that income threshold.

The most popular insurance products in both the private market and HealthCare.gov are Silver plans, though these plans have a much higher market share on Healthcare.gov. Nearly

70% of consumers on HealthCare.gov purchase a Silver plan, while less than half of consumers on the private online marketplace do so.

The preference for Silver plans is closely related to consumer income, both because premium subsidies target the affordability of Silver plans and because households earning less than 250 percent of FPL receive cost-sharing reductions when purchasing Silver plans (see Section I for more details). To show the importance of selection on income, we also compute demographic distributions in the private online marketplace after re-weighting to match the geographic and income distribution of consumers who purchase insurance through HealthCare.gov in the same states. We compute this re-weighting using 50 percentage point increments of household income as a percent of FPL and the rating area. The results are presented in the final column of Table 1.

After re-weighting the estimation sample by income, the market share of each insurance product in the private online marketplace is nearly identical to that in HealthCare.gov.[9] The re-weighting is also associated with some changes to the age distribution of households, but these changes are relatively small. This suggests that income is an important aspect of selection into the private online marketplace and that consumers appear to have roughly similar preferences, conditional on income.

In Table 2, we demonstrate the variation in premiums and choices across demographic groups. Due to the age rating curve, households with the oldest members face premiums that are 2 to 3

[9] The similarity is driven almost entirely by weighting on income. We can also re-weight the sample based on geography alone, but the distributions are not substantially different from the estimation sample.

times higher than those with young members. And due to premium subsidies, the lowest-income households, earning less than 250 percent of FPL, face premiums that are roughly $200 per month cheaper than higher-income, subsidy-ineligible households. Older and younger households select very similar plans, but income is associated with large differences in plan selection. A large majority of the lowest-income households select Silver plans, while only 29 percent of subsidy-ineligible households do so. This is partially due to the premium subsidies, but also the additional cost-sharing reductions that are provided to the lowest-income households enrolled in Silver plans.

In Figure 1a, we plot the share of insured individuals relative to the median premium of a Silver plan in the household's rating area, unadjusted for the age, income, or size of the household. Each dot on the graph displays the average share of consumers who are insured based on $20 increments in the median premium of a Silver plan, and the size of the dot represents the number of consumers. Higher-income, subsidy-ineligible consumers are more likely to be insured, and consumers facing a higher price of insurance are less likely to be insured. However, the relation between the insurance rate and price is not very strong. For consumers who earn at least 250 percent of FPL, a $100 increase in the monthly premium is associated with only a 5 percent decline in the share of insured. For consumers who earn less than 250 percent of FPL, the share of insured is flat relative to the insurance premium.

We also show that consumers facing a higher relative premium of Silver plans are less likely to select Silver plans. In Figure 1b, we plot the share of consumers who choose a Silver plan among households that purchase either Silver or Bronze relative to the difference between the

median Silver plan premium and the median Bronze plan premium. Each dot on the graph displays the average Silver share for $5 increments in the difference between the Silver and Bronze premium, and the size of the dot represents the number of consumers. Subsidy-eligible consumers are more likely to select Silver plans over Bronze plans. This is especially true for consumers who earn less than 250 percent of FPL and receive cost-sharing subsidies that increase the value of Silver plans. However, it remains true even for income groups above 250 percent of FPL that face the same level of Silver plan cost-sharing.

This variation is central to the demand estimation that follows. The parameters will be identified using variation in the monthly premium for the same products across households, as well as similar households in different geographic areas that face different prices. A detailed description of the model follows in the next section.

III. Empirical Model

In this section, we present the empirical discrete choice model of health insurance demand. We first present a general description of the environment and model, and then provide details on the estimation and identification of the key parameters.

A. MODEL DESCRIPTION

There are R rating areas, indexed by r, and J_r health insurance plans offered to the households in each region. Product characteristics are the annual premium (p); observed attributes (X) which include the annual insurance deductible and the maximum allowed out-of-

pocket spending; and unobserved product quality (ξ).

The choice set for each household consists of up to five available metal levels (See Section II.C). Every household has at least three options: Bronze, Silver, and Gold. Only some rating areas have Platinum plans available and only households with no one over the age of 30 may purchase Catastrophic plans.

Household demographics (Z) include age, income, and household size. The premium of each insurance plan depends on these characteristics. In particular, a household has an associated age rating factor (a) that depends on the size, age composition, and smoking status of the household, and a premium subsidy (b) that depends on the size and income of the household. The premium paid by household i for plan j in region r is a linear function of the base premium of the insurance plan, $p_{ijr} = a_i p_{jr} - b_{ir}$, with a minimum allowable premium of $0.

If a household does not select an insurance plan and instead decides to be uninsured, the household is still charged a price equivalent to the mandate penalty (m), which also depends on household characteristics. Since the mandate penalty was in place during 2015, the time period of our data, we maintain this price in estimation.

Households choose a plan j or the outside option of being uninsured to maximize utility. The indirect utility of household i in rating area r selecting plan j or the outside option ($j = 0$) is given by:

$$u_{ijr} = \gamma_i + \alpha_i p_{ijr} + \beta_i' X_{jr} + \xi_{jr} + \epsilon_{ijr}$$

$$u_{i0r} = \alpha_i m_{i0r} + \epsilon_{i0r}$$

where γ captures how the mean value of insurance depends on household demographics, β represents preferences for observed product attributes, and α is the utility value of money, which applies equivalently to the mandate penalty and the insurance premium. We do not allow individuals to value the mandate differently than premium dollars. Ericson and Kessler (2016) find in an experiment that, in the wake of media coverage of the mandate penalty articulated as a tax, consumers viewed the mandate plainly as a monetary fine. There is some evidence that individuals may be more sensitive to a mandate because of a preference for compliance (Saltzman 2019), but individuals appear to be less sensitive to the statutory level of the mandate penalty because they pay less than the full value of the penalty (Lurie, Sacks, and Heim 2019).

B. ESTIMATION STRATEGY

Our estimation strategy captures two features of the demand for health insurance. First, we quantify the extent to which being uninsured is a close substitute for insurance products. Roughly half of the households in each market area we analyze are uninsured, a much higher rate than the market share for any particular product. Such a large market share would imply substantial substitution between each insurance product and uninsurance, unless insurance products are closer substitutes to one another than to uninsurance.

To capture this pattern of substitution, we estimate a nested logit specification that allows individuals to have unobserved idiosyncratic preferences for insurance. Formally,

$$\epsilon_{ijr} = \zeta_{ir} + (1-\sigma)\omega_{ijr} \quad \forall\, j \neq 0$$

where ω_{ijr} is distributed by type I extreme value, and ζ_{ir} is distributed such that ϵ_{ijr} is type I extreme value. The parameter σ governs the substitutability of insurance products and uninsurance. The case of equal substitutes if given by $\sigma = 0$, with $\sigma \to 1$ implying that the products are not substitutes for uninsurance.

The second important feature of demand is that premiums may be correlated with unobserved product quality, leading to potentially biased estimates of the premium elasticity of demand and the effect of the insurance mandate penalty. We identify the premium elasticity by using the regulated, non-linear price function itself as a source of variation in premium uncorrelated with the unobserved quality of insurance (Tebaldi 2020; Drake 2019; Saltzman 2019).

We aggregate consumers into demographic groups θ, which consist of three-year age brackets, household size (up to 4 members), and 50 percentage point increments in household income relative to the federal poverty limit (up to 400 percent).[10] From Berry (1994), we can write the nested logit model as a linear equation:

[10] Due to small samples, we do not further divide the observations by smoking status, but the smoking prevalence in each demographic group is incorporated into the calculation of $p_{\theta jr}$. The observed smoking status rates do not demonstrate strong trends across age and income groups and do not vary substantially.

$$\log(S_{\theta jr}) - \log(S_{\theta 0r}) = \gamma' Z_{\theta r} + \alpha_\theta (p_{\theta jr} - m_{\theta 0r}) + \beta'_\theta X_{jr} + \sigma \log\left(\frac{S_{\theta jr}}{1 - S_{\theta 0r}}\right) + \xi_{\theta jr} \quad (1)$$

where $S_{\theta jr}$ is the market share of product j in market r among demographic group θ, $p_{\theta jr}$ is the average premium of the product for consumers in demographic group θ in market r, and $m_{\theta 0r}$ is the average mandate penalty. We model the preference for insurance as $\gamma_i = \gamma' Z_{\theta r}$. The vector of demographic characteristics $Z_{\theta r}$ includes 4 age categories (18-30, 31-40, 41-50, and 51-64), three income categories (less than 250 percent of FPL, 250 to 400 percent of FPL, and more than 400 percent of FPL), and whether the household is a single consumer or not. The premium elasticity parameter and the preferences for observed attributes also depend on these same household demographic characteristics, i.e. $\alpha_\theta = \alpha' Z_{\theta r}$. Importantly for identification, these fixed effects and interaction terms are coarser than the demographic groups that define θ.

We allow the unobserved quality, $\xi_{\theta jr}$, of each insurance product to potentially be different across demographic groups. We use two sets of fixed effects to assist in the identification of the premium elasticity. First, we use fixed effects for the rating areas in which the products are offered and a dummy variable for whether or not the insurance plan exceeds an actuarial value of 80%. Second, we use the Cartesian product of these two categorical variables. Through these fixed effects, we allow the quality of an insurance product to be market-specific, and in the most flexible specification, we allow variation in quality across markets to be different for the most and the least comprehensive insurance products.

The key assumption to identify price sensitivity is that $p_{\theta jr}$ and $\xi_{\theta jr}$ are uncorrelated. Our data contain two plausibly exogenous sources of variation in the insurance premium. The first

source is due to the premium regulation formula, which adjusts the premium paid for the same product according to the age, income, and size of the household. While we allow preferences for insurance to depend on coarse measures of these demographics, the remaining variation within each coarse demographic groups provides identification. This argument is outlined in detail in Tebaldi (2020). The second source of identification comes from geographic variation in insurance premiums, driven in part by the large variance in health care costs (Finkelstein, Gentzkow, and Williams 2016; Dartmouth Atlas). After controlling for variation in the quality of insurance in each rating area through fixed effects, some identifying variation remains in the relation between consumer behavior and the relative prices of different insurance products in each area. In Section II.D, we describe the variation in the data that drives this identification.

Intuitively, the nesting parameter σ is identified through the relation between the value of the nest of inside goods, i.e. the different types of insurance products, and the share of consumers who decide to purchase any insurance at all. While the value of the inside good is only a model concept, the same variation in price across households can identify this parameter through its association with the likelihood that those households are insured.

The nesting term in the regression, $\log\left(\frac{S_{\theta jr}}{1-S_{\theta 0r}}\right)$, requires an instrument to deal with the easily apparent endogeneity. We use a measure of the "leave-out share"—the average inside market share of the same metal level for households of type θ in all states other than the state that includes region *r*. This is similar to other methods that have been used to identify both

the nesting and premium elasticity terms (e.g. Panhans 2017), and it targets the key exogenous variation in the data that motivates both the identification of the nesting and price elasticity terms.

Finally, we estimate two versions of our demand equation with different sets of observed product characteristics. The first estimation uses two measures of the out-of-pocket risk of the insurance plan: the deductible and excess out-of-pocket risk after the deductible. The deductible measures the amount consumers must spend before receiving any insurance benefits. The excess out-of-pocket risk measures the difference between the deductible and the maximum allowable out-of-pocket spending in the insurance plan. In the second estimation, we use the actuarial value of the insurance plan to summarize the observable attributes. The actuarial value represents the average expected insurance coverage rate of the plan. For example, an insurance plan with an actuarial value of 70% is expected to cover 70% of a consumer's medical expenses, on average.

Since the demographic groups are small, the data contain some observations where $S_{\theta jr}$ is equal to zero. We follow the common solution to this problem by adding a small constant to each share. We estimated the demand equation with a wide range of constant values and present results for the value of 10^{-6}. At this magnitude, the results are robust to changes in the value of the constant.

IV. Estimation Results

In Table 3, we present the results of our demand estimation. Columns 1 and 2 present results

for the demand specification with deductible and excess out-of-pocket risk with each set of fixed effects. Columns 3 and 4 present results for the demand specification with the actuarial value of insurance with each set of fixed effects. The nesting parameter (σ) and the premium sensitivity parameters are consistent across all of these specifications. The willingness to pay for insurance, while measured differently in each specification, is also consistent. We will focus in this section on the specification with actuarial value and the finest set of fixed effects, presented in column 4.

Table 4 presents the results from the first stage regression that identifies the nesting parameter. The leave-out share is a statistically significant predictor of the endogenous nesting term, and the F-statistic is larger than 10,000. Therefore, there is little concern for a weak instrument.

Demographics are an important determinant of premium sensitivity, which decreases with age, household size, and income. Income has the largest effect on premium sensitivity. Consumers that earn less than 250 percent of FPL are roughly twice as elastic to the premium as consumers earning more than 250 percent of FPL. Families are less premium sensitive that single households, and while households with older members are slightly less premium sensitive, these findings are not statistically significant.

Consumer preferences for additional insurance are declining with income. In Table 5, we display consumer willingness to pay for 10 additional percentage points of actuarial value, analogous to moving from a Bronze plan to a Silver plan. For nearly every demographic

group of age and household size, the willingness to pay for insurance is decreasing in income. In the youngest age group, the willingness to pay for additional insurance declines from $164 per month and $212 per month for single and non-single households earning less than 250 percent of FPL to $30 per month and $95 per month for single and non-single households earning more than 400 percent of FPL. This large decline is typical of each age group and monotonic across the three income categories, with the exception of non-single households with a member over the age of 50.

These results suggest that higher-income households are more willing to self-insure their out-of-pocket costs. This is consistent with other research that shows protecting against medical expenses is a motivation for wealth accumulation (De Nardi et al. 2010). When insurance is sold at a premium above its actuarial value, households that have the ability self-insure against medical risks may choose to do so rather than pay the markup for additional insurance. In this environment, households have more affordable access to insurance that covers all catastrophic medical events that exceed about $6,600 during the year, which likely makes the choice to self-insure for the remaining out-of-pocket expenses more attractive.

While consumers earning less than 250 percent of FPL have higher willingness to pay for additional insurance, they on average get a lower total surplus from the extensive margin decision of purchasing insurance. Average consumer surplus for Bronze plans, the cheapest available option, is -$140 per month among all potential consumers earning less than 250 percent of FPL, and the average consumer surplus for the more generous Silver plans is $1.6 per

month.[11] While consumer surplus for Silver plans is positive, it is still quite small when compared to consumers earning more than 400 percent of FPL, where the average surplus potential consumers is $514 and $536 per month for Bronze and Silver plans, respectively. Low-income consumers in our data have relatively high rates of uninsurance, and conditional on purchasing insurance, they select generous insurance products. This data drives the finding of both a high willingness to pay for additional insurance and low total consumer surplus.

In Table 6, we present the estimated cross-product semi-elasticities for each income group. These semi-elasticities represent the percentage change in the aggregate market share of each metal level (rows) as a result of a $100 increase in the annual premium of another metal level (columns). We also include the semi-elasticity with respect to a $100 increase in the mandate penalty. The nested logit demand model assumes that semi-elasticities between products within each nest are identical, which appears for Bronze, Silver, and Gold plans. The semi-elasticities for Catastrophic and Platinum plans differ slightly because they are offered only to a subset of consumers. Our estimated nesting parameter σ governs the extent to which these within-nest semi-elasticities are similar to the semi-elasticity of uninsurance.

We find that the lowest-income consumers who earn less than 250 percent of FPL are the most premium elastic, especially for products other than the Silver plans. A $100 increase in

[11] We calculate average consumer surplus as the mean value of the monetary value of utility, $E[u_{ij}/\alpha_{ij}]$, where we use the Euler–Mascheroni constant as the expected value of the idiosyncratic shock. This is not conditional on the consumers that end up purchasing Silver plans, who have idiosyncratically higher surplus.

the annual Bronze, Gold, or Platinum premiums leads to a 34.3% to 38.4% reduction in the plan's market share among the lowest-income consumers. The semi-elasticity with respect to the the Silver plan premium is much lower at -15.6, which results from the popularity of Silver plans plans among these consumers. Households with incomes above 250 percent of FPL are less premium elastic, with semi-elasticities between -7.2 and -14.7.

Our estimates suggest that these consumers are highly premium elastic, similar to some estimates in the literature. Drake (2018) finds semi-elasticities that range from about -1 to -24 depending on the demographic group and Shepard (2016) finds an average semi-elasticity of -25, both in similar populations. Other recent work on the demand for insurance in this market finds average semi-elasticities that range from -1 to -10 (K. M. M. Ericson and Starc 2015; Saltzman 2019; Tebaldi 2020; Chan and Gruber 2010). One possibility is that consumers who use third party brokers to purchase health insurance are among the more premium elastic consumers in the market.

The semi-elasticities among insurance products reflect, in part, that consumers consider the different categories of insurance products to be close substitutes. Consumers are less likely to forgo insurance entirely. The semi-elasticity with respect to the insurance mandate penalty is -7.0, -1.8, and -4.2 among households earning less than 250 percent of FPL, 250 percent to 400 percent of FPL, and more than 400 percent of FPL respectively. This implies that a $100 increase in the annual penalty for being uninsured will decrease the uninsurance rate by 7.0% among households earning less than 250 percent of FPL and 2 to 4 percent among consumers earning more than 250 percent of FPL.

In Table 7, we present another view of consumer substitution toward being uninsured: the diversion ratios. This table displays the probability that a marginal consumer who leaves their current plan due to a premium increase will choose to become uninsured. The interquartile range of diversion ratios among all consumers is 13.0 to 30.7. The diversion ratios make clear that, while the lowest-income consumers are the most price elastic, consumers who earn less than 250 percent of FPL and consumers who earn up to 400 percent of FPL are similarly likely to forgo purchasing any insurance if they leave their current plan, with median diversion ratios of 22.3 and 23.4. Consumers who earn more than 400 percent of FPL are about half as likely to decide to be uninsured with a median diversion ratio of 11.7.

Our findings on the extensive premium elasticity of insurance are comparable to those in the literature. Finkelstein, Hendren, and Shepard (2019) exploit discontinuity in the premiums offered to low-income consumers (between 1.5 and 3 times FPL) in Massachusetts and find extensive margin semi-elasticities of -5 to -9 percent. Jaffe and Shepard (2017) look at the effects of implementing a mandate penalty in Massachusetts in 2008 on the same population and find an extensive margin semi-elasticity of -9.7. Hackmann, Kolstad, and Kowalski (2015) use a similar method to Jaffe and Shepard (2017), but focus on the higher-income, non-subsidy-eligible population, and find an elasticity of -2.5. These results mirror our finding that the elasticities are larger among low-income consumers than those with middle to high incomes.

Outside Massachusetts, other work has found the elasticity of insurance coverage to be

quite low. Our findings are within the range of those that use a similar demand estimation approach. Studies using choice data from the individual markets in California and Washington find semi-elasticities from -0.08 to -3.7 (Tebaldi 2020; Saltzman 2019). Others use variation in the mandate penalty itself to identify the elasticity. Frean, Gruber, and Sommers (2017) exploit geographic variation in the incidence of the mandate during the implementation of the ACA and find economically negligible effects. Lurie, Sacks, and Heim (2019) use tax data from a national sample and discontinuities in the mandate formula and find a semi-elasticity of -0.5 with respect to the statutory penalty. However, they find that there is substantial underpayment and that the semi-elasticity with respect to the actual mandate penalty paid is -1.6.

V. Conclusion

This paper analyzes consumer demand in a segment of the individual health insurance market that has previously been opaque: consumers who do not purchase insurance through the government-run web portals. These consumers tend to have higher incomes and constitute large segments of the market. We find that the own-premium semi-elasticities with respect to a $100 increase in the annual premium range between -7 and -38 across insurance products and income ranges. Low-income consumers are the most premium-elastic, and they also have the highest willingness to pay for additional insurance. Single households earning less than 250 percent of FPL are willing to pay from $164 to $177 per month, and non-single households from $194 to $231 per month, for an additional 10 percentage points of actuarial value across demographic groups. Higher-income single-person households earning more than 400 percent of FPL are willing to pay between $30 and $46 per month, and higher-income non-single households are willing to pay $82 to $135 per month.

Our estimates are consistent with the higher end of elasticity estimates in the literature (Drake 2019; Shepard 2016), which suggests that the consumers who use third-party platforms to select insurance plans may be among the more price-sensitive consumers. We also find that the willingness to pay for more generous insurance among higher-income consumers is relatively low compared to the lowest-income consumers. Both these features of consumer demand are of first-order concern as policy makers consider options to make health insurance more affordable and more valuable to consumers who are currently ineligible for premium subsidies.

The demand for health insurance on the extensive margin is relatively less price elastic, with semi-elasticities that range from -1.8 to -7.0 across income groups. Our estimates are consistent with the literature on insurance demand. Studies of higher-income populations across broad areas of the US find semi-elasticities from -1 to -4 (Tebaldi 2020; Sacks, Lurie, and Heim 2018; Hackmann, Kolstad, and Kowalski 2015; Saltzman 2019). Studies of low-income populations in Massachusetts have found much larger semi-elasticities, which supports our conclusion that income is an important determinant of demand (Finkelstein, Hendren, and Shepard 2019; Jaffe and Shepard 2017).

Our results suggest that the intensive margin may be the most important aspect of adverse selection. While consumers do not consider uninsurance to be a close substitute for insurance products, consumers are price elastic when choosing products within the market. This suggests that policies such as risk adjustment are likely critical in mitigating the effects of

adverse selection among insurance plans (Geruso et al. 2019).

REFERENCES

Abraham, Jean, Coleman Drake, Daniel W. Sacks, and Kosali Simon. 2017. "Demand for Health Insurance Marketplace Plans Was Highly Elastic in 2014-2015." *Economics Letters* 159: 69–73.

ASPE. 2016. "Health Insurance Marketplace 2016 Open Enrollment Period: Final Report." *Department of Health and Human Services, Office of the Assistant Secretary for Planning and Evaluation*, 1–48.

Atherly, Adam, and Bryan Dowd. 2009. "Should Healthy Medicare Beneficiaries Postpone Enrollment in Medicare Part D?" *Health Economics* 18 (8): 921–31. https://doi.org/10.1002/hec.1413.

Berry, Steven T. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *The RAND Journal of Economics* 25 (2): 242–62. https://doi.org/10.2307/2555829.

Chan, David, and Jonathan Gruber. 2010. "How Sensitive Are Low Income Families to Health Plan Prices?" *American Economic Review* 100 (2): 292–96. https://doi.org/10.1257/aer.100.2.292.

Claxton, Gary, Matthew Rae, Michelle Long, Nirmita Panchal, and Anthony Damico. 2015. "Kaiser 2015 Employer Health Benefits Survey." Menlo Park, CA.

Dafny, Leemore, Jonathan Gruber, and Christopher Ody. 2015. "More Insurers Lower Premiums: Evidence from Initial Pricing in the Health Insurance Marketplaces." *American Journal of Health Economics* 1 (1): 53–81.

Dartmouth Atlas. n.d. "Dartmouth Atlas of Health Care."

Drake, Coleman. 2018. "What Are Consumers Willing to Pay for a Broad Network Health Plan?:

Evidence from Covered California." *Becker Friedman Institute Working Paper* 2018 (23).

———. 2019. "What Are Consumers Willing to Pay for a Broad Network Health Plan?: Evidence from Covered California." *Journal of Health Economics* 65: 63–77.

Ericson, Keith M Marzilli, and Amanda Starc. 2015. "Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange." *Review of Economics and Statistics* 97 (3): 667–82.

Ericson, Keith Marzilli, and Judd B. Kessler. 2016. "The Articulation of Government Policy: Health Insurance Mandates versus Taxes." *Journal of Economic Behavior & Organization* 124 (April): 43–54. https://doi.org/10.1016/J.JEBO.2015.09.021.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams. 2016. "SOURCES OF GEOGRAPHIC VARIATION IN HEALTH CARE: EVIDENCE FROM PATIENT MIGRATION." *Quarterly Journal of Economics* 131 (4): 1681–1726. https://doi.org/10.1093/qje/qjw023.

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. https://doi.org/10.1257/aer.20171455.

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 (May): 72–86. https://doi.org/10.1016/j.jhealeco.2017.02.004.

GAO. 2012. "CHILDREN'S HEALTH INSURANCE: Opportunities Exist for Improved Access to Affordable." *Government Accountability Office,* 1–45.

Geruso, Michael, Timothy Layton, Grace McCormack, and Mark Shepard. 2019. "The Two Margin Problem in Insurance Markets." Cambridge, MA. https://doi.org/10.3386/w26288.

Hackmann, Martin B, Jonathan T Kolstad, and Amanda E Kowalski. 2015. “Adverse Selection and an Individual Mandate: When Theory Meets Practice.” *American Economic Review* 105 (3): 1030–66. https://doi.org/10.1257/aer.20130758.

Heiss, Florian, Daniel McFadden, and Joachim Winter. 2009. “Regulation of Private Health Insurance Markets: Lessons from Enrollment, Plan Type Choice, and Adverse Selection in Medicare Part D.” 15392. Cambridge, MA. https://doi.org/10.3386/w15392.

Jaffe, Sonia, and Mark Shepard. 2017. “Price-Linked Subsidies and Health Insurance Markups.” *National Bureau of Economic Research Working Paper Series* No. 23104. https://doi.org/10.3386/w23104.

Lurie, Ithai Z., Daniel W. Sacks, and Bradly Heim. 2019. “Does the Individual Mandate Affect Insurance Coverage? Evidence from Tax Returns.”

Nardi, Mariacristina De, Eric French, John B Jones, Jerome Adda, Kartik Athreya, Gadi Barlevy, Marco Bassetto, et al. 2010. “Why Do the Elderly Save? The Role of Medical Expenses.” *Source: Journal of Political Economy* 118 (1): 39–75. https://doi.org/10.1086/651674.

Panhans, Matthew. 2019. “Adverse Selection in ACA Exchange Markets: Evidence from Colorado.” *American Economic Journal: Applied Economics* 11 (2): 1–36. https://doi.org/10.1257/APP.20170117.

Planalp, Colin, Julie Sonier, and Brett Fried. 2015. “State-Level Trends in Employer-Sponsored Health Insurance: A State-by-State Analysis.” *Robert Wood Johnson Foundation*, 1–83.

Sacks, Daniel;, Ithai; Lurie, and Bradley Heim. 2018. “Does a Larger Individual Mandate Penalty Increase Insurance Coverage?” Atlanta.

Saltzman, Evan. 2019. “Demand for Health Insurance: Evidence from the California and Washington ACA Marketplaces.” *Journal of Health Economics* 63: 197–222.

Shepard, Mark. 2016. “Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange.” 22600. *NBER Worker Paper*. Vol. 000186. NBER Working Paper. https://doi.org/10.3386/w22600.

Tebaldi, Pietro. 2020. “Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA.”

Tables

Table 1. Comparing the Private Marketplace to Other Data Sources

	ACS	Healthcare.gov	Private Marketplace *Estimation Sample*	Private Marketplace *Re-weighted Sample*
*Maximum Household Age**				
Under 18	0%	9%	0%	0%
18 to 25	16%	11%	8%	12%
26 to 34	23%	17%	27%	27%
35 to 44	21%	17%	25%	19%
45 to 54	23%	21%	23%	22%
55 to 64	17%	24%	16%	20%
Household Income				
Less than 250% FPL	61%	75%	30%	75%
250% to 400% FPL	19%	16%	13%	16%
More than 400% FPL	20%	9%	57%	9%
Metal Level Selections				
Catastrophic		1%	3%	1%
Bronze		23%	38%	24%
Silver		67%	44%	68%
Gold		7%	12%	6%
Platinum		2%	3%	1%
Sample Size*	118,427	4,084,834	92,502	92,502

Notes: The private marketplace data is higher income and more concentrated in bronze plans than other data sources, but after controlling for income, plan choices are similar across data sources. The first column contains data from households in the American Community Survey (ACS) that do not get insurance through the government or an employer. The second column come from the public use files describing enrollment through the federal Healthcare.gov platform. The final two columns describe the data from a private insurance marketplace used in this paper. The first of these is the sample used for estimation, and the second is a sample that has been re-weighted to match the income distribution of the Healthcare.gov data.
* The maximum age distribution and the sample size refer to households, with the exception of the Healthcare.gov data where household information is unavailable. The statistics for Healthcare.gov are computed from individuals.

Table 2. Premiums and Market Shares Across Demographic Groups

	Bronze	Silver	Gold	Platinum
		Average Monthly Premium		
Maximum Household Age				
18 to 25	131	167	212	293
26 to 34	167	207	257	348
35 to 44	180	222	276	373
45 to 54	209	266	341	471
55 to 64	261	347	460	661
Household Income				
Less than 250% FPL	59	111	179	306
250% to 400% FPL	183	242	320	459
More than 400% FPL	260	308	370	475
		Share of Plan Choices		
Maximum Household Age				
18 to 25	31%	55%	11%	3%
26 to 34	36%	46%	14%	4%
35 to 44	41%	42%	13%	4%
45 to 54	42%	43%	11%	3%
55 to 64	42%	47%	9%	2%
Household Income				
Less than 250% FPL	19%	76%	5%	1%
250% to 400% FPL	43%	44%	10%	2%
More than 400% FPL	50%	29%	16%	5%

Notes: This table displays the average monthly premium and the share of plan choices for each metal level (excluding Catastrophic) by household age and household income. The average monthly premium is calculated as the average of the 25th percentile of premiums available to a particular household within a particular metal group, after adjustments to the premium due to age, income, and family size. This matches the methodology used in estimation, but the variation across groups and plans is similar if we use the median or the mean of premiums available to a particular household.

Table 3. Demand Estimates

	(1)		(2)		(3)		(4)	
	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
Household Premium	-1.016	0.051	-1.059	0.052	-1.111	0.053	-1.114	0.053
Interaction Terms								
Age 31 - 40	-0.066	0.058	-0.066	0.058	-0.117	0.060	-0.120	0.060
Age 41 - 50	-0.030	0.053	-0.019	0.053	-0.045	0.055	-0.048	0.055
Age 51 - 64	0.070	0.050	0.101	0.049	0.049	0.052	0.050	0.052
Family	0.187	0.022	0.181	0.022	0.188	0.023	0.183	0.023
250 - 400% FPL	0.619	0.036	0.584	0.036	0.767	0.038	0.759	0.038
More than 400% FPL	0.495	0.031	0.484	0.031	0.621	0.032	0.622	0.032
Deductible	-1.299	0.033	-1.394	0.033				
Interaction Terms								
Age 31 - 40	-0.006	0.036	-0.012	0.036				
Age 41 - 50	0.001	0.035	-0.006	0.034				
Age 51 - 64	-0.006	0.032	-0.007	0.032				
Family	-0.061	0.023	-0.091	0.024				
250 - 400% FPL	1.195	0.034	1.137	0.034				
More than 400% FPL	1.165	0.031	1.170	0.031				
Additional OOP Risk	-1.003	0.035	-1.084	0.035				
Interaction Terms								
Age 31 - 40	0.149	0.041	0.145	0.041				
Age 41 - 50	0.002	0.040	-0.002	0.040				
Age 51 - 64	-0.066	0.038	-0.068	0.038				
Family	-0.059	0.027	-0.080	0.027				
250 - 400% FPL	0.892	0.035	0.848	0.035				
More than 400% FPL	0.998	0.033	0.997	0.033				
Actuarial Value					21.702	0.564	21.978	0.568
Interaction Terms								
Age 31 - 40					0.695	0.601	0.721	0.600
Age 41 - 50					0.420	0.581	0.491	0.581
Age 51 - 64					0.699	0.548	0.669	0.549
Family					1.277	0.392	1.729	0.398
250 - 400% FPL					-19.873	0.576	-19.601	0.578
More than 400% FPL					-19.555	0.515	-20.187	0.519
σ	0.616	0.010	0.606	0.011	0.621	0.010	0.614	0.010
Fixed Effects								
Age, Family, Income	Y		Y		Y		Y	
Market	Y				Y			
Product Category	Y				Y			
Market-Category			Y				Y	
N	62,384		62,384		62,384		62,384	

Notes: This table displays the results from the demand estimation. All dollar values used in estimation are denominated in thousands, including annual base premium, annual household premium, deductible, and additional out-of-pocket (OOP) risk. The number of observations is equal to the product of the number of demographic groups (15,460) and the size of the choice set (4.04 choices on average). Each panel of the table displays the base value for a given characteristic followed by the demographic interaction terms. This same set of demographic variables make up the age, family, and income fixed effects.

Table 4. First Stage of Demand Estimation

	Log Market Share
Log Leave-Out Share	0.831
	(0.007)
Constant	Y
N	62,384
F-statistic	12,550
R^2	0.1679

Notes: This table displays the results from the first stage of demand estimation. The instrument is used to identify the nesting parameter in the demand specification.

Table 5. Willingness to Pay for Additional Insurance

	Household Income (% of FPL)		
	< 250%	250% - 400%	>400% FPL
Single Households			
Age 18 - 30	164	56	30
Age 31 - 40	153	54	34
Age 41 - 50	161	59	35
Age 51 - 64	177	83	46
Non Single Households			
Age 18 - 30	212	199	95
Age 31 - 40	194	138	82
Age 41 - 50	206	175	94
Age 51 - 64	231	326	135

Notes: Higher-income consumers have a lower willingness to pay for additional insurance than low-income consumers. This table displays consumer willingness to pay for a 10 percentage point increase in the actuarial value of an insurance plan in dollars per month. The results are computed from specification 4 in Table 3 by dividing the demographic specific coefficient for actuarial value by the demographic specific coefficient for price.

Table 6. Semi-Elasticity By Product Category and Income Levels

	Annual Price					
	Earning Less Than 250% of FPL					
	Catastrophic	Bronze	Silver	Gold	Platinum	Man. Penalty
Market Share						
Catastrophic	**-12.2**	1.6	7.7	0.4	0.1	2.5
Bronze	0.0	**-34.3**	24.2	1.4	0.3	8.3
Silver	0.0	5.5	**-15.6**	1.4	0.3	8.3
Gold	0.0	5.5	24.2	**-38.4**	0.3	8.3
Platinum	0.0	5.0	23.0	1.3	**-37.4**	8.1
Uninsured	0.0	1.3	5.4	0.3	0.1	**-7.0**
	Earning Between 250% and 400% of FPL					
	Catastrophic	Bronze	Silver	Gold	Platinum	Man. Penalty
Market Share						
Catastrophic	**-1.1**	0.3	0.4	0.1	0.0	0.3
Bronze	0.0	**-7.6**	4.0	0.9	0.2	2.5
Silver	0.0	3.6	**-7.2**	0.9	0.2	2.5
Gold	0.0	3.6	4.0	**-10.3**	0.2	2.5
Platinum	0.0	3.4	3.7	0.9	**-10.4**	2.4
Uninsured	0.0	0.8	0.8	0.2	0.0	**-1.8**
	Earning More Than 400% of FPL					
	Catastrophic	Bronze	Silver	Gold	Platinum	Man. Penalty
Market Share						
Catastrophic	**-1.5**	0.5	0.4	0.2	0.1	0.3
Bronze	0.4	**-9.2**	4.0	2.1	0.6	2.0
Silver	0.4	7.0	**-12.1**	2.1	0.6	2.0
Gold	0.4	7.0	4.0	**-14.1**	0.6	2.0
Platinum	0.4	6.5	3.8	2.0	**-14.7**	1.9
Uninsured	0.1	2.1	1.2	0.6	0.2	**-4.2**

Notes: Higher-income consumers are less price-elastic than low-income consumers. This table displays average own-price and cross-price semi-elasticities with respect to the annual premium. Each elasticity represents the percent change in market share (rows) in response to a $100 increase in the annual premium (column). The last column indicates the percent change in market share with respect to a $100 increase in the individual mandate penalty.

Table 7. Diversion to Uninsured

	By Income		
	25th Percentile	Median	75th Percentile
All	13.0	20.4	30.7
< 250% FPL	16.0	22.3	32.7
250 - 400% FPL	16.8	23.4	34.4
>400% FPL	5.9	11.7	22.1
	By Product		
	25th Percentile	Median	75th Percentile
Catastrophic	11.4	17.0	24.4
Bronze	15.0	23.6	38.6
Silver	15.3	32.2	100.0
Gold	11.5	17.3	24.0
Platinum	12.0	17.2	23.0

Notes: Higher-income consumers are less likely than low-income consumers to decide to be uninsured when switching insurance products. This table displays the distribution of diversion ratios towards uninsurance. Each value represents the probability that a marginal consumer will become uninsured, rather than purchase another product. The top panel displays the distribution diversion ratios among consumers within each income category and across all products. In the bottom panel, the distributions are conditional on consumers purchasing a product of a particular metal level.

Figures

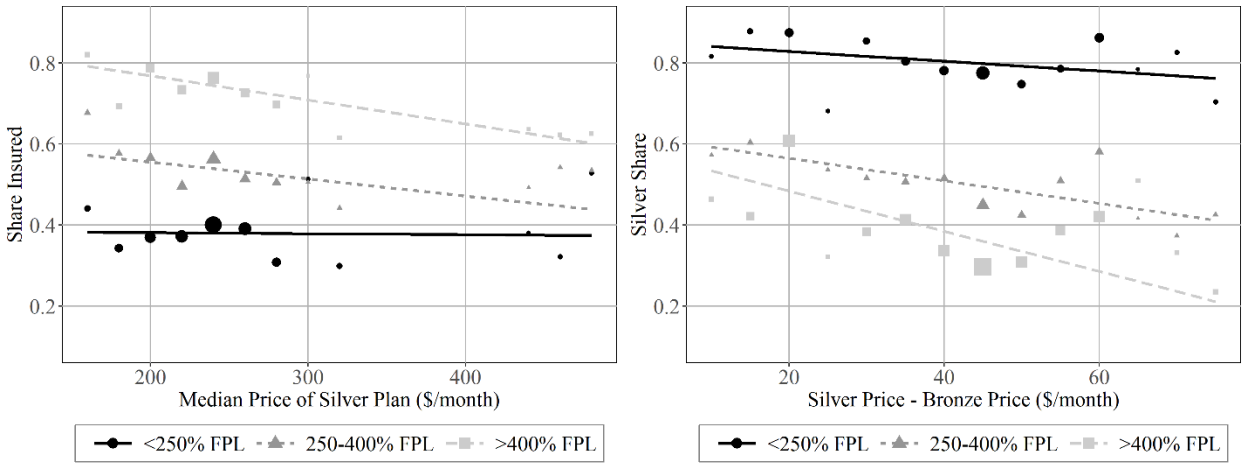

(a) (b)

Figure 1: Demand and Price

Note: This figure displays the raw relationship between the quantity sold of insurance products and the price of insurance. The left panel shows the share of individuals that purchase insurance with respect to the median price of a silver plan for those households. The right panel shows the share of households that purchase a Silver plan among households that purchase either Silver or Bronze relative to the difference between the median Silver plan premium and the median Bronze plan premium. Source: ACS, Private Online Marketplace, HIX 2.0

Figure Information:

Figure 1a: A data point on this graph represents the average share of consumers that purchase insurance given the median price for a Silver plan in their local market. Light gray squares and the long-dash line represent households that earn more than 400 percent of FPL. Dark gray triangles and the short-dashed line represent households that earn between 250 and 400 percent of FPL. Black circles and the solid line represent households that earn less than 250 percent of FPL. The size of the squares, triangles, and circles is proportional to the number of consumers in the sample that make up that

average, and the best fit lines are the weighted according to these sample sizes.

Figure 1b: A data point on this graph represents the average share of consumers that purchase a Silver plan among all consumer that purchase either a Silver plan or Bronze plan, given the difference between the median Silver plan price and the median Bronze plan price in their local market. Light gray squares and the long-dash line represent households that earn more than 400 percent of FPL. Dark gray triangles and the short-dashed line represent households that earn between 250 and 400 percent of FPL. Black circles and the solid line represent households that earn less than 250 percent of FPL. The size of the squares, triangles, and circles is proportional to the number of consumers in the sample that make up that average, and the best fit lines are the weighted according to these sample sizes.