The Effect of an Insurance Mandate: Evidence from an Online Exchange

By CONOR RYAN, ROGER FELDMAN, AND STEVE PARENTE*

We use a novel data set from a private online marketplace to estimate the demand for individual health insurance and the efficacy of the individual mandate across 126 geographic markets. We find that the own-price semi-elasticities for an insurance product range between -16 and -22, and the semi-elasticity of insurance coverage with respect to the mandate penalty is -1.5. Using a stylized model of supply, we find that the individual mandate modestly increased the share of insured individuals in 2015 by 2.9 percentage points – from 46.6 percent to 49.5 percent.

* Ryan: University of Minnesota, 1925 S 4th Street, 3-163, Minneapolis, Minnesota, 55454 (email: ryan0463@umn.ed); Feldman: 15-210 Phillips-Wangensteen Building, Minneapolis, MN, 55454 (email: feldm002@umn.edu); Parente: 321 19th Street South, 3-122, Minneapolis, Minnesota, 55454 (email: stparente@umn.edu); We would like to thank seminar participants at the American Society of Health Economists and the Congressional Budget Office for helpful comments.

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 prices 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, referred to as the individual mandate. In addition to directly assisting with insurance affordability, these policies—in particular the mandate penalty—are intended 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).

The necessity and efficacy of the subsidy and penalty design depends on the price elasticity of demand for health insurance, as well as the supply response to shifting demand. Since the ACA was implemented, the number of uninsured has declined, and insurance has become much more affordable for households that qualify for government assistance. However, the coverage rate among the 40 million individuals in the non-group market has reached a plateau of only about 50 percent. This raises the question of whether the policies intended to incentivize coverage are insufficient or ill-targeted.

In this paper, we estimate the demand for health insurance and the efficacy of the individual mandate using a novel data set that provides evidence on the elasticities of consumer demand across a broad sample of the individual market. A large and growing literature estimates discrete choice demand models for individual health insurance (Ericson and Starc 2015; Shepard 2016; Jaffe and Shepard 2017; Tebaldi 2017; Saltzman 2017; 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 states provide opportunities to study many aspects of market design, they may not be the ideal setting for answering questions about the national response to market regulation. The rates of insurance, management of uncompensated care, and institutional efforts to support and advocate for state-run health insurance marketplaces in Massachusetts and California are not typical of many other states.

Our data, which come from a large private online insurance broker, have two advantages over the sources used in the existing literature. First, they cover a broad geographic area that complements existing insurance demand studies. Our estimation sample has 126 geographic markets in 18 states. Second, our data contain an equal distribution of low-income individuals who are eligible for subsidies and high-income individuals who are not. Since government-run exchanges are the primary source of subsidies and the target of enrollment efforts, most exchange consumers are low-income—98 percent earn less than 400% of the federal poverty level (ASPE 2016)—and a substantial amount of previous work has found that income is important in explaining health insurance choices (e.g., Tebaldi 2016; Saltzman 2017). This same income group represents only 69% of the national market.[1] Roughly half of the households in our data earn more than 400% of the federal poverty level.

As is typically the case with choice data, our data represent individual choices on the intensive margin: the insurance product a consumer selects, given that some product is chosen. We use uninsurance rates from the American Community Survey (ACS), matched to the choice data through geography and demographics, to identify the extensive margin of insurance choice—whether or not to purchase any insurance (Tebaldi 2017; Saltzman 2019). In Section II.B, we explain how we constructed households and determined program eligibility from the ACS for our study of health insurance purchasing decisions.

We complement our demand estimates with a stylized model of the supply of health insurance that incorporates estimates of adverse selection: the relationship between the quantity of insurance sold and the average cost of insurance. We use adverse selection estimates from Hackmann, Kolstad, and Kowalski (2015), who exploit the 2006 Massachusetts health reform to determine the effect of the insurance coverage boost on insurance costs, and Panhans (2017), who derives similar cost effects of the ACA insurance coverage expansion in Colorado. For robustness, we also use several calibrated average cost curves from the Medical Expenditure Panel Survey.

Our primary contribution is to a growing literature on the effects of the individual mandate and similar penalties. The data and methods used to estimate the effects of a national mandate are varied, and the resulting estimated impacts range from negligible (Frean, Gruber, and Sommers 2017) to essential (Ericson and Starc 2015). We find that the individual mandate increased the share of insured individuals in 2015 by 2.9 percentage points – from 46.6 percent to 49.5 percent. The increase in insured individuals is concentrated among higher-income enrollees who are not eligible for price-linked subsidies. In this population, the mandate increased the share of insured individuals from about 62 percent to 70 percent. Our demand estimates imply that the semi-elasticity of insurance coverage with respect to a $100 increase in the mandate penalty is -1.5, i.e. a $100 increase in the annual penalty for going without insurance

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

leads to an increase of 1.5 percent in the share of insured. This estimate is on the low end of extensive margin elasticities for health insurance in the literature, which range from -1 to -10. Many of the estimates on the high end of this range are based on low-income households in Massachusetts (Finkelstein, Hendren, and Shepard 2019; Jaffe and Shepard 2017), while our estimates are more similar to those in higher-income populations and outside Massachusetts (Tebaldi 2017; Sacks, Lurie, and Heim 2018; Hackmann, Kolstad, and Kowalski 2015; Saltzman 2019).

We also contribute to a broad literature on the demand for individual market health insurance. A large portion of the recent literature uses variation induced by the ACA premium regulations to identify the price elasticity of demand in this market (Tebaldi 2017; Drake 2019; Saltzman 2019). We compare this methodology to an alternative instrumental variable approach that uses a measure of costs in out-of-state products to instrument for the base ACA premium. We find that the extensive margin elasticities are similar in the two methods, but that the instrumental variable approach yields slightly higher elasticities among insurance products. Using variation in the ACA premium formula, we find the own-price semi-elasticities range from -11 to -22 across Bronze, Silver, and Gold plans. With the instrumental variable approach, own-price semi-elasticities range from -11 to -26.

Section I describes the institutional details of the post-ACA individual health insurance landscape. In Section II, we describe the third-party data. In Section III, we outline the demand model. In Section IV, we present the results of the demand estimation. Section V presents a stylized model of adverse selection to simulate the equilibrium effects of the individual mandate.

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 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.

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. In most states, insurance companies 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 declines to 0.765 for children under 14. 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 a premium subsidy if they earn between 100 and 400 percent of the federal poverty level (FPL), an income threshold that accounts for household size. 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. 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 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 family'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.[2] 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.

[2] In many states, this is Healthcare.gov. However, a number of states run their own insurance portal, e.g. CoveredCA.com in California.

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 online marketplace. In addition to “on-exchange” plans, the online marketplace offers some “off-exchange” plans as well. For insurance firms that offer some plan types on and off the exchange, this distinction is minimal.[3] 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 exchange.[4] The data contain the choices of subsidized and unsubsidized consumers in 48 states.

We do not investigate why an individual may decide to enroll via a third-party broker instead of the government-run websites, which are the most popular mechanism for enrollment since 2014. One important dimension of selection is income. Consumers who are eligible to receive subsidies may be more likely to enroll via the government websites, where many of the advertisements that market the availability of subsidies direct consumers. However, nearly half of the households in our data are subsidy-eligible.

In Table 1, we summarize the private online market and compare it to two other references on the individual insurance market: the 2015 American Community Survey (ACS); and the Office of the Assistant Secretary for Planning and Evaluation at the US Department of Health and Human Services (APSE). The ACS offers the broadest depiction of the health insurance market across all purchasing platforms. ASPE publishes detailed descriptive statistics for health insurance purchases through HealthCare.gov, which allows us to make plan choice comparisons with our data. Compared with the ACS, HealthCare.gov enrollment is weighted heavily toward low-income, subsidy-eligible consumers. As a result, the plan type market shares reported by ASPE are weighted heavily toward Silver plans, which have extra cost-sharing benefits at low

[3] 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.
[4] The government-run marketplaces allowed some third-party marketplaces to interface with the back-end of the government enrollment process in order to administer subsidies to their consumers.

incomes. The private online market data represent disproportionately higher-income and younger enrollees. The last panel shows plan type market shares conditioned on earning at least 400% of FPL. The choices are roughly similar with higher enrollment in Bronze plans through our data.

The data contain information on the consumer's age, the first three digits of the consumer's zip code, the plan purchased by the consumer, and the subsidy received. The data are detailed, but because of plan choice aggregation (explained in Section II.C), the only information we use from the data is the metal level of the selected plan. 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. We assume that the smoking status refers only to the head of household. We also assume that any household that contains more than one individual contains two adults of the same age, and any additional persons are children under the age of 14.[5]

We observe or can compute the income of nearly every individual who receives a subsidy. However, we do not know 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

[5] We observe prices paid for a subset of observations and can impute the household-level age rating using the median firm metal-level base premium for the purchased product. The correlation between our simple age rating rule and the age rating in the imputed sample is 0.89. We experimented with more detailed rating imputations based on the observed sample of prices paid, and found it does not significantly affect our results.

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 to ensure that the sample of choices is not limited to only a portion of the market. We discuss the the choice sets in Section II.C.

After dropping additional observations because of missing data, the remaining data set includes roughly 88,000 individual and family health insurance choices across 18 states and 126 rating areas.[6]

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 using the 2015 American Community Survey (ACS). Since we only observe a sample of product choices, we use the ACS to construct market shares for uninsurance among households with similar observable characteristics. 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 or not the household is eligible for a subsidy, and whether the household has one, two, or at least 3 members.

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

[6] 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.

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. This approach abstracts from demand differences across insurance firms, which is not the focus of our paper.

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 regions, 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 3 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 family 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 the household is a linear function of the base premium of the insurance plan, $p_{ijr} = a_i p_{jr} - b_i$, 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:

$$u_{ijr} = \gamma' Z_i + \alpha_i p_{ijr} + \beta' X_{jr} + \xi_{jr} + \eta_{ijr}$$

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

where γ captures how the mean value of insurance depends on household demographics, β represents preferences for observed product attributes, and α_i 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. There is some evidence that individuals are more sensitive to a mandate because of a preference for compliance (Saltzman 2019), but individuals may appear to be less sensitive to the statutory level of the mandate penalty because they pay less than the full value of 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 estimate two specifications of demand with two approaches to obtain unbiased estimates of α, the parameter governing price elasticity.

First, we use an instrumental variable approach that exploits market shares across all consumers who face the same choice set. In this approach, we are able to group consumers into large enough groups that all product categories have a non-zero market share. Because some products are offered only to some consumers, we divide markets into demographic buckets (θ), which consist of sets of consumers facing the same choice set in a particular rating area. We divide households into two income buckets according to whether households are eligible for cost-sharing reductions: those earning 2.5 times the federal poverty limit (FPL) or less, and those earning more.[7] Since only individuals under the age of 30 can purchase catastrophic plans, we also divide households based on whether the head of household is less than or equal to 30 years old. We estimate the demand for insurance products as a function of only the base premium, but

[7] The cost-sharing reductions vary in degree by 50 percentage point increments for households earning between 100 and 250 percent of FPL, and our results are not substantially different if we further divide income into each step of cost-sharing reduction. We use the broader income specification for better comparison to household-level estimates.

we allow the premium sensitivity to vary with the demographic characteristics of the market. From Berry (1994), we can write the nested logit model as a linear equation:

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

where $S_{\theta jr}$ is the market share of product j in market r among demographic group θ and $p_{\theta jr}$ is the average premium of the product, after subsidies and age-rating, in market r. The vector of demographic characteristics $Z_{\theta r}$ includes the categorical age and income variables that we use to define θ, and we let the price elasticity parameter α depend on these same demographics. We assume that the unobserved product qualities ξ are normally distributed and possibly correlated with base premiums.

Identification requires two instruments: one to address the endogeneity of price and another to address the endogeneity of the nesting term, $\log\left(\frac{S_{\theta jr}}{1-S_{\theta 0r}}\right)$. We use a measure of "leave-out cost" to identify α (Panhans 2017). From rate filing data, we calculate the average cost for a particular metal level in every state other than the state in which the product is located. We interact the leave-out cost instrument with the demographic characteristics in $Z_{\theta r}$ to account for price differences across demographic groups. To identify σ, we use 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 specification exploits variation in the data that identifies the nesting parameter: the share of consumers with insurance in each demographic group and the price of insurance. However, the model has several potential weaknesses. First, the instrument may not be exogenous. In health insurance markets with adverse selection, the cost of providing insurance is positively correlated with the preference for insurance. For this reason, we use the "leave-out cost" instead of the within-market cost of each insurance product. However, there may still be concern that variation in cost across products in other states is determined by similar demand behavior. Second, the price faced by consumers is a non-linear function of age, income, and household size. Our aggregate specification, which relies on an average across large demographic groups and a linear first stage instrument, may not capture the richness of premium responses that result from the non-linear insurance pricing function.

An alternative method of identifying the premium elasticity is to use the non-linear price function itself as a source of variation in price uncorrelated with the unobserved quality of insurance (Tebaldi 2017; Drake 2019; Saltzman 2019). To exploit this variation, we aggregate consumers into smaller 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).[8] We define the rating area as the market. Our estimation equation for this household-level specification is slightly modified:

$$\log(S_{\theta jr}) - \log(S_{\theta 0r}) = \gamma' Z_{\theta r} + \alpha_{\theta} p_{\theta jr} + \beta' X_{jr} + \xi_{jr} + \nu_{\theta jr} + \sigma \log\left(\frac{S_{\theta jr}}{1 - S_{\theta 0r}}\right) \qquad (2)$$

where S_{jr} is the market share of product j in market r among demographic group θ and $p_{\theta jr}$ is the average premium of the product for consumers in demographic group θ in market r. The vector of demographic characteristics $Z_{r\theta}$ includes 4 age categories (18-30, 31-40, 41-50, and 51-64), three income categories (less than 200% of FPL, 200% to 400% of FPL, and more than 400% of FPL), and whether the household is a single consumer or a family. We let the price elasticity parameter α depend on these same demographics.

In the household-level specification, the unobserved quality of each insurance product is given by $\xi_{jr} + \nu_{\theta jr}$, where ξ_{jr} is the average unobserved product quality in a rating area and $\nu_{\theta jr}$ is a normally distributed and mean 0 deviation in product quality for each demographic bucket θ. We group the metal levels into two product categories (high and low) and control for ξ_{jr} using fixed effects for each market-product category. The "high" category includes all metal levels that cover at least 80 percent of expected expenditures (Platinum, Gold, and two reduced cost-sharing variants of Silver); all other metal levels are in the "low" category.

The key identifying assumption is that $p_{\theta jr}$ and $\nu_{\theta jr}$ are uncorrelated. Importantly, the regulated price formula addresses both the cost of providing insurance and affordability but not

[8] 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.

demand. Were prices set solely to address the cost of insurance, adverse selection between demand and cost might cause an endogeneity problem. However, the age component of the premium does not fully capture average costs by age (see Appendix A.2 for more discussion) the income component creates price variation unrelated to costs. This identification argument is outlined in detail in Tebaldi (2017).

In the household-level specification, we maintain the same instrumental variable for the nesting parameter. We construct the measure of "leave-out share" separately for each demographic group θ in a manner analogous to the instrumental variable specification.

IV. Estimation Results

In Table 2, we present the results of our demand estimation. The results for each model are consistent across fixed-effect specifications and across our identification strategies for price elasticity, which suggests that price elasticities are not particularly sensitive to the choice of specification. The results for the first stage of the instrumental variable regressions are presented in the Table 3.

Our estimates of the nesting parameter (σ) are consistently close to 0.9 across all specifications, which indicates that insurance products are much closer substitutes for each other than for uninsurance. Demographics are an important determinant of price sensitivity, which decreases with age and family size. In the IV specifications, we find that low-income consumers are slightly more price sensitive than high-income consumers. In contrast, the household level results find that higher-income enrollees are significantly more price sensitive. However, the result of the IV specification are less precise, possibly because of the measurement error in the average premium. In our preferred specification (HH-2), the youngest consumers are nearly twice as price sensitive as the oldest consumers. Families with middle-aged heads of household are about half as price sensitive as single middle-aged adults, and subsidy-ineligible households are between 1.5 to 2 times as price sensitive as the lowest-income households.

In Table 4, we present the implied semi-elasticities of demand for the instrumental variable specification (Panel A) and the micro specification (Panel B). 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.

In both specifications, we find very similar and relatively little substitution toward the outside good. The mean semi-elasticity of uninsured with respect to a $100 price increase in the mandate penalty is similar in both specifications: -1.4 in the instrumental variable approach and -1.5 in the household-level estimates. A semi-elasticity of -1.5 implies that the fraction of uninsured would decline by 1.5 percent in the event of a $100 increase in the mandate penalty for all households. Similarly, the cross-elasticities imply that a $100 increase in the price of each metal level results in less than a 1 percent change in the fraction of insured. The instrumental variable specification leads to higher estimates of between-insurance elasticities in some product categories, but very similar estimates of extensive margin elasticities. For all remaining results in this section and Section V, we use the most detailed household-level specification (HH-2), which allows for more detailed household-level heterogeneity.

Our findings are comparable to those in the literature, with the exception of studies of the low-income population in Massachusetts. 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. These consumers may be considerably more price-elastic for two reasons. First, low-income consumers who face credit constraints may be more likely to forgo insurance out of necessity when faced with premium increases. 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. Second, Massachusetts has a system for providing care for individuals without insurance that may lead consumers to consider uninsurance to be a closer substitute for insurance than in other areas of the country.

Outside Massachusetts, our semi-elasticity estimates are in the range reported in the literature. 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. Studies using choice data from the individual markets in California and Washington find semi-elasticities from -0.08 to -3.7 (Tebaldi 2017; Saltzman 2019).

In Table 5, we present another view of substitution toward the outside good: the diversion ratio. This table displays the probability that a marginal consumer who leaves their current plan when price increases will choose to become uninsured. Our estimates imply that the interquartile range of diversion ratios among all consumers is 4.1 percent to 10.0 percent. Marginal consumers who receive subsidies are more likely to switch to uninsurance in the event of a price increase (Panel A), with the median diversion ratio of 7.2 for subsidized consumers and 3.8 for unsubsidized consumers. The median diversion ratios for consumers in each product category (Panel B) are between 5.5 percent and 11.2 percent. The tail of large diversion ratios for consumers enrolled in Silver plans is driven by consumers who are offered very generous silver plans for very low prices, but without other affordable options. These consumers are almost certain to become uninsured if they leave the Silver plan. This substitution pattern also exists to a lesser extent for Bronze plans among higher-income consumers who are eligible for subsidies but not cost-sharing reductions.

V. Policy Analysis

The policy objective of an individual mandate is not only to encourage participation in the insurance market, but to improve affordability. Because adverse selection creates a link between the quantity of insurance sold and the price of insurance, a mandate can reduce insurance prices through broad participation in the market. In extreme cases of adverse selection, the market could completely unravel in the absence of a strong insurance mandate. Additionally, the mandate is not the only tool to guard against adverse selection. The price-linked subsidies shield subsidized consumers from higher price levels across the entire market, which provides another backstop against unravelling.

To interpret the effect of the penalty for remaining uninsured, we must combine our demand results with estimates of adverse selection. In this section, we present a stylized model of supply that incorporates estimates from the literature on the relationship between quantity sold and average costs in the individual insurance market to determine the effect of the individual mandate.

A. Supply Model

We assume the insurance market is perfectly competitive and premiums are risk-adjusted. In each region r, the prices of J_r products are set so that average revenue is equal to average covered costs, and every insurance product is priced as if the average total cost of its enrolled pool is identical to that of the entire market.

A detailed exposition of the model and definition of equilibrium is presented in Appendix A. More briefly, an equilibrium set of prices, $p = \{p_j\}_{j=1}^{J_r}$, solves two equations:

$$p_{jr} = P(I_r, A_{jr}; AV_{jr}) \tag{1}$$

$$I_r = \sum_{j \in J_r} S_{jr}\,(p) \tag{2}$$

Equation (1) captures how the price of each insurance plan depends on the insurance rate in a market (I_r), the age distribution enrolled in that insurance plan $\left(A_{jr}\right)$, and the actuarial value of the plan (AV_{jr}). In the presence of adverse selection, the price of insurance is negatively correlated with the insurance rate $\left({}^{\partial P}/_{\partial I} < 0\right)$, and the strength of this relationship governs the degree to which an enrollment boost encouraged by the individual mandate translates into lower prices for consumers.

Because we lack data to estimate the relation between the size of a market and the average cost of the insured, we use two estimated log-linear curves from the literature on the degree of adverse selection in the individual market: Panhans (2017) and Hackmann, Kolstad, and Kowalski (HKK, 2015). These authors estimate (around observed equilibrium) the slope of a log-linear cost curve with respect to insurance enrollment, and the normalized curves are pictured in Figure 1.[9]

We supplement these estimates with selection scenarios estimated from the Medical Expenditure Panel Survey (MEPS), which follow from a stylized thought experiment. Suppose that the marginal consumer— the first consumer to leave given an increase in price—is ex-post the cheapest consumer. Then, if the insurance rate is 50 percent, the average cost of the market is given by the average cost of the most costly 50 percent of consumers in the market. This is an extreme case of selection which we can estimate from MEPS by ranking consumers by total ex-post cost. Because this scenario is unrealistic, we model a selection curve with a slope that is 5 percent of the extreme ordering (low) and 30 percent of the extreme ordering (high). For robustness, we also present the results for 100 percent of the extreme ordering.

The individual market in the rating areas that we study is predominantly low-income. While this population makes up roughly half of our data, they comprise more than 80 percent of the population, according to the ACS. To account for this discrepancy, we re-weight our data according to the weights in the ACS for the same demographic categories used to match the rates of insurance (see Section II.B). After re-weighting, the estimated model precisely matches the uninsurance rate in the data when the structural demand errors are incorporated into the prediction.

[9] See Appendix Section A.2 for more detail on how we compute these selection curves and incorporate the estimates into the price functions.

B. Equilibrium Analysis

Table 5 shows equilibrium results in the absence of the individual mandate using six cost scenarios, ordered by the degree of adverse selection—i.e. the elasticity of average cost with respect to enrollment. The first scenario has no adverse selection, and accounts only for the demand effect.[10] The next four scenarios include a range of reasonable adverse selection curves. In the absence of the individual mandate in 2015, we find that the share of insured would be between 2.8 and 2.9 percentage points lower. The decline in the share of insured is concentrated among the consumers who are not eligible for subsidies. The fraction insured would decline by between 6.8 and 7.8 percentage points in this population. As an upper bound, we also include the most extreme adverse ordering from MEPS, which would imply a 3.3 percentage point increase in the total uninsured rate and an 8.7 percentage point increase in the uninsured rate among the subsidy-ineligible in the absence of the mandate. We exclude the 100% MEPS scenario in our discussion of the range of effects.

The effects of the individual mandate are mitigated by the presence of price-linked subsidies. The average premium paid among the subsidy-eligible either remains constant or declines in the absence of the individual mandate, despite modest increases in the average total premium. This is due to an increase in the value of the subsidy in response to higher overall premiums (Saltzman 2019). Marginal consumers also tend to have higher incomes and face higher premiums, even among the subsidized population. This effect is evident in the demand response, which predicts a lower average premium paid in the absence of a mandate despite an increase in the average total premium and no response in the supply of insurance.

The effects on average premiums paid among the subsidy-ineligible are small but non-negligible – from a 0.2 percent increase to a 1.3 percent decline. Since the subsidized population comprises a large portion of the market—83% in our sample—the overall supply response is small. Importantly, we model perfectly competitive firms that do not internalize the effect of the price-linked subsidy policy on consumer demand. If firms have sufficient market power, this may have an effect on the price response (Jaffe and Shepard 2017).

[10] We do allow for cost differences coming from the changing distribution of age, but these have a negligible effect.

In the final column, we present the effect on average premiums holding the distribution of enrollees fixed to the baseline distribution. This removes any composition changes or substitution effects and reflects only the implied price increase facing consumers for their preferred plan when the individual mandate is in place. These price increases range from 0.4 percent to 2.3 percent in the presence of adverse selection.

In our most plausible scenarios using estimates from the literature, total premium increases in the absence of the mandate are small. Prior to the ACA implementation, Eibner and Price (2012) predicted that premiums would increase between 10 and 25 percent if the mandate were repealed. More recently, the Congressional Budget Office (2017) estimated that premiums would increase by 10 percent. Other studies suggest that the market would completely unravel in the absence of a mandate (Handel, Hendel, and Whinston 2015; Azevedo and Gottlieb 2017; Ericson and Starc 2015).

One reason that our estimates of premium increases are small is that we do not allow for changes in overall insurance participation to affect sorting between plans. Geruso et al. (2019) show that an individual mandate may worsen adverse selection on the intensive margin. Using estimates from Finkelstein, Hendren, and Shepard (2019) and Hackmann, Kolstad, and Kowalski (2015), Geruso et al. (2019) find that a similar reduction in the magnitude of a mandate penalty would lead to premium increases of 5 to 8 percent.

VI. Conclusion

This paper analyzed consumer demand in a segment of the 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 found that the own-price semi-elasticities of insurance products with respect to a $100 increase in the annual premium range between -11 and -22.

However, the demand for health insurance on the extensive margin is relatively priceinelastic, with a mean semi-elasticity of -1.5. Our estimates are consistent with the literature on the relatively high-income population and the population outside Massachusetts, which find semi-elasticities from -1 to -4 (Tebaldi 2017; 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 suggest that income and the institutional environment are important (Finkelstein, Hendren, and Shepard 2019; Jaffe and Shepard 2017).

To determine the effect of the individual mandate penalty, we presented a stylized model of supply that incorporates estimates from the literature on adverse selection. We found that the individual mandate led to a 2.9 percentage point increase—from roughly 46 percent to 49 percent—in the fraction of insured in the individual market in 2015. Under plausible estimates of adverse selection, insurance premiums would have been between 0.4 and 2.3 percent higher in the absence of the mandate.

Our results, together with the expanding body of literature on the individual mandate, suggest that the individual mandate penalty was not critical to the existence of the individual market for health insurance for large segments of the population, even after the regulatory policies of the ACA increased traditional channels of adverse selection. This is likely for several reasons. First, the ACA subsidies provide premium assistance that greatly exceeds the mandate penalty for low-income households—the most price-elastic with respect to coverage (Finkelstein, Hendren, and Shepard 2019; Saltzman 2019).

Second, the extensive margin may not be the most important margin of selection. In this paper, we abstracted from selection among the available insurance plans, and the ways in which a mandate penalty interacts with selection among insurance products (Geruso et al. 2019). 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 like risk adjustment are likely more critical than the mandate penalty in mitigating the effects of adverse selection.

REFERENCES

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.
Azevedo, Eduardo M., and Daniel Gottlieb. 2017. “Perfect Competition in Markets With Adverse Selection.” *Econometrica* 85 (1): 67–105. https://doi.org/10.3982/ECTA13434.
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.
Claxton, Gary, Matthew Rae, Michelle Long, Nirmita Panchal, and Anthony Damico. 2015. “Kaiser 2015 Employer Health Benefits Survey.” Menlo Park, CA.
Congressional Budget Office. 2017. “Repealing the Individual Health Insurance Mandate: An Updated Estimate.” Vol. 2027.
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.
Eibner, Christine, and Carter C Price. 2012. “The Effect of the Affordable Care Act on Enrollment and Premiums, With and Without the Individual Mandate.” *Rand Health Quarterly*. Vol. 2. http://www.ncbi.nlm.nih.gov/pubmed/28083224%0Ahttp://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4945284.
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. http://www.nber.org/papers/w18089.
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. https://www.gao.gov/assets/600/591797.pdf.
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.
Handel, Ben, Igal Hendel, and Michael D Whinston. 2015. “Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk.” *Econometrica* 83 (4): 1261–1313.
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.

Kautter, John, Gregory C. Pope, and Patricia Keenan. 2014. “Affordable Care Act Risk Adjustment: Overview, Context, and Challenges.” *Medicare & Medicaid Research Review* 4 (3): 1–11. https://doi.org/10.5600/mmrr.004.03.a02.
Lurie, Ithai Z., Daniel W. Sacks, and Bradly Heim. 2019. “Does the Individual Mandate Affect Insurance Coverage? Evidence from Tax Returns.” https://dansacks.files.wordpress.com/2019/09/hslmandate.pdf.
Panhans, Matthew T. 2017. “Adverse Selection in ACA Exchange Markets: Evidence from Colorado.” https://doi.org/10.2139/ssrn.2920049.
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. https://www.rwjf.org/en/library/research/2015/01/state-level-trends-in-employer-sponsored-health-insurance.html.
Sacks, Daniel;, Ithai; Lurie, and Bradley Heim. 2018. “Does a Larger Individual Mandate Penalty Increase Insurance Coverage?” Atlanta. https://ashecon.confex.com/ashecon/2018/webprogram/Paper5585.html.
Saltzman, Evan. 2017. “Demand for Health Insurance: Evidence from the California and Washington ACA Marketplaces.” *Working Paper*, 1–46. https://faculty.wharton.upenn.edu/wp-content/uploads/2017/02/ACA-Demand-Paper-2-20-2017.pdf.
———. 2019. “Demand for Health Insurance: Evidence from the California and Washington ACA Exchanges.” *Journal of Health Economics*. https://doi.org/10.1016/j.jhealeco.2018.11.004.
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. 2017. “Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA.” 05. *Becker Friedman Institute for Research in Economics Working Paper*. Becker Friedman Institute for Research in Economics Working Paper. Chicago. http://www.pietrotebaldi.com/uploads/8/1/3/6/81364040/eehix_aug14_17.pdf.

Tables

Table 1. Data Description

	Private Marketplace	ACS	ASPE
		Age Distribution	
Under 18	0%	1%	8%
18 to 26	18%	19%	11%
26 to 34	24%	17%	17%
35 to 44	20%	17%	17%
45 to 54	20%	20%	22%
55 to 64	18%	27%	25%
		Income Distribution	
Under 200% FPL	28%	36%	68%
200% to 400% FPL	22%	33%	31%
Over 400% FPL	51%	31%	2%
		Metal Level Market Shares	
Catastrophic	5%		1%
Bronze	36%		22%
Silver	46%		67%
Gold	10%		7%
Platinum	3%		3%
		Metal Level Market Shares (Income > 400% FPL)	
Catastrophic	8%		7%
Bronze	47%		35%
Silver	26%		32%
Gold	15%		19%
Platinum	5%		8%

Notes: This table displays demographic and choice distributions in each data source. The American Community Survey numbers come from heads of household that are insured through the individual insurance market. The ASPE numbers come from the 2015 Open Enrollment Report for enrollments through HealthCare.gov. The age numbers are not adjusted for head of household. Source: ACS 2015, ASPE 2015

Table 2. Demand Estimates

	IV Specification		HH-level Specification	
	(IV-1)	(IV-2)	(HH-1)	(HH-2)
Price Coefficients				
Base Premium	-0.437	-0.449		
	(0.073)	(0.074)		
Age 31 - 64	0.196	0.197		
	(0.033)	(0.033)		
Greater than 250% FPL	0.089	0.092		
	(0.043)	(0.043)		
Household Premium			-0.245	-0.275
			(0.026)	(0.026)
Age 31 - 40			0.013	0.017
			(0.028)	(0.027)
Age 41 - 50			0.071	0.079
			(0.025)	(0.025)
Age 51 - 64			0.175	0.192
			(0.024)	(0.024)
Family			0.118	0.122
			(0.010)	(0.010)
250 - 400% FPL			-0.103	-0.101
			(0.015)	(0.015)
Greater than 400% FPL			-0.184	-0.186
			(0.014)	(0.014)
Deductible	-0.043	-0.037	-0.059	-0.059
	(0.017)	(0.018)	(0.009)	(0.009)
OOP Max	-0.048	-0.041	-0.310	-0.369
	(0.012)	(0.015)	(0.009)	(0.009)
σ	0.900	0.907	0.889	0.878
	(0.028)	(0.029)	(0.009)	(0.009)
Fixed Effects				
Age, Income, Family	Y	Y	Y	Y
Market	Y	Y	Y	
Product Category		Y	Y	
Market-Category				Y
N	2,718	2,718	64,384	64,384
R^2	0.718	0.718	0.384	0.391

Notes: This table displays the results from demand estimation. All dollar values used in estimation are denominated in thousands, including annual base premium, annual household premium, deductible, and out-of-pocket maximum. The displayed R-squared is unadjusted. Standard errors are displayed in parenthesis.

Table 3. First Stage of Demand Estimation

	Instrumental Variable		Household-Level
	Inside Share	Base Premium	Inside Share
Log Leave-Out Share	1.15		0.89
	(0.021)		(0.006)
Leave-Out Cost		0.343	
		(0.017)	
Age 31 - 64		0.284	
		(0.020)	
Greater than 250% FPL		0.154	
		(0.020)	
Constant	Y	Y	Y
State Fixed Effect		Y	
Age, Income FE		Y	
N	2,718	2,718	64,384
R^2	0.695	0.779	0.282

Notes: This table displays the results from the first stage of demand estimation. The instrumental variable estimates are used in specifications IV-1 and IV-2 in Table 2. The household-level estimates are used in specifications HH-1 and HH-2 in Table 2. The R-squared is unadjusted. Standard errors are in parenthesis.

Table 4. Semi-Elasticity Across Metal Levels

Panel A: Instrumental Variable Specification (IV-2)						
	Catastrophic	Bronze	Silver	Gold	Platinum	Mandate Penalty
Catastrophic	**-1.9**	0.7	0.8	0.3	0.1	0.1
Bronze	0.2	**-22.7**	18.7	1.9	0.5	1.4
Silver	0.2	6.7	**-10.8**	1.9	0.5	1.4
Gold	0.2	6.3	17.2	**-25.5**	0.5	1.3
Platinum	0.2	4.9	12.7	1.5	**-20.4**	1.1
Uninsured	**0.0**	**0.3**	**0.9**	**0.1**	**0.0**	**-1.4**

Panel B: Household-level Specification (HH-2)						
	Catastrophic	Bronze	Silver	Gold	Platinum	Mandate Penalty
Catastrophic	**-7.1**	1.4	4.7	0.4	0.1	0.5
Bronze	0.3	**-16.6**	12.5	2.0	0.5	1.4
Silver	0.3	7.0	**-11.1**	2.0	0.5	1.4
Gold	0.3	7.0	12.5	**-21.6**	0.5	1.4
Platinum	0.3	6.4	11.8	1.8	**-21.7**	1.3
Uninsured	**0.0**	**0.5**	**0.8**	**0.1**	**0.0**	**-1.5**

Notes: This table displays average own-price and cross-price semi-elasticities. Each elasticity represents the percent change in market share with respect to a $100 increase in the annual premium. The last column indicates the percent change in market share with respect to a $100 increase in the individual mandate penalty.

Table 5. Diversion to Uninsured

Panel A	By Income		
	25^{th} Percentile	Median	75^{th} Percentile
All	4.1	6.5	10.0
Subsidized	5.2	7.2	11.3
Unsubsidized	1.9	3.8	7.1
Panel B	**By Product**		
	25^{th} Percentile	Median	75^{th} Percentile
Catastrophic	3.6	5.5	7.9
Bronze	4.8	7.6	14.3
Silver	4.9	11.2	93.5
Gold	3.7	5.5	7.6
Platinum	3.8	5.5	7.3

Notes: 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. In Panel A, it displays the distribution diversion ratios among consumers within each income category and across all products. In Panel B, the distributions are conditional on consumers purchasing a product of a particular metal level.

Table 6. Equilibrium Effects of the Mandate Penalty

Selection Scenario	Fraction Insured (%) Total	Fraction Insured (%) Subsidy Eligible	Fraction Insured (%) Subsidy Inelgible	Average Premium Paid Subsidy Eligible	Average Premium Paid Subsidy Inelgible	Average Total Premium Equilibrium Selection	Average Total Premium Baseline Selection
Baseline	49.5	45.1	70.3	135	324	328	328
Demand Effect	46.6	43.1	63.1	134	325	330	328
5% MEPS	46.6	43.1	63.0	134	325	331	330
HKK	46.6	43.1	62.9	134	327	333	332
Panhans	46.5	43.1	62.7	134	328	334	333
30% MEPS	46.5	43.1	62.5	134	329	336	335
100% MEPS	46.2	43.1	61.6	135	336	345	346

Notes: This table displays the effects repealing the mandate penalty under different scenarios of adverse selection. All premium values represent average per-person monthly household premiums. Fraction uninsured is displayed in percent. Baseline values incorporate the structural demand errors and pricesely match data values.

Figures

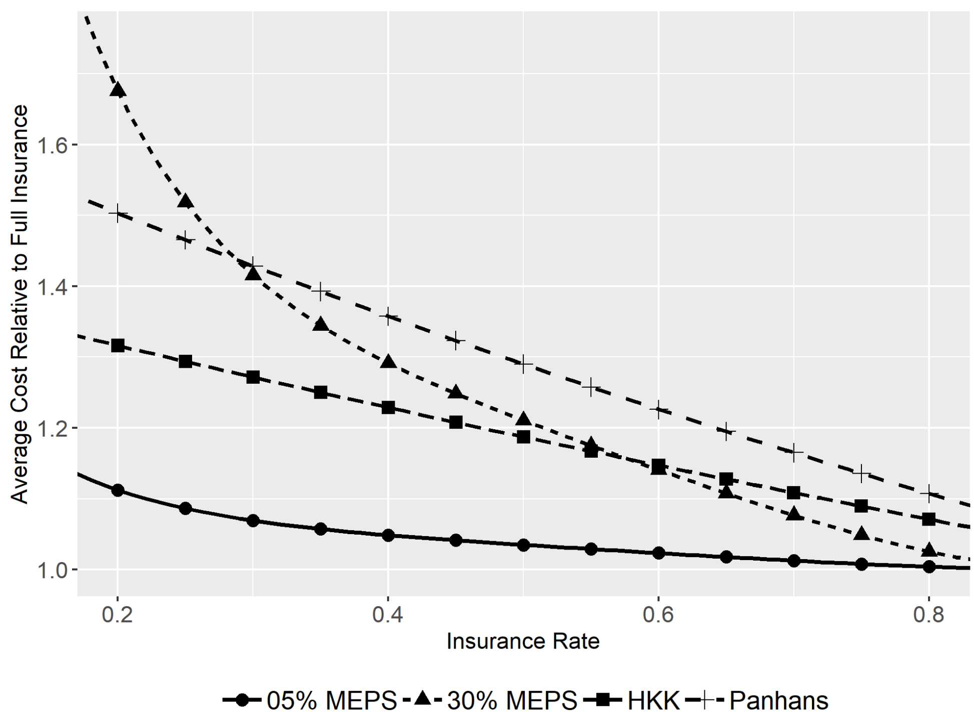

Figure 1: Normalized Cost Functions

Note: This figure displays calibrated average cost curves. Each point represents the ratio of average cost given a particular insurance rate to the average cost when the insurance rate is 1.0. Source: MEPS 2015, Hackmann, Kolstad, and Kowalski (2015), and Panhans (2017)

Appendix A

A.1 Supply Model

A consumer of observable type θ in region r has an expected total cost given by $c_{\theta r}$. In the case of adverse selection, where cost and demand are positively correlated, the marginal consumer is likely to have a lower cost than the average consumer, conditional on observable characteristics. To account for this possibility, we specify cost as:

$$c_{\theta r} = a_{\theta}^{c} C_r(I_r)$$

where a_{θ}^{c} is an age-cost factor that reflects the average cost of consumers of a given age, and $C_r(\cdot)$ is a region-specific cost function that depends on the insurance rate of the market, I_r.

Each region r contains J_r insurance products which are offered by a perfectly competitive and risk-adjusted market. Prices are set so revenue is equal to cost, and every insurance product is prices as if the average cost of its enrolled pool is identical to that of the entire market.

A plan j covers a percentage of costs equal to the regulated actuarial value of the plan, AV_j. For example, a Silver plan will cover 70 percent of total costs, and a gold plan will cover 80 percent.[11] In addition to actuarial value, plan cost also varies by a factor, γ_j, which captures other non-risk related factors that may lead to differences in cost across plans, such as medical utilization differences that result from lower effective prices. We assume that the premium differences do not reflect differences in underlying average cost (perfect risk adjustment), but they may reflect differences in the age-mix of the enrolled population, following the risk adjustment formula of the ACA (Kautter, Pope, and Keenan 2014).

An equilibrium in a region r is defined as a set of prices, $\boldsymbol{p_r} = \{p_j\}_{j=1}^{J_r}$, such that average plan revenue is equal to average plan cost:

[11] Cost sharing reductions are offered to low-income purchasers of silver plans. However, these reductions were subsidized by the government in 2015, and thus did not affect plan costs. While these subsidies are not currently in place, we maintain this assumption to isolate the effect of the mandate.

$$p_j \bar{a}_j = AV_j \gamma_j \bar{a}_j^c C_r(I_r) \tag{1}$$

$$I_r = \sum_{j \in J_r} S_{jr}\,(\boldsymbol{p_r}) \tag{2}$$

where $\bar{a}_j$ is the average age rating of a plan j and $\bar{a}_j^c$ is the average age-cost factor of plan j. Both $\bar{a}_j$ and $\bar{a}_j^c$ depend on the full vector of prices ($\boldsymbol{p_r}$). In Section V, we describe the price function $P(\cdot)$ which can be written as

$$P\big(I_r, A_{jr}; AV_{jr}\big) = \frac{V_j \gamma_j \bar{a}_j^c C_r(I_r)}{\bar{a}_j}.$$

A.2 Model Calibration

We estimate the cost function using data from the 2015 Medical Expenditure Panel Survey (MEPS) and estimates of adverse selection from the literature. First, we estimate the age-cost rating (a_θ^c) from the average total expenditures of all insured individuals by age across all payers. Figure A.1 shows the estimated age-cost curve (solid) as well as the regulated age-rating price curve under the ACA (dotted). We use a fourth-order polynomial of average total spending by age, and then we normalize relative to the average spending of a 21 year-old.

We specify the region-specific cost function as:

$$C_r(I_r) = b_r g(I_r)$$

where b_r is a base cost factor that varies by region and $g(\cdot)$ is a normalized cost index where $g(1.0) = 1$.

Because we lack data to estimate the relation between the size of a market and the average cost of the insured, we construct two indices using estimated log-linear curves from the literature on the degree of adverse selection in the individual market: Panhans (2017) and

Hackmann, Kolstad, and Kowalski (HKK, 2015). These authors estimate (around observed equilibrium) the slope of a log-linear cost curve with respect to insurance enrollment.

HKK exploit the coverage expansion during the Massachusetts state health reform in 2006 through 2008. They find that a coverage expansion that increased the fraction insured from 0.70 to 0.97 led to an 8.7 percent decline in the average cost of insurance. We use this finding to find the slope of the average cost curve with respect to the insurance rate. Panhans uses geographic boundary discontinuities to estimate the effect of premium changes in the average available Silver plan on the local insurance rate and average Silver plan cost. We use combine these two elasticities to compute the implied slope of the average cost curve with respect to the insurance rate. The formulas for computing each slope are below, and more details on the origin of each estimate can be found in the respective papers. In both cases, we compute the normalized cost index as $g(I) = C(I)/C(1.0)$.

$$\frac{\partial \log(C)}{\partial I} = \left(\frac{\log(5{,}270) - \log(4811)}{.703 - .968}\right) \qquad \text{(HKK)}$$

$$\frac{\partial \log(C)}{\partial I} = \frac{\partial \log(C) \big/ \partial \log(p)}{I \left(\partial \log(I) \big/ d\log(p)\right)} = \frac{0.747}{0.57 * -2.572} \qquad \text{(Panhans)}$$

To compute the alternative selection scenarios from MEPS, we compute the average cost of the population that exceed a point in the distribution in health expenditures. For example, we compute $C(0.2)$ as the average total expenditures of all enrollees that exceed the 20^{th} percentile of the expenditure distribution. We hold fixed the age distribution by computing the relevant expenditure percentile separately for 5-year age brackets. We fit the implied average cost estimates with a fourth order polynomial of the expenditure percentile, which we then interpret as the insurance rate. In order to create more reasonable scenarios that represent 5% and 30% of this extreme case of selection, we modify the predicted average cost ratios by the respective percent reductions. If $\hat{g}^{100}(I)$ is the full estimated, normalized curve with $\hat{g}^{100}(1.0)$ normalized to 1, then

$$\hat{g}^{x}(I) = 1 + (\hat{g}^{100}(I) - 1) * \frac{x}{100}$$

With the cost function estimates and the age characteristics of enrollment simulated from our demand estimation, we calibrate γ_j and b_r so the observed data constitute an equilibrium. Because there is one free parameter, we normalize γ_j to equal 1 for the Bronze plan in each market. With this normalization, the calibration is exactly identified. In solving for equilibrium, we also adjust the price of the benchmark Silver plan to allow for subsidies to adjust with the price level.

Appendix Figures

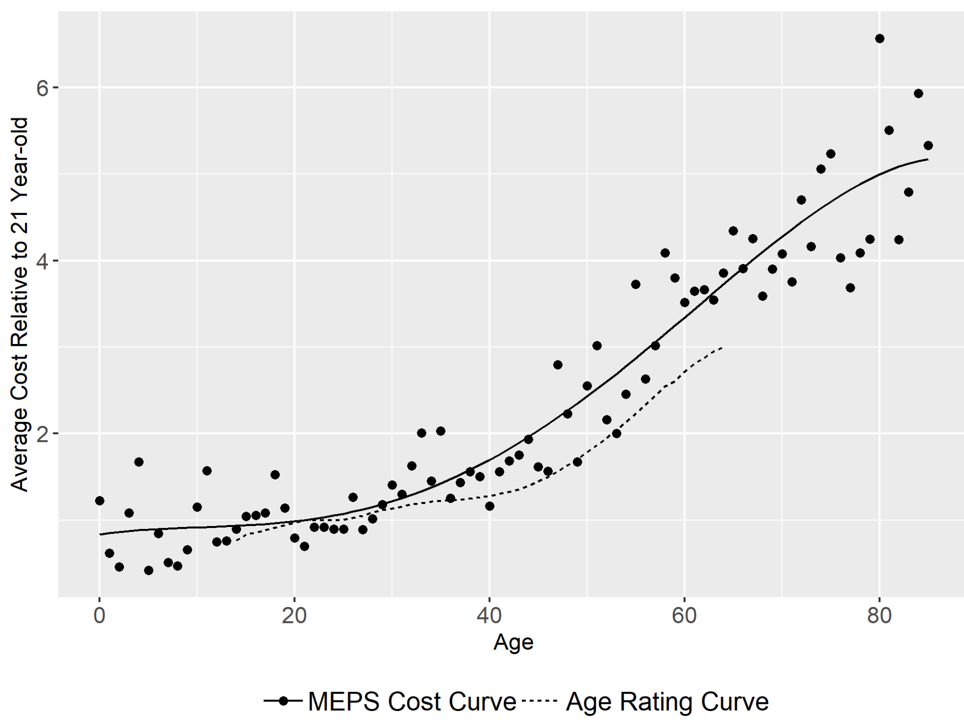

Figure A.1: Age – Cost Curve

Note: This figure displays the relationship between age and average spending. Each point represents the ratio of average spending given a particular age relative to the average spending of a 21 year old. The age rating curve implemented by the ACA is also displayed for reference. Source: MEPS 2015.