How Do the Mandate and the Premium Tax Credit Affect the Individual Insurance Market?*

Daniel W. Sacks

May 3, 2017

Abstract

The Affordable Care Act's (ACA) community rating and guaranteed issue provisions require that insurers sell insurance at the same price to all enrollees, regardless of health. To limit adverse selection concerns, the ACA also mandates coverage, to make coverage affordable, a it subsidizes insurance with a premium tax credit which is more generous for low-income people To understand the effects of the mandate and premium tax credit, I develop and estimate a model of the individual insurance under the Affordable Care Act. The model allows for adverse selection through an industry cost curve that depend on price: as prices rise, healthy people exit the market, driving up average costs. I estimate the demand and cost curves using exogenous variation in the effective price of insurance, which arises from differential eligibility for the premium tax credit. I find modestly elastic demand and an upward sloping cost curve, pointing to adverse selection. I use the demand and cost curves to simulate the effect of repealing the mandate and premium tax credit, or replacing them with the age based tax credits of the American health Care Act. Repealing the mandate and premium tax credits without replacement would raise prices, reduce coverage, and lower welfare. Replacing them with age based tax credits would largely undo these effects.

JEL codes: D12, G22
Key words: Health insurance, adverse selection, mandate

*Indiana University, Kelley School of Business, email: dansacks@indiana.edu. I am grateful to Ray Kuntz of the AHRQ/CFACT data center for assistance with the data, and to Alex Gelber, Rick Harbaugh, Martin Hackmann, Brad Heim, Sean Lyons, Kosali Simon, Jeff Prince, Amanda Starc, Matthijs Wildenbeest, and seminar participants at Indiana University, the University of Wisconsin, the IU-VU-UL miniconference on health economics, and the IIOC for helpful comments. The research in this paper was conducted at the CFACT Data Center, and the support of AHRQ is acknowledged. The results and conclusions in this paper are those of the author and do not indicate concurrence by AHRQ or the Department of Health and Human Services. Data Center staff reviewed the results to avoid disclosing confidential information. This study was designated exempt by the Indiana University IRB under protocol #1609359160. Jordan Keener provided excellent research assistance.

1 Introduction

A major goal of the Affordable Care Act (ACA) was to reform the individual insurance market, so that people without public or employer-sponsored health insurance could obtain coverage. Before the ACA went into effect in 2014, many feared that the individual insurance market was adversely selected. Adverse selection arises when individuals have private information about their expected medical costs, but insurers cannot price on this information. A large theoretical literature has demonstrated that, with adverse selction, a competitive equilibrium may not exist, and when it does, it is inefficient: prices are too high, too few people have insurance, and people with insurance have too little coverage (e.g. Akerlof (1970); Rothschild and Stiglitz (1976); Hendren (2013); Handel et al. (2015); Azevedo and Gottlieb (2017)). To reform the individual insurance market, the ACA included three key provisions. First, community rating and guaranteed issue require insurers to offer the same price regardless of health, and to offer coverage to all who demand it. Second, the individual mandate penalizes people who do not purchase insurance. Third, the premium tax credit subsidizes insurance, making it affordable for low-income Americans. The architects of the law view these provision as a "three-legged stool:" without any of them, the individual insurance market might fail (Gruber, 2010). The mandate, however, is among the least popular parts of the ACA, and the recently proposed American Health Care Act would repeal the mandate and replace the premium tax credit, which is tied to income and premiums, with an age-based tax credit, which is independent of premiums and only loosely related to premiums.

In this paper, I develop and estimate an equilibrium model of supply and demand in the individual insurance market. I use the model to simulate the effects of repealing the mandate or the premium tax credit, or replacing them with age-based tax credits. These simulations are valuable for understanding the possible consequences of alternative repeal or replacement plans, not only of the American Health Care Act but more generally of plans that eliminate the mandate or change the structure of premium subsidies. Fully understanding these consequences requires an equilibrium model because the mandate and the premium subsidies affect insurance coverage decisions not only directly, as implicit subsidies, but indirectly, through their effect on selection and equilibrium prices.

I use the model of the insurance market developed by Einav et al. (2010), but extended to incorporate the pricing and subsidy structure of the individual insurance. Insurers set a single, market-wide price for insurance. People self-select into insurance on the basis of expected costs and benefits, including any tax credits for purchase and penalties for non-purchase. If individuals' expected health care costs are correlated with their willingness to pay for insurance, then at any price, the average cost of the insured is higher than the overall population average cost, and the market is adversely selected. As the price rises, the average cost of the insured also rises. Thus the industry average cost curve summarizes adverse selection. With adverse selection, many people would like to buy insurance at their person-specific actuarially fair price, but not at the market price, leading to a welfare loss. In this case, a mandate or subsidy can raise welfare by shifting up the demand curve, drawing healthy people into the market, and lowering prices. The effect of a mandate or a subsidy on prices, coverage, and welfare can be simulated given estimates only of demand and cost curves. Because social surplus is the area between the demand and cost curves, these objects are also sufficient statistics for welfare under alternative policies.

I estimate demand and cost curves using data from the Medical Expenditure Panel Survey (MEPS), matched to a hand-collected data set on insurance premiums. The MEPS is well-suited to this exercise because it is a nationally representative sample, designed to measure health care spending; it contains information on family structure, insurance coverage, medical spending, and health.[1] To identify the demand and cost curves, I require exogenous variation in premiums, i.e. variation unrelated to underlying demand or costs. I rely on differential exposure to the premium tax credit. This credit went into effect in 2014, but it is available only to families with income between 100 and 400% of the federal poverty line, and only in counties with relatively high premiums. My identification strategy relies on the differential change in insurance coverage from 2013 to 2014, for people eligible for the subsidies, relative to ineligible people, and adjusting for differential trends by income level. Validation tests and placebo specifications support this strategy.

I find that the demand for insurance is modestly elastic. A $1,000 increase in annual premiums reduces the insurance coverage rate by about 6.2 percentage points. This implies an average elasticity of about -0.3. Because I identify this effect from differential exposure to the premium tax credit, this elasticity also demonstrates that the credit has been effective in expanding insurance coverage. The elasticity estimate is towards the lower end of previous estimates of the elasticity of demand for insurance, which range from -0.2 to -0.9 (for example, Cutler and Reber (1998); Royalty and Solomon (1999); Einav et al. (2010); Ericson and Starc (2015)). However, most estimates of the demand for health insurance focus on the intensive margin of plan generosity, rather than

[1] In Appendix C, I replicate the main results from the MEPS data using the CPS, which is larger in size but lacks health care spending information, and has incomplete geographic identifiers. The results are nonetheless quite similar.

the extensive margin of whether to have any coverage at all, which I estimate. Extensive margin elasticities are likely to be smaller than intensive margin elasticities, as no coverage and coverage are not close substitutes.

To test for adverse selection, I estimate the average cost curve. Adverse selection requires that, as the premium rises, so too does the average cost of the people continuing to buy insurance. This is exactly what I find: a $1,000 increase in premiums causes the expected costs of the insured to rise by about $440, or put differently, a price increase that causes a 1 percent decrease in the insurance rate raises the average costs of the insured by about $70. The premium tax credit draws relatively healthy, low cost people into the market, bringing down average costs.

Because the market is adversely selected, it is possible that the mandate and the tax credits can improve welfare To measure the consequences of the mandate and the tax credits, I simulate equilibrium under alternative policies. Relative to 2017 baseline, eliminating the mandate would reduce coverage by about 6 percentage points. This overall effect includes both a direct and an equilibrium effect. The direct effect is that, at fixed prices, repealing the mandate reduces demand. The equilibrium effect is that when the mandate is repealed, premiums rise as healthy people exit the market, driving up prices and further reducing coverage. This equilibrium effect accounts for about a 25% larger than the direct effect alone, as prices rise by about 14% after repeal. Overall the mandate is welfare improving; repealing it reduces welfare by about $230 per capita. Interestingly, the premium tax credit mitigates some of the consequences of repealing the mandate. Because the tax credit is linked to prices, it rises when the mandate is removed, which eases the upward pressure on prices from adverse selection. Eliminating the premium tax credit instead of the mandate has similar, but larger effects on prices and coverage. Eliminating both the mandate and the premium tax credits has a large effect on the market: coverage falls by about a third, prices rise by about a third, and welfare welfare falls by $420 per capita.

As a final simulation, I find equilibrium if the mandate and premium tax credits are replaced by age-based tax credits as recently proposed in the American Health Care Act. These age based tax credits are, on average, generous, and so they are an effective replacement for the mandate and premium tax credit. Equilibrium coverage, premiums, and welfare would all be essentially unchanged if the mandate and premium tax credit were replaced by the AHCA's age-based tax credits.

These simulations, of course depend on the modelling assumptions and empirical approach. Two caveats in particular apply. First, the counterfactuals assume perfect competition. I show in a robustness test, however, that assuming the polar opposite—a monopolist insurer—produces fairly similar results for prices, coverage, and welfare. Second, the repeal or replace simulations hold fixed many aspects of the ACA. In particular, all the estimates and counterfactuals assume continued Medicaid expansion, as well as guaranteed issue and community rating; the simulations therefore do not give the full effect of the proposed American Health Care Act.

These results contribute to the large empirical literature on adverse selection in insurance markets. Much of this literature has studied selection in employer sponsored insurance (e.g. Carlin and Town (2008); Einav et al. (2010); Bundorf et al. (2012); Handel (2013)). Some studies papers have detected adverse selection in the individual insurance market (e.g. Buchmueller and DiNardo (2002), Simon (2005), and Sasso and Lurie (2009)) prior to the Affordable Care Act. It is unclear, however, how selection has changed with the ACA's many reforms.

This paper is most closely related to recent papers on the impact of a mandate or tax credits on coverage and welfare in the individual insurance market, Frean et al. (2017) and Hackman et al. (2015).[2] Frean et al. (2017) estimate the effect of the premium tax credit subsidies and mandate penalties on coverage, post-ACA. Their identification strategy is similar to my strategy for estimating price sensitivity. They find that the premium tax credit has an important effect on insurance coverage but, surprisingly, the mandate does not. I differ from their approach in two ways. First, standard theory implies suggests that people respond symmetrically to a mandate penalty or to a tax subsidy. I impose this constraint, but Frean et al. do not. Because coverage responds to the premium tax credit, I find in simulations that repealing the mandate also affects coverage and prices.[3] A second important difference is that Frean et al.'s estimates of the effect of the mandate and the premium tax credit take the price of coverage as given, but I simulate the equilibrium effects of these policies, allowing them to affect prices. This distinction is important, as the equilibrium effects are roughly 25 percent larger than the direct effects.

Hackman et al. (2015) study the role of the mandate in the Massachusetts reform. Using the framework of Einav et al. (2010). They find that the mandate raises coverage and welfare. I build on their work in two ways. First, I document adverse selection in the post-ACA insurance market,

[2] A broader literature covers the effects of health insurance reform on the individual insurance market. Kowalski (2014) looks at the effects of state-specific implementation decisions, such as Medicaid expansions or state-specific Marketplaces. Ericson and Starc (2015), Tebaldi (2016), and Jaffe and Shepard (2017) study the interaction between imperfect competition and pricing distortions induced by regulation—limits on age based pricing, or the price-linked subsidies implicit in the premium tax credit. These papers do not estimate the effects of the mandate or premium tax credit on coverage, prices, or welfare.

[3] Unlike Frean et al., I do not attempt to estimate directly the effect of the mandate on coverage. This is a challenging task because because the mandate penalty is nearly linear in income, making it difficult to cleanly identify its effect.

for the nation as a whole. Second, Hackman et al. (2015) do not study the role of premium tax credits, because they focus on the unsubsidized population in Massachusetts. I estimate the effect of both the premium tax credit and the mandate, and I simulate an alternative, plausible replacement policy. It is important to account for the premium tax credits because the tax credits mitigate the effect of repealing the mandate, and because they are important in their own right; removing them but keeping the mandate has a large effect on coverage, prices, and welfare.

2 Background on the Individual Insurance Market

The Patient Protection and Affordable Care Act (the ACA) made major changes to both health care delivery and health insurance provision in the United States. It expanded Medicaid, mandated that employers provide insurance and individuals obtain it, and changed the pricing regulations in the individual insurance market. Here I describe the most important aspects of these reforms for the individual insurance market. All of the provisions described below went into effect January 1, 2014.

2.1 Market establishment and price regulations

The ACA created centralized marketplaces known as the Health Insurance Marketplaces (sometimes called exchanges). In principle each state was free to establish its own exchange, but many states use `healthcare.gov`, including 28 states that use Federally Facilitated Marketplaces, letting the federal government perform all Marketplace functions. The exchanges differ from the traditional individual insurance market because all products and prices are displayed in a single place. Exchange plans are grouped into different metal levels, reflecting generosity, from bronze to platinum, with silver (the second least generous) by far the most popular tier. Insurers are allowed to continue selling products off-exchange, principally through brokers. ACA regulations require that insurers charge the same price off-exchange and on-exchange. Insurers selling on the exchange must pay a fee to do so, but enrollees buying on the exchange are eligible for premium subsidies (the premium tax credit, described below). Many Americans obtained health insurance through the Marketplaces. During the open enrollment period in 2014, roughly 8 million people signed up for a plan; sign-ups increased to 12 million in 2015, 13 million in 2016 (Uberoi et al., 2016).

Under the ACA, prices are community rated, and adjusted only for age and tobacco use. This means that in each rating area—typically a set of counties—each insurer chooses a single price, p^b.

This is the posted price for a single 21 year-old for health insurance, provided she does not use tobacco, and gross of any subsidies. I refer to this price as the base price. For other ages, prices are set according to an age curve, which is a collection of factors f_a that scale the base price, so a single, non-tobacco user of age a faces a price of $f_a p^b$. The law requires a 3:1 age band, meaning that $f_{64}/f_{21} \leq 3$. States may permit a tobacco surcharge t of up to 50%, so a tobacco users of age a pays $(1+t)f_a p^b$. Children younger than 21 have a single factor f_{child}. Typically households buy health insurance together, and there are some economies of scale from doing so. Family prices are determined by adding up the prices that individual family members face, but children beyond the third do not contribute to family premiums. A family with two parents, aged 35 and 40, and three or more children would face a price of $P = p^b(f_{35} + f_{40} + 3f_{child})$, ignoring tobacco use.[4] Thus insurers can price only on age, tobacco use, and location. To prevent insurers from only accepting patients with low expected costs, the ACA also mandates "guaranteed issue", which means that insurers must issue insurance coverage to anyone wishing to buy it at the posted price.

2.2 Premium stabilization programs

To help combat adverse selection, and to assist the market in coming to equilibrium in its first few years, the ACA created three "premium stabilization" programs, sometimes called the "three R's:" risk adjustment, risk corridors, and reinsurance. Risk adjustment is permanent, but the risk corridor and reinsurance program were temporary, and in effect only in 2014-2016.[5]

The risk adjustment program transfers money among insurers in the individual market, based on the observed riskiness of their enrollees. Each enrollee is assigned a risk score, equal to the ratio of their expected spending to average spending. Expected spending depends on coded diagnoses, such as for diabetes or hypertension. Insurers with below average risk scores—with healthy enrollees—make a transfer to insurers with above average risk scores. These transfers resemble those used in Medicare Advantage, where they reduce cream skimming, i.e. competition among insurers to enroll patients with low costs relative to payment (Brown et al., 2014; Curto et al., 2015). In the individual insurance market (in contrast to Medicare Advantage), the risk adjustment payments mechanically sum to zero. As a results, while risk adjustment discourages plans from cream skimming against each other, it cannot compensate insurers if the entire pool of enrollees is sicker than expected.

[4] These rules apply to all states except New York and Vermont. In those states, the age curve is flat—the same price applies to everyone—and there are separate factors for single parent coverage (i.e. a plan for a single parent and any number of children) and for family coverage.

[5] My description in this subsection draws heavily on Cox et al. (2016), who provide a clear overview of these programs.

Risk adjustment does not affect the industry average cost curve.

While risk adjustment provides insurance against taking on patients with high expected costs, risk corridors and reinsurance provide some insurance against high realized costs. The risk corridor program provides a payment to insurers whose total allowable costs—claims costs and some other allowable expenses—are greater than 83% of the their premium revenue. Plans with allowable costs below 77% of premium revenue make a comparable payment. The risk corridor program thus subsidized insurers with high overall costs, and penalized insurers with low costs.

By contrast, the reinsurance program insures plans against very high costs for individual enrollees. In particular, if an enrollee has incurred claims above $45,000, the program pays a reinsurance rate of r on the claims above this amount and below $250,000. The program is financed with fees on all insurers, but only insurers in the individual market are eligible for it. It therefore represents a transfer from other insurance markets to the individual insurance market. In 2014, reinsurance payouts totalled $7.9 billion, or about $500 per enrollee in the individual insurance market.[6] Like the risk corridors, the reinsurance program was phased out gradually. In 2014, the reinsurance rate was set to be 80%, but higher than expected revenues meant that it paid out 100% of claims. In 2015 and 2016, it was set to be 40%, and it was eliminated in 2017.

I largely abstract from the premium stabilization programs in my analysis, which focuses on the consequences relative to 2017 policy of repealing or replacing the mandate and/or the premium tax credit. This is because risk adjustment and risk corridors have little effect on competitive equilibrium, and reinsurance expired in 2017. When simulating 2014 equilibrium, however, I account for reinsurance, treating it as a pure subsidy in the individual insurance market.

2.3 Mandates, tax credits, and cost-sharing subsidies

The ACA contains two provisions that encourage health insurance purchase, the "shared responsibility payment," i.e. the mandate; and the premium tax credit. The shared responsibility payment is a tax payment required for each month of non-coverage. The payment is the larger of a fixed percent of family income or a flat dollar amount, per person-month of non-coverage in the tax family. The payment is capped at the national average price of a bronze plan, for that family

[6] The denominator in this calculation is not the 8 million people with on-exchange plans , but the 15.6 million people in the individual insurance market, all of whom are eligible for reinsurance payments (Levitt et al., 2015).

structure. For a family in which no one buys any insurance, the exact amount is

$$SRP = \min \{NBP, \max \{pct(AGI - thresh), flat(N_a + N_c/2)\}\}, \tag{1}$$

where SRP is the shared responsibility payment, NBP is the national average price of a bronze plan, AGI is adjusted gross income, $thresh$ is the tax filing threshold, N_a and N_c are the number of adult and children, and pct and $flat$ are the percent of income or the flat dollar amount, which vary from year to year. pct is 1% in 2014, 2% in 2015, and 2.5% in 2016 and onward. $flat$ is \$95 in 2014, \$325 in \$695 in 2016; beyond 2016 is adjusted inflation. Families with income below the filing threshold are exempt from the mandate, as are some low income families, families without affordable coverage, and Native Americans.

The mandate is a tax incentive for all households to obtain insurance, and its bite is weakly increasing in household income. Some households are also eligible for the premium tax credit, which helps make insurance affordable. To receive the credit, households must have income between 100 and 400% of the federal poverty line (FPL), be unable to purchase affordable health insurance through employer sponsored insurance, and be ineligible for Medicaid. Only insurance purchased in a health insurance marketplace qualifies for the premium tax credit; off-exchange purchases do not. Households that qualify for the premium tax credit have it paid throughout the year; it goes directly to the insurance company from which they bought coverage.

The premium tax credit works as follows. First, the law specifies an expected household contribution, which depends on household income (modified adjusted gross income, Y) relative to the federal poverty line; call this $EC(Y)$. The expected contribution is a *percent* of household income, and this percent is weakly increasing in the *level* of household income. The expected contribution is 2% of income for people with income between 100-138% of FPL, jumps to 3% at 138% of FPL, and then rises steadily to 9.5% of income at 300% of FPL, where it remains until 400% of FPL, when eligibility ends.

In practice, the actual price of health insurance may exceed a family's expected contribution. The premium tax credit covers the difference. In particular, it is exactly equal to the difference between a households expected contribution, and their benchmark price, the price for them of the second lowest cost silver plan in that family's county, denoted P_f. Net of the premium tax credit, the price of the benchmark health insurance plan for a family with income Y is

$$\bar{P} = \begin{cases} \min\{Y * EC(Y), P_f\} & \text{if } 1 \leq Y/FPL \leq 4 \\ P_f & \text{otherwise.} \end{cases} \tag{2}$$

Thus, viewed as a function of income, $\bar{P}$ is kinked or discontinuous, given P_f. If $P_f > 0.095(4)FPL$, then $\bar{P}$ has a discontinuity at 400% of FPL. Otherwise $\bar{P}$ has a kink where the expected contribution is exactly equal to the benchmark price. Figure 1 illustrates these kinks and discontinuities for a 25, 35, or 45 year-old single person in 2014 in Monroe County, Indiana. It shows the expected dollar contribution, given income, and the benchmark price of insurance. For 25 and 35 year-olds, $\bar{P}$ is kinked in income. For a 45 year-old, it is discontinuous in income.

In addition to subsidies for insurance purchase, the ACA also includes subsidies for cost-sharing reductions (CSR). The law requires that insurers offer plans with reduced cost-sharing requirements, but at the same price as their silver plan. The government finances the difference between the baseline coverage and the enhanced coverage. Only people with low income are eligible to purchase these plans: people with income below 150% of FPL may purchase the most generous plan, which provides 94% actuarial value (relative to 70% of a standard silver plan), people with income between 150 and 200% of FPL may purchase a plan with 87% actuarial value, and people within income between 200 and 250% of FPL may purchase a plan with 73% actuarial value. DeLeire et al. (2016) show that the cost-sharing reduction subsidies influence plan choice on the intensive margin, but they do not affect enrollment in the Marketplace. I account for them in some of my empirical specifications by including dummies for household income below 200% of FPL, where the generous CSR subsidies apply.

3 Model

I begin by developing a competitive model of supply and demand in a market for insurance when insurers must charge the same price to all enrollees, regardless of health status, closely following Einav et al. (2010). I then extend the model to incorporate many of the ACA's relevant institutional features: an age-curve for pricing, a penalty for non-purchase, and a price-linked subsidy.

3.1 Basic model

Set-up I consider an insurance market with a single, fixed product available. People differ according to a vector of heterogeneity ξ_i, distributed according to F. This heterogeneity completely describes both risk preferences and riskiness. It also includes all information relevant to market regulations, like age and family income. Each person has expected costs $c(\xi)$, which represents the expected cost to the insurer of providing coverage to a person of type ξ, given the fixed insurance contract. $c(\xi)$ differs across people because of differences in health characteristics, including age, sex, and pre-existing conditions, as well as differences in taste for health and health care. $c(\xi)$ differs from total expected spending in two ways: it represents the insurer's costs, so it does not include third party payments (such as reinsurance) or out-of-pocket payments; and it includes non-medical costs, often called loads, such as processing costs, commissions or other marketing costs, and other administrative costs that scale with the number of people covered.[7]

People also differ in their willingness-to-pay for insurance, $v(\xi)$.[8] Willingness-to-pay differs across people because of differences in expected costs, but also because of differences in risk preferences. In the absence of heterogeneity in risk preferences, costs and willingness to pay would be perfectly (rank) correlated, as high expected costs are the only reason for high willingness-to-pay. At any given price, the set of people willing to buy health insurance would have higher expected costs than the population as a whole, leading to adverse selection. But heterogeneous risk preferences can reverse this correlation if, for example, highly risk averse people engage in many healthy behaviors and so have low expected costs but high willingness to pay, leading to advantageous selection (de Meza and Webb, 2001). This positive correlation has been documented in many settings; for example in long-term care insurance (Finkelstein and McGarry, 2006) and Medigap insurance (Fang et al., 2008). A key feature of the modelling approach here is that it imposes no restriction on F, and so it does not restrict the joint distribution of expected costs and willingness to pay.

Demand and costs At any price of insurance coverage p, demand is the measure of people with willingness to pay above p:

$$D(p) = \int_{\xi} 1\left\{v(\xi) \geq p\right\} dF(\xi). \tag{3}$$

[7] Note that the model imposes no restrictions on moral hazard. To see this, let $c^{total}(\xi)$ measure total costs under the contract (out-of-pocket and insurer) and let $c_0^{total}(\xi)$ measure total cost in the absence of insurance. Moral hazard—the effect of insurance on incurred costs—is $c^{total}(\xi) - c_0^{total}(\xi)$, and I place no restrictions on this object, because $c_0^{total}(\xi)$ is unrestricted.

[8] Here I implicitly assume away any income effects. This assumption is exactly true for CARA utility, and approximately true for small price changes.

The industry average cost curve is equal to the average expected costs of people who demand insurance at p:

$$AC(p) = \frac{1}{D(p)} \int_{\xi} c(\xi) 1 \{v(\xi) \geq p\} \, dF(\xi) \tag{4}$$

As long as $c(\xi)$ and $v(\xi)$ are not independent, $AC(p)$ will be different from the unconditional average of $c(\xi)$ in the population, $\bar{AC}$. When $c(\xi)$ and $v(\xi)$ are positively correlated, $AC(p)$ will be higher than $\bar{AC}$, indicating adverse selection.

It is straightforward to define the marginal cost curve, $MC(p)$, which gives the expected costs of people who are just indifferent between buying insurance and not at p:

$$MC(p) = E[c(\xi)|v(\xi) = p]$$

In some cases—for example, without moral hazard and when the only cost to the insurer are medical costs—willingness to pay for insurance is always above cost, and so the marginal cost curve everywhere lies below the (inverse) demand curve. In those cases, full coverage is optimal, but it may not be an equilibrium outcome, because insurers cannot charge different prices to different people.

Equilibrium I consider competitive equilibrium. Assume there are $n \geq 2$ homogenous firms simultaneously choosing prices p to maximize revenue net of costs. As the products are identical, enrollees are indifferent between the firms and always buy from the lowest price one. If both firms charge the same price, enrollees are allocated randomly between them. The model accounts for guaranteed issue and community rating by requiring that firms charge the same price to all enrollees, and offer insurance to all who wish to buy.

As the higher price firm always makes zero profit, and the lower price firm makes a positive profit as long as $p > AC(p)$, the equilibrium price is simply the lowest break-even price:

$$p^{eq} = \min \{p : AC(p) = p\}, \tag{5}$$

and equilibrium demand is $D(p^{eq})$.

Figure 2 illustrates equilibrium. The average cost AC curve is downward sloping in q, indicating adverse selection: as more people enroll in insurance, average costs fall, so the most expensive people to insure are the first to buy insurance. The marginal cost curve MC lies below the demand curve D_0, meaning that everyone would like to buy insurance that is actuarially fair for them.

The equilibrium quantity is below 1, however, because of adverse selection. Some people have willingness to pay below $\bar{AC}$, so they refuse to buy insurance at the break-even price for the whole population. Their exit from the market drives up cost and prices, driving more people out and further driving up cost. As the figure shows, to find equilibrium, we only need to know $D(p)$ and $AC(p)$, and not F, the underlying distribution of heterogeneity. These objects can be estimated from data on insurance demand and costs, given sufficient variation in premiums.

3.2 Incorporating ACA institutions

The model so far abstracts from several features of the ACA that are likely to be quantitatively important in equilibrium prices. Here I describe how to integrate the age curve, the premium tax credit, the mandate, and the reinsurance subsidy into the model. I introduce the age curve by assuming that insurers continue to set a base price, now denoted p^b, but the price a person of type ξ face is $f(\xi)p^b$, where $f(\cdot)$ is the age factor for ξ, for example 1 for a 21 year-old and 3 for a 64 year-old. I assume that the premium tax credit and the mandate affect demand by changing the effective price that people pay. I define the net price for a person of type ξ when the base price is p^b as

$$\tilde{p}(p^b, \xi) = f(\xi)p^b - S(\xi, p^b) - \tau(\xi), \tag{6}$$

where $S(\xi, p^b)$ is the subsidy for a person of type ξ when the base price is p^b,[9] and $\tau(\xi)$ is the mandate penalty for a person of type ξ.

Demand and costs As before, people buy insurance if their willingness to pay exceeds their price, but price is now type-specific. Expressed as a function of the base price, demand is

$$D(p^b) = \int_{\xi} 1\left\{v_i(\xi) \geq \tilde{p}(p^b, \xi)\right\} dF_{\xi}. \tag{7}$$

The average cost of the insured, net of reinsurance can likewise be expressed as a function of the base price, but now with two new parameters

$$AC(p^b) = D(p^b)^{-1} \int_{\xi} c_i(\xi) 1\left\{v_i(\xi) \geq \tilde{p}(p^b, \xi)\right\} dF_{\xi} - r. \tag{8}$$

Here r is the per-person expected value of the reinsurance subsidy, which I separate out from expected costs because it is a policy variable that affects equilibrium prices.

[9] $S(\xi, p^b) = \max\left\{EC(\xi) - f(\xi)p^b, 0\right\}$, where $EC(\xi)$ is the expected dollar contribution for ξ.

With age-based pricing, average revenue is no longer equal to p^b, because revenue depends on the set of people who buy insurance at p. Instead average revenue given the base price p^b is

$$AR(p^b) = D(p^b)^{-1} \int_\xi p^b f(\xi) 1 \left\{ v_i(\xi) \geq \tilde{p}(p^b, \xi) \right\} dF_\xi. \tag{9}$$

While demand depends on the net price $\tilde{p}$, revenue depends on the full price $f(\xi)p^b$.

Equilibrium Equilibrium continues to be characterized by the lowest break-even price.

$$p^{eq} = \min_{p^b} \left\{ p^b : AR(p^b) = AC(p^b) \right\}. \tag{10}$$

The subsidy and the mandate are intended to limit adverse selection by encouraging more people to purchase insurance. Figure 2 illustrates the effect of demand subsidy of size s (or equivalently a mandate with penalty $-s$). In the absence of the subsidy, adverse selection prevents low-cost people from purchasing insurance. With the subsidy, the market is fully covered, and the equilibrium price net of the subsidy falls by $p_0 - p_1$ as the average cost of the insured falls as the subsidy draws more people into the market. The effect of the subsidy and mandate on prices and coverage ultimately depends on the slope of the cost and demand curves, and the empirical goal of the paper is to estimate these objects.

Welfare Imposing a mandate, or subsidizing coverage, can improve welfare. Welfare at a given equilibrium base price p is always consumer surplus less total costs:

$$\begin{aligned} W(p) &= \int_\xi (v(\xi) - c(\xi)) 1 \left\{ v(\xi) \geq \tilde{p}(p, \xi) \right\} dF(\xi) \\ &= \underbrace{\int_\xi v(\xi) 1 \left\{ v(\xi) \geq \tilde{p}(p, \xi) \right\} dF(\xi)}_{\text{consumer surplus}} - \underbrace{AC(p)D(p)}_{\text{total cost}}, \end{aligned} \tag{11}$$

The change in welfare from a given policy that changes equilibrium price from p^0 to p^1 is simply $W(p^1) - W(p^0)$. Figure 2 illustrates the welfare consequences of a mandate or subsidy. In the absence of a subsidy, demand is given by D_0 and equilibrium quantity, Q_0, is inefficient because everyone has demand above marginal cost, but not everyone has willingness to pay above $\bar{AC}$. A subsidy that shifts the demand curve to D_1 leads to complete coverage. The welfare gain is given by the shaded area, which represents in the increase in consumer surplus among the newly insured, less their costs. This is identical to the welfare gain derived by Hackman et al. (2015), in the absence of markups. As the figure suggests, knowledge of D and MC are sufficient to calculate this welfare gain. Because MC can be recovered from AC, only D and AC are necessary to estimate welfare (Einav et al., 2010).

These welfare calculations assume that demand reveals welfare-relevant preferences. A growing literature, however, notes that demand for insurance suffers from several biases, including excess aversion to deductibles (Sydnor, 2010), information frictions (Abaluck and Gruber, 2011; Handel and Kolstad, 2015) and switching costs (Handel, 2013), and so may not reflect true preferences. Thus these welfare estimates should be interpreted cautiously. They indicate whether the subsidy or mandate induces "too much" coverage by encouraging take up among people with (revealed) willingness-to-pay below costs, but they do not account for behavioral biases or other choice frictions, and they may understate the welfare gains from increasing coverage (Spinnewijn, 2017).

Discussion of modelling choices The modelling approach here makes two important simplifying assumptions: binary insurance choice, and perfect competition. The assumption of binary insurance choice is common in the literature (e.g. Cutler and Reber (1998); Einav et al. (2010); Hackman et al. (2015)). It simplifies the analysis because equilibrium can be graphically represented, and it eases the identification requirements. The price, quantity, and welfare consequences of a mandate or subsidy can be inferred from a single demand curve, without having to estimate cross-price elasticities.

Focusing on binary insurance choice is limiting, however, because it shuts down a dimension of adverse selection. In particular, while binary insurance choice permits adverse selection to affect the extensive margin (how many people have insurance coverage) it does not allow for intensive margin effects (how generous is that coverage). Rothschild and Stiglitz (1976) establish theoretically that adverse selection reduces equilibrium coverage generosity (if an equilibrium exists), and Handel et al. (2015) use simulations to show that these effects can be large. Empirical research confirms that adverse selection can drive the most generous plans out of the market (Cutler and Reber, 1998; Buchmueller and DiNardo, 2002), and reduce their network size (Shepard, 2016). In the ACA marketplaces this intensive margin adverse selection is limited by regulation which establishes minimum generosity (e.g. 60% actuarial value, with caps on out-of-pocket spending). The ACA also requires that Marketplace plans maintain adequate network size, although it does not offer definite standards for doing so (Giovannelli et al., 2015). Thus, although although the modelling approach here cannot speak to intensive margin adverse selection, it is still valuable because it helps address a key policy question: how do alternative policies affect the equilibrium fraction of people with creditable insurance coverage, as well as the price of that coverage.

It is also a limitation that the model assumes perfect competition. With an estimated demand and cost curve, it is possible to simulate equilibrium under any market structure, however (at the cost of assuming that consumers view all insurers as perfect substitutes). In robustness exercises, I consider the effect of of policy changes with a monopolist insurer, rather than a competitive equilibrium. This is an extreme assumption, but it gives a sense of how allowing for imperfect competition could affect the policy simulations, and it turns out that the comparative statistics are quite similar across market structures (although of course the level of prices, coverage, and welfare are quite different). An alternative approach would be to directly model the demand for differentiated insurance products, as Ericson and Starc (2015); Tebaldi (2016) and Shepard (2016) do. The disadvantage of that approach, relative to the one taken here, is that it would require more identifying cross price elasticities (i.e. intensive margin responses) as well as own price elasticities.

4 Identification strategy and data

4.1 Identification strategy

The empirical goal of this paper is to estimate the effect of premiums on demand for insurance and the average cost of the insured. The main empirical challenge is the endogeneity of premiums. Prior to 2014, premiums are typically unobserved, and they depend in an unknown way on individual health status. Even in 2014 when premiums are observed, they are a deterministic function of the local benchmark premium, family composition, and family income. These variables are likely to independently affect insurance demand and to be correlated with health care spending, but it is impossible to control for them in a fully flexible way, because doing so removes all price variation. One solution to this problem is to use a regression discontinuity design or a regression kink design, taking advantage of the discontinuities and kinks in the tax subsidies (as a function of income). This approach is not feasible with the survey data I study, because there is likely considerable measurement error in family income, which is self-reported.

Instead, I rely on an instrumental variables approach that takes advantage of the nonlinearities in the subsidy schedule, without requiring a precise measure of family income. To understand the approach, note that we can classify people (i) by whether their base premium is high enough that their price is a discontinuous function of income ("discontinuity sample") and (ii) by whether they receive a subsidy or not ("subsidized"). For example, as Figure 1 shows, single 45 year-olds in Monroe County, IN are in the discontinuity sample, but single 35 year-olds are not; a 35 year-old making \$35,000 would receive a subsidy, but 35 year-old making \$45,000 would not.

Using this classification, I define two variables, $aboveKink$ and $aboveDisc$, which are indicators for people whose income is high enough that they are above the kink or discontinuity on their subsidy schedule, and so receive no subsidy. It is possible that higher income people have high demand for insurance or low costs. It is also possible they have a differential trend in demand or costs from 2013 to 2014. However my identification assumption is that higher income people in high premium areas (who are more likely to be subsidy eligible) have the same trend in demand and costs as higher income income people in low premium areas. This is essentially a triple difference approach. I operationalize this identification assumption with the following specifications:

$$\begin{aligned} y_{it} =& \beta_0 + \beta_1 \tilde{p}_{it} + \beta_2 aboveKink_{it} + \beta_3 aboveDisc_{it} + \beta_4 discSample_i \\ &+ \beta_5 post_{it} + \beta_6 discSample_{it} \times post_{it} \\ &+ \beta_7 inc_{it} + \beta_8 inc_{it}^2 + \beta_9 inc_{it} \times post_{it} + \beta_{10} inc_{it}^2 \times post_{it} + x_{it}\beta + \varepsilon_{it} \end{aligned} \tag{12}$$

$$\begin{aligned} \tilde{p}_{it} =& \alpha_0 + \alpha_1 aboveKink_{it} \times post_{it} + \alpha_2 aboveDisc_{it} \times post_{it} \\ &+ \alpha_3 aboveKink_{it} + \alpha_4 aboveDisc_{it} + \alpha_5 discSample_i \\ &+ \alpha_7 inc_{it} + \alpha_8 inc_{it}^2 + \alpha_9 inc_{it} \times post_{it} + \alpha_{10} inc_{it}^2 \times post_{it} + x_{it}\alpha + \varepsilon_{it} \end{aligned} \tag{13}$$

Here y_{it} is either a dummy variable for insurance coverage, or expected costs among the insured. The coefficient of interest is β_1, the effect of prices on insurance coverage or expected costs. The specification allows for different levels of insurance demand for people above their kink or discontinuity. Since the level of the kink or discontinuity depends on age, family composition, and local prices, this specification controls for permanent differences in y_{it} that are related to these variables.I also control for them age, family composition, and prices in the vector x. The specification controls for any general change in y_{it} in 2014 (with the variable $post$), and it allows for a differential trend for people in the discontinuity sample (via the interaction $discSample \times post$). The specification allows for differential trends by income levels because it controls for a quadratic in income, interacted with $post$. In some specifications, I also control for a full set of state-by-year fixed effects, to account for any possible state-specific policy changes (e.g. Medicaid expansion, use of the federally facilitated marketplace) that may affect the individual insurance market. The excluded instruments are the interactions of $aboveKink$ with $post$ and $aboveDisc$ with $post$. Price sensitivity, β_1, is identified by any differential change in insurance coverage among those receiving a subsidy, scaled by the change in the net price of insurance in those groups, and adjusting for differential trends by income. Note that the endogenous regressor, $\tilde{p}_{it}$, measures prices net of the mandate as well as the premium tax credit. Thus the estimates account for the fact that subsidy eligible people have relatively low income and face a lower mandate penalty as a result.[10]

A limitation of my data is that I do not observe true prices prior to 2014.[11] Instead I impute prices pre-2014 as the gross price of insurance (not including the subsidy or the mandate penalty) in 2014. This imputation creates substantial measurement error in prices prior to 2014, but I show in Appendix A that this is not a problem for identification. The intuition is that I identify price sensitivity by using shocks to *subsidies* rather than to prices. Because these subsidies are known to be zero prior to 2014, I can identify responses to subsidy changes even if I do observe not price levels.

The identification strategy also helps address the possibilities that subsidies may be mismeasured in 2014. Subsidies depend on family income, which is likely measured with error. Nonetheless the instrumental variables are free of measurement error as long as the measurement error in income does not move a person over their kink or discontinuity. To guard against this possibility, I limit the estimation sample to people whose family income is far from their kink or discontinuity, in particular at least 20% away from it.[12] Excluding people in this way also guards against the possibility that people with high demand for insurance manipulate their income to move under the 400% cutoff (Heim et al., 2016).

These instruments address the endogeneity of prices because they isolate policy-induced price variation that is plausibly unrelated to unobserved determinants of insurance demand or health care costs. I test the validity of the instruments in two ways. First, if demand or costs are evolving differentially with the instrument, then the instruments should be correlated with predetermined variables that are related to insurance demand or costs; I show below that the instruments predict no effect on predetermined variables. This test rules out an important potential threat to identification: in principle low income people may be in poor health, and therefore excluded from the individual insurance market. However the instruments appear uncorrelated with health. As a second overall

[10] This is important for scaling price sensitivity, but not for identification. The mandate penalty is roughly linear in income, and I control for income×post, so this differential mandate penalty is controlled for.

[11] It might seem that an adequate proxy for price is insurer revenue per member per month, which is reported by the NAIC. This measure of course reflects the average price of the insured, but it is not useful for measuring the price of the uninsured; because of risk based pricing and underwriting restrictions, the uninsured faced pre-ACA faced very different prices than the insured, if they could buy insurance at all.

[12] I chose this threshold under the hypothesis that individual income could be rounded by as much as 10%, and so family income might be off by 20%. I show below that the results are not sensitive to the exact cutoff.

test of instrument validity, I also present a placebo test that treats 2013 as the treatment year and 2012 as the control year; I find no effect of the instruments prior to the main ACA policy changes.

This identification strategy is similar to the approach used by Frean et al. (2017). They identify the effect of the subsidy on insurance coverage by regressing coverage on the imputed subsidy amount, and subsidy×post-2014 (as well as many other controls and policy variables). Their approach and mine rely on fundamentally similar variation: we both use the change in insurance coverage among people who receive the subsidy, relative to people who do not, adjusting for general trends by income level. I differ from them, however, in that my approach is more robust to the likely possibility of measurement error in imputed subsidy (deriving from measurement error in family income). Hackman et al. (2015) also estimate the effect of a mandate, but my identification strategy differs from theirs. My approach is to use exogenous price variation to estimate the demand and cost curves, and use these curves plus an equilibrium assumption to infer the effect of policy changes on coverage and premiums. By contrast, Hackman et al. (2015) use a policy change itself—the imposition of the mandate in Massachusetts—to empirically measure the impact of that policy on coverage and premiums, and use the empirical effects of the policy to infer its welfare consequences.

4.2 Data

Estimation requires market-wide data on insurance demand and costs, which necessitates nationally representative data on insurance coverage, medical spending, and prices. I use the Medical Expenditure Panel Survey (MEPS), the only nationally representative data set I am aware of with health care spending information. The MEPS is an overlapping panel data set: new households enter the panel every year, stay in it for two years, and have five interviews. The survey includes questions about health insurance offers and coverage, as well as medical spending, health status, and income, and a health insurance family identifier. I work with an annual summary file which records total health care spending over the course of the year. The survey does not contain information about offered premiums, so I supplement it with data from `healthcare.gov` and hand-collected data on the benchmark silver premium for each county in the United States in 2014 (except in Kentucky, where I was unable to obtain data). Additional information about the construction of the benchmark premiums is in Appendix B.[13] Because the MEPS is relatively small, I also repli-

[13] I matched observations to their premiums using county identifiers, which are not in the public use MEPS file. The merge was conducted by AHRQ contractors, who returned a dataset with premiums and encrypted county identifiers. I performed all data analysis at the AHRQ data center.

cated the analysis using the Current Population Survey. Appendix C provides further details and all replication results.

Variable creation I define insurance coverage using a dummy variable for any insurance coverage in the past year. This definition pools together on- and off-exchange purchases (because I exclude people with Medicaid, or eligibility for ESI or Medicare), but this is appropriate because insurers must charge the same price to both groups. Thus I measure insurance demand on the extensive margin only, abstracting from spells of non-coverage or quality choices. This choice is consistent with the modeling approach, where the coverage choice is on the extensive margin and there is no quality differentiation.

I calculate family income as the sum of total personal income within the insurance family, and convert this to a percent of FPL using year and family-size specific poverty levels. I define the gross family price as the price for family coverage of their benchmark silver plan. Given gross prices and income, I calculate the premium tax credit and the mandate penalty, and define the per-person net price as the gross price, less any tax credits for purchase or penalty for non purchase, divided by the number of people in the health insurance unit. In 2012 and 2013, the net price is the gross price.

I use family income and gross prices to define the instruments. Families are in the discontinuity sample if their gross price is at least 38% of their poverty line (i.e. 9.5% of 400%); all other families are in the kink sample. I define $aboveKink$ and $aboveDisc$ as dummy variables equal to one if the family does not receive a subsidy. The instruments are interactions of $aboveKink$ and $aboveDisc$ with $post$, a dummy variable for 2014 and later.

Of course, many plans are offered in in the individual insurance market, and the benchmark premium is not the only price. Nonetheless, there is little loss from ignoring this additional price variation. My instrumental variable analysis isolates price variation that shifts the price of any insurance plan relative to no insurance, but it induces very little variation in the relative price of different insurance plans.[14] This price variation is valuable for detecting adverse selection for the market as a whole; it reveals whether, as insurance overall becomes cheaper, low-cost people become more likely to buy insurance.

To estimate average cost curves, we require a measure of expected costs that varies across individuals, so that it can be linked to subsidy eligibility, prices, and insurance coverage. The

[14] If a family that qualifies for a premium tax credit of \$X, the price of a given plan net of the premium tax credit becomes $\min\{P - X, 1\}$. The premium tax credit does not change the relative price of plans that are sufficiently expensive.

MEPS measures actual costs, but these are likely to be measured with considerable error. It is also unlikely that people know their exact costs at the time they sign up for insurance. I therefore develop an alternative measure of expected costs based on self-reported health status and demographics. Aizawa and Kim (2015) also use self-reported health status to impute expected health care costs, in their study of selection in Medicare Advantage. Using insured people in the Marketplace sample in the 2014 MEPS, I calculated insurer costs as total medical spending less out-of-pocket spending.[15] I then regressed insurer costs on a set of demographic variables and dummy variables for each level of self-reported health. These variables likely represent much (though surely not all) of the private information available to people at the time they make their health insurance decisions.

I define expected costs as the predicted values from this regression. The imputation regression estimates are reported are in Appendix Table C.1. Health status is highly predictive of spending. People in fair health have an additional $5,500 in spending, and people in poor health and additional $8,800 in spending. Older people also have higher expected spending. Thus, this measure of imputed spending is well suited for detecting whether younger people or people in good health—the "young invincibles"—are especially likely to select out of the individual insurance market in 2014, an oft-cited adverse selection concern.

Sample selection Using the MEPS data, I impose three restrictions to create a sample of people eligible to participate in the health insurance Marketplaces; I call this the Marketplace sample. First, I exclude all family-years that contain at least one person who is offered ESI in each month of the year; people who are offered affordable ESI are ineligible for subsidies.[16] Second, I exclude all person-years who report having any Medicaid coverage, and all people with family income below 138% of FPL, because people who are Medicaid eligible are ineligible for subsidies.[17] Third, I exclude all people aged 65 and older, who are ineligible for subsidies because they qualify for Medicare. These restrictions reduce the sample from 104,743 person-years in 38,102 families to 15,636 person-years in 8,656 families. The sample is small because most Americans receive insurance from their employer or the government. Finally, in the IV estimation, I further exclude people with household income within 20% of their household's cutoff for receiving a subsidy for purchasing income. For households in the discontinuity sample, this cutoff is always at 400% of FPL.

[15] I use the 2014 MEPS only so that the cost-sharing underlying the imputation reflects 2014 insurance plans. The results are not sensitive to this choice, however, nor to the exact form of the imputation.

[16] Here and throughout, "family" refers to the health insurance family, a MEPS-defined concept identifying the set of people eligible to purchase coverage together, roughly married adults and their dependents.

[17] Medicaid eligibility extends to all adults with income up to 138% of FPL in the expansion states; children and some adults are eligible at higher levels of income.

For households in the kink sample, the location of this cutoff depends on the family premium.[18] Appendix Table C.2 compares the full sample, Marketplace sample, and estimation sample. Relative to the full sample, the estimation sample is similar in age, but contains fewer children and more high-income households (due to the exclusions on Medicaid enrollees and household income). The Marketplace sample and estimation sample are quite similar.

Summary statistics by insurance status Table 1 provides summary statistics by year and insurance status for the marketplace sample. The insurance coverage rate is roughly flat between 2012 and 2013, then jumps to 67% in 2014. Part of this increase is due to the reduction in net prices as the mandate and premium tax credit went into effect in 2014. The gross price averages about $3,700; the average price net of the subsidy and mandate is lower by nearly a third. Of course, the increase in insurance also reflects the many other ACA provisions that went into effect in 2014, including the Health Insurance Marketplaces, community rating and guaranteed issue, and the imposition of 3:1 age bands. Surprisingly, however, despite the ban on risk-based pricing, there is no evidence that older people, people with chronic conditions, or generally unhealthy people (i.e. people with high expected insurer costs) were differentially likely to gain insurance in 2014.[19] Less healthy people in a given year are more likely to be insured (as the insured always have higher expected costs), but this difference is constant over time.

To show this, the final column of the table reports the "difference-in-difference", as well as its standard error. This difference-in-difference is, for each variable x,

$$\begin{aligned}\Delta\Delta \equiv & (E[x|insured, t = 2014] - E[x|uninsured, t = 2014]) \\ & - (E[x|insured, t < 2014] - E[x|uninsured, t < 2014])\end{aligned}$$

If the ban on underwriting worsened the health of the insured population, then we should see $\Delta\Delta$ large and positive for characteristics that correlate with health care spending. In fact $\Delta\Delta$ is small and insignificant; there is no evidence of a worsening risk pool. The insured and uninsured were similarly different in 2014 and in early years, in terms of age, family income, or the presence of chronic conditions. Women were especially likely to gain insurance coverage, but the overall difference in expected insurer cost—a dollarized measure of selection—remained stable in 2014.

[18] Calculating the exact cutoff is nontrivial because the expected family contribution is a nonlinear function of income when income is between 138 and 300% of FPL. It is well approximated by a quadratic function over this range, however, and I use the quadratic approximation to find the cutoff points.

[19] Any chronic condition is an indicator equal to one if a person reports having had any of the following chronic conditions: angina, heart attack, other heart disease, stroke, emphysema, cholesterol disorder, cancer, diabetes, arthritis, asthma.

Despite the many changes in the individual insurance market, the risk pool did not substantially worsen.

4.3 Replication in the Current Population Survey

Before arriving at my main estimating equations, I considered alternative (though similar) specifications based on different sample restrictions, instrument groupings, and variable definitions. The results are robust to across these specifications, as I show below. As further reassurance that the reported estimates are not the result of specification searching, I also replicate all results using data from the CPS. Appendix C provides more details, including a discussion of the similarities and differences between the CPS and MEPS data. All results are similar in the CPS and MEPS, although of course the exact estimates differ because of sampling error and because of difference in variable definition and sampling frame.

5 Estimates

5.1 Validity and strength of the instruments

I begin by establishing the validity of the identification strategy. My exclusion restriction requires that the instruments—$aboveKink \times post$ and $aboveDisc \times post$—affect insurance demand and average costs only by changing the effective price. A necessary condition is that they be uncorrelated with observed determinants of insurance demand and costs, conditional on the controls. Table 2 provides evidence in support of this condition. It shows the "effect" of premiums on an indicator for self-reported poor health, an indicator for any chronic condition, and expected insurer costs, estimated with Equation 12.

I selected these variables because they are closely related to insurance demand and costs, but it is unlikely that insurance coverage has any direct effect on them (at least in the short run). Thus a significant effect here indicates a failure of the identification assumption. Instead, across all variables, the effect of price is small and insignificant. This table shows that the instruments are likely uncorrelated with the main non-price determinants of insurance demand and costs.

This establishes the validity of the instruments. Panel A of Table 3 shows their strength. It reports the first stage regression of effective price per person on the instruments and controls. In the base specification, being above the kink raises the effective price by about $1,100, and being above the discontinuity raises it by about $2,600. These estimates are highly significant individually and jointly, with an F-statistic of 379. The remaining columns of the table show robustness to alternative specifications: using different cutoffs for near the threshold, controlling for state-year fixed effects, or controlling for dummy variables for eligibility for cost-sharing. The final column shows the first stage from a placebo regression which excludes 2014 and treats 2013 as the year the subsidies and mandate go into effect.

5.2 Main estimates

Demand Table 3, Panel B, presents estimates of the demand for insurance. In the basic specification, the price coefficient is -6.18. The average net price is about $3,200, so a price increase of 10% of the average price would reduce insurance demand by about 2 percentage points, or about 3% of the baseline rate of 60%; the estimates imply an elasticity of about -0.3. This elasticity is identified from the change in insurance status among subsidy eligible people, scaled by their 2014 subsidy amount, and adjusting for differential trends by eligibility and by income. The key to identification is that insurance coverage grew faster among the subsidy eligible than the ineligible, even adjusting for their differential income.

The elasticity estimate is is on the low end of the literature, but comparable to past results. For example, in the ESI context, Cutler and Reber (1998) report elasticities in the range -0.3 to -0.6, Royalty and Solomon (1999) report elasticities of -0.2 to -0.8, and the estimates in Einav et al. (2010) imply elasticities on the order of -0.4. Ericson and Starc (2015) estimate elasticities from the Massachusetts of about -0.82 for people age 45 and above, and -0.27 for younger people. These estimates are all obtained, however, by studying intensive margin plan choice—between more and less generous insurance plans. The estimates in this paper are from the extensive margin choice of any insurance or no insurance, so it is not surprising that the elasticity estimate is smaller. Hackman et al. (2015) also estimate extensive margin responses, and they find a larger elasticity, about -0.8. Their identification strategy, however, focuses on high-income adults (who face a mandate but are not subsidized).

Columns (2)-(6) show that the estimated price sensitivity is robust to alternative specification choices. In columns (2) and (3) I change the distance cutoff used to select the estimation sample, from 20% at baseline to 10% or 30%. Using a narrower cutoff decreases the estimated coefficient slightly, and using a wider cutoff increases it, consistent with measurement error in family income. In column (4), I control for a full set of state by year fixed effects. These controls adjust for any state-specific trends, arising for example because of differences in state Marketplaces or Medicaid policies. In column (5), I add a dummy for income below 200% of FPL, as well as its interaction with post-2014. 200% of FPL is the cutoff for eligibility for generous CSR plans. People eligible for these plans are also generally subsidy eligible, so these controls ensure that I do not conflate price sensitivity insurance generosity sensitivity. Across all these specifications, the estimated price sensitivity is stable.

Costs As the price rises, fewer people buy insurance. These remaining buyers also have higher costs on average, as revealed by the the the estimated cost curve, shown in Table 4.[20] The average cost curve is upward sloping, indicating some adverse selection. The point estimates are around 0.44, indicating that a $1,000 increase in premiums increases the average cost of the insured by $440. This result demonstrates adverse selection in the individual insurance market. At lower prices—induced by exposure to the premium tax credits—relatively healthy are more likely to self select into the insurance market, bringing down expected costs. The estimated slope of the cost curve is robust to choosing alternative distance threshold, as well as controlling for state-year fixed effects or eligibility for cost-sharing reductions.

Placebo tests The estimated demand and cost curve are therefore robust to using alternative distance cutoffs or different sets of controls. Nonetheless it is possible t hat differential trends by subsidy eligiblity could explain the results. To rule out this possibility, the final column of Table 3 and Table 4 shows the result of a placebo test, in which I take 2013 to be the treatment year, and 2012 the control year. If differential trends by eligibility drove the result, then we would expect to see similar estimates in the placebo specification. In fact the placebo estimates are an order of magnitude smaller than the main estimate, statistically insignificant and, in the case of the cost curve, wrong signed.

6 Counterfactuals

I use the estimated demand and cost curves to simulate equilibrium, under baseline policy in 2014 and in 2017, and then under alternative repeal or replacement plans: with the mandate or the premium tax credit repealed, or with them replaced by the age-based tax credits of the American Health Care Act of 2016.

I follow the model of Section 3 in these counterfactuals. Recall that equilibrium in the market is characterized by a base price p^b such that $AR(p^b) = AC(p^b)$, where $AR(p^b)$ and $AC(p^b)$ give

[20] Because the sample is limited to the insured, the first stage estimates are slightly different, and I report these estimates in Panel A.

the average revenue and average cost to the insurance company of the insured.[21] Average revenue need not equal p^b because insurers set a base price but revenue is determined by the ages and family structure of enrollees. $AC(p^b)$ differs from the estimated cost curve because the estimated cost curve reflects all non-out-of-pocket medical costs, but $AC(p^b)$ includes loads and reinsurance payments. I model reinsurance as a cost-side subsidy of s_t, equal to \$500 per enrollee in 2014 and zero in 2017. onward.[22] I account for non-medical costs by assuming a load of \$480 per enrollee, roughly equal to the average mark up of premiums over medical costs in 2012.[23] The choice of loads affects the level of prices but has little effect on the counterfactual change in prices as policies change.

All counterfactuals are inherently extrapolations, and so some skepticism is warranted. It is reassuring, however, that the underlying demand and cost curves are based on large price changes, of \$1,100-\$2,600. The counterfactual policy changes result in price changes within this range or not far outside it. Indeed, the effect of removing the premium tax credit is approximated by the comparison of subsidy eligible and ineligible people.[24] The counterfactuals reported here are therefore largely within the span of the data and identifying variation.

6.1 Equilibrium under existing and counterfactual policies

Before turning to counterfactual policies, I first ask whether the model can match the observed prices and quantities in 2014. Although 2014 is in-sample, estimation did not impose any equilibrium conditions, so there is no guarantee that observed prices or quantities equal their equilibrium values. Column (1) of Table 5 reports actual prices and coverage in 2014, and column (2) reports simulated equilibrium values, which are reassuringly close to the data. This concordance suggests that the model can accurately predict equilibrium under alternative policies as well, and so I use it for other simulations.

The table also reports welfare, defined as consumer surplus less total costs. To calculate welfare, I use Equation 11, but I must account for the fact that net prices vary across people because the mandate penalty and the premium tax credit are person specific. To do so, let α_i be the demand

[21] This is a nonlinear equation in p^b. I solve it numerically, using a Newton-type algorithm to search for p^b. In practice the revenue and cost curves are close to linear, so the solution is easy to find.

[22] The legislation called for the reinsurance rate to be cut by 50% in 2015, and eliminated in 2017. I translate the rate to a dollar amount using the average 2014 payout.

[23] I calculate this using revenue and cost data reported to CMS. Insurers must report this information to certify that they have complied with the 80% medical loss ratio target required by the ACA. The data are available from `https://www.cms.gov/CCIIO/Resources/Data-Resources/mlr.html`.

[24] The counterfactuals differ from the estimated effect of subsidies or mandate coverage for a given individual because those effects hold fixed the equilibrium prices, whereas the counterfactuals allow them to adjust.

intercept for person i and γ_i be her cost intercept, assuming linear demand and cost curves. α_i and γ_i vary across people because of the controls included in the estimating equation equations for β and γ. If $\tilde{p}_i$ is i's net price, then her consumer surplus is

$$CS_i(\tilde{p}_i) = 1/2(-\alpha_i/\beta + \tilde{p}_i)(\alpha_i + \beta\tilde{p}_i) \tag{14}$$

and her expected cost to the insurer is

$$AC_i(\tilde{p}_i) = (\gamma_i + \delta\tilde{p}_i + load)(\alpha_i + \beta\tilde{p}_i). \tag{15}$$

Given equilibrium prices, I use these expressions to calculate expected welfare person-by-person as $CS_i(\tilde{p}_i) - AC_i(\tilde{p}_i)D_i(\tilde{p}_i)$; market welfare is the sum of individual welfare. Total welfare in 2014 was about \$2,410 per capita.

I now consider the effect of alternative policies on coverage, prices, and welfare. I define this effect relative to the 2017 baseline policy, which differs from 2014 policy because the mandate penalty grew from 1% of income to 2.5%, and the reinsurance subsidy was eliminated. Column (3) shows predicted coverage and prices under 2017 policy. The net effect of these two changes is that the gross price rises (as the reinsurance subsidy is eliminated), but the price net of the premium tax credit and mandate penalty falls, and coverage rises by about 3 percentage points.

In columns (4), I simulate equilibrium with the mandate removed. Relative to 2017 baseline, this has three effects on coverage. First, there is a direct effect that, for people on the margin of buying insurance, eliminating the mandate makes insurance less appealing, so coverage drops immediately. On average the mandate penalty is about \$900 per person, so the direct effect implies a drop in coverage of about 5.5 percentage points, given the estimated price sensitivity of 6.2 percentage per \$1000. Second, there is a selection effect: when these people exit the market, costs and prices rise, further driving people out and driving prices up. The new equilibrium gross price (equal to the average cost of the insured) is \$540 higher. The third effect of repealing the mandate is that, as gross prices rise, the premium tax credit also rises for subsidy-eligible households. It increases by \$310; this offsets some of the price and coverage effects of the mandate. Thus the overall effect of repealing the mandate on effective prices is to raise them by \$1,130, or about 25% more than the direct effect of just eliminating the mandate. In equilibrium without the mandate, coverage falls about 6 percentage points as marginal enrollees exit the market. Eliminating the mandate reduces welfare per capita also falls, by about \$230 per capita.

In column (5), I instead eliminate the premium tax credit but keep the mandate. Like eliminating the mandate, eliminating the PTC has a direct effect and a selection effect. The average PTC amount is $1,280, so the direct effect of eliminating the PTC is to reduce coverage by about 8 percentage points. In equilibrium, however, gross prices also rise by $560, as healthy people exit the market, worsening selection. The coverage effect of eliminating the PTC is larger than the effect of eliminating the mandate: coverage falls by 11 percentage points when it is eliminated. The PTC has a larger effect than the mandate for two reasons. First, the premium tax credit is larger than the mandate penalty; second, when the PTC is eliminated, there is no offsetting rise in the mandate penalty. Despite the larger fall in coverage, the welfare consequences of eliminating the PTC are not as severe as the consequence of eliminating the mandate. This is because the PTC is most generous for low income families, who have relatively low willingness-to-pay for insurance, meaning that the premium tax credit induces some families to purchase insurance even though their willingness-to-pay is below their cost.

In column (6), I eliminate both the mandate and the PTC. Relative to just eliminating the mandate, gross prices rise by a further $800 and effective prices rise by a further $2,400, nearly triple their 2017 levels. As a result of this large price increases, coverage falls by 21 percentage points relative to baseline (nearly a third) , or 15 percentage points beyond the level with the mandate repealed. Welfare falls further, by about $420 relative to baseline. This is larger than the sum of the welfare losses from eliminating the mandate or PTC in isolation. This is because the welfare loss from adverse selection is not linear but approximately quadratic in the price change change required to correct it.[25]

6.2 Effect of replacement legislation

Thus, repealing without replacement would lead to higher prices, lower coverage, and lower welfare. I consider a possible replacement: the age-based tax credits proposed in the American Health Care Act of 2017 (the AHCA).[26] This plan repeals the mandate and the premium tax credit, and creates a new age based tax credit: $2000 for people younger than 30, $2,500 for 30-39 year-olds, $3,000 for 40-50 year olds, $3,500 for 50-60 year-olds, and $4,000 for people older than 60. Total credit per household is capped at $14,000, and the credits are reduced by 10% for each dollar

[25] To see this, note that the size of the shaded region in Figure 2 is roughly $\Delta W = (Q_1 - Q_0)(P_0 - P - 1)$, where (Q_0, P_0) is the competitive allocation and (Q_1, P_1) the optimal quantity and the price that implements that quantity. With linear demand with slope β, this is $-\beta(p_0 - p_1)^2$.

[26] The full text of the AHCA, as introduced by Rep. Diane Black on March 20, 2017, may be seen here `https://www.congress.gov/bill/115th-congress/house-bill/1628`.

of household income above $75,000 ($150,000 if filing jointly). These tax credits are refundable and, if they exceed the premium of the chosen plan, the excess may be deposited in a health savings account. I model these tax credits as subsidies for insurance purchase.[27]

The AHCA has many other features. For example, it changes the "age-curve" from the 3:1 ratio of the ACA to a 5:1 ratio. It also eliminates the cost-sharing reductions and Medicaid expansions of the ACA. The modelling framework here is well-suited to studying policies that change prices paid by individuals or received by insurers, but it is not suitable for simulating changes to insurance generosity or to the composition of potential insurance purchasers. I therefore abstract from the non-price aspects of the AHCA, although I consider simulations that let the age curve change.[28] Thus these simulations show the effect of replacing the PTC and the mandate with the AHCA's age-based tax credit. They do not show the total effect of replacing the ACA with the AHCA.

I present the replacement simulations in Table 6. Column (1) shows simulated baseline equilibrium in 2017, and column (2) shows the effect of repeal without replacement. Column (3) shows the effect of replacing the PTC and mandate with age-based credits, but without changing the age curve. These credits are fairly generous: on average, nearly $2,400 per person, or $1,100 more than equilibrium premium tax credits. This extra tax credit more than outweighs the consequences of eliminating the mandate. Thus the AHCA replacement on net lower the effective price of insurance, raising coverage by about 2 percentage point. Welfare per capita rises by about $90, in part because of the increase in coverage and in part because the tax credits are now better targeted towards people with high willingness to pay for insurance. Moving to a 5:1 age curve does not substantially change this conclusion. Interestingly, just changing the age curve (but holding fixed the mandate and premium tax credit) reduces the gross price by about $200, consistent with the findings of Tebaldi (2016) that tilting the age curve towards the young can reduce prices by encouraging them to enter the market.

It might be surprising that I find that the age-based tax credits of the AHCA are more generous than the premium tax credits of the ACA. Many recent analyses report that Exchange enrollees would face much higher premiums under the AHCA than under the ACA.[29] These analyses report

[27] This approach treats health savings account balances as though they were cash. These balances accumulate interest free, and after age 65, they can be used for non-health care spending without penalty, so they are valuable even for someone with no health care spending.

[28] I assume that the new age factor, f_{AHCA}, is related to the old age factor f_{ACA}, by the simple transformation $f_{AHCA} = 2f_{ACA} - 1$, following Eibner and Saltzman (2015). I continue to set $f_{child} = 0.635$.

[29] For example, see "House GOP Health Plan Cuts Tax Credits, Raises Costs by Thousands of Dollars for Californians," `http://www.cbpp.org/research/health/house-gop-health-plan-cuts-tax-credits-raises-costs-by-thousands-of-dollars-for` or "Analysis: GOP plan to cost Obamacare enrollees $1,542 more a year," `https://www.vox.com/the-big-idea/2017/3/7/`

the effect of the AHCA on net premiums for *current enrollees* in the Health Insurance Marketplaces. This group is likely to have above-average average tax credits under the ACA, because these credits encourage purchase and because people who buy insurance off-exchange, for example through a broker, are not eligible for the premium tax credit. The overall average amount of the AHCA credits, averaged over all people without employer sponsored insurance or Medicaid, is much larger.

There are several important limitations of these simulations. First, I emphasize that the simulations give the effect of the AHCA's age-based tax credits, not other aspects of the AHCA such as removing the cost sharing reduction subsidies, ending Medicaid, or allowing insurers to price based on pre-existing conditions. These changes could have large effects on coverage or welfare. Second, these simulations give the *average* effects of the AHCA age-based credit. This effect is of course highly heterogeneous by both income and age. Low income people receive a relatively generous subsidy from the premium tax credit, but they would receive much less under the AHCA. Older people face higher base premiums under the ACA (and higher still under the AHCA), so they too receive a relatively large credit under the ACA. They would likely face higher effective prices under the AHCA. Third, these simulations do not permit adjustment along many potentially important margins. In particular, they do not allow income to respond, which could be important as families attempt to reduce their income to below 400% (Heim et al., 2016), and they do not allow insurers to change the characteristics of contracts offered, although Geruso et al. (2016) note that insurers can tailor their formularies to attract high-profit enrollees.

6.3 Robustness

These counterfactual simulations are of course sensitive to both estimation results and arbitrary modelling choices. In Appendix Table C.3 and Appendix Table C.4, I show that the counterfactual simulations are largely robust to alternative approaches.

Robustness of repeal counterfactuals Appendix Table C.3 shows the robustness of the counterfactual effects of repealing the mandate or premium tax credit. I begin by showing that the choice of insurance load—the only parameter that I do not estimate directly from demand and cost data—does not much affect the results. In Panel B, I present simulation results with the load set to $1,000. This number is the largest plausible load, since it implies that at baseline, loads are about 22% of average revenue in 2017; the ACA's medical loss ratio regulation allows for loads only up to 20% of premium revenue. A higher load raises predicted prices and lowers coverage and

14843632/aca-republican-health-care-plan-premiums-cost-price.

welfare, but does not attenuate the coverage or welfare effects of repeal. This is important because subsidizing coverage can lead to welfare losses if there are large non-medical costs of costs (which enrollees may not value).

In Panel C, I consider how the results might change with an alternative market structure. Although it is difficult to incorporate a realistic model of imperfect competition, with product differentiation and Bertrand competition, I simulate an extreme form of imperfect competition: pure monopoly. I find prices and coverage by assuming that a monopolist insurer maximizes profit subject to the ACA's minimum medical loss ratio requirement, which requires that medical costs be at least 80% of claims revenue. In the absence of this constraint, a monopolist could extract unbounded profit, because subsidies increase one for one with posted prices (for subsidy eligible people). Under monopoly, prices are of course higher, and coverage and welfare are lower. However, repealing the premium tax credit and mandate continue to raise prices, by a roughly similar amount. In sum, across the baseline specification and the alternatives, I find that full repeal of the mandate and premium tax credit raises base prices by 35-50%, reduces coverage by about 20 percentage points, and reduces welfare by about $400.

Robustness of AHCA age credits counterfactuals Appendix Table C.4 shows the robustness of the counterfactual effects of repealing the mandate or premium tax credit. I consider a higher load, of $1000, or a monopolist insurer. The simulated effect of the AHCA tax credits or age curve is quite similar across these alternative modelling choices. In particular, a robust result is that replacing the mandate and PTC with the age-based credits of the AHCA undoes most of the effects of repeal. Replacing the mandate and PTC with the AHCA's age-based tax credits credits has a small effect on coverage. With a high load, the combined effect of the age curve and the AHCA tax credits is to reduce the coverage rate by about 5 percentage points.

Counterfactual simulations in the CPS I also considered robustness of the counterfactuals to an alternative data set, the Current Population Survey (CPS). Appendix Tables C.9 and C.10 shows the counterfactual simulations using data and estimates from the CPS. Because I estimate a lower price sensitivity in the CPS, I find smaller, although still substantial, changes in prices, coverage, and welfare in the CPS than in the MEPS. Eliminating just mandate reduces coverage by 3 percentage points and welfare by about $80 per capita; eliminating just the premium tax credit reduces coverage by 6 percentage points and welfare by about $10, and eliminating both reduces coverage by 9 percentage pottns and welfare by about $130 per capita.

7 Conclusion

A central goal of the Affordable Care Act was to improve the operation of the individual insurance market. Prior to the ACA, people were denied coverage because of pre-existing conditions, so the law required both guaranteed issue and community rating. These changes on their own could exacerbate adverse selection, so the ACA also mandated that all individuals buy coverage, and it provided a premium tax credit to make coverage affordable. These reforms are often viewed as a “three-legged stool;” without one of them, the individual insurance market would collapse.

In this paper, I have provided a new assessment of this view. I began by showing that the individual insurance market is adversely selected. As premiums rise, relatively healthy people drop out of the market, driving up the average cost of the insured. Thus it is possible that the mandate increases coverage and welfare. Indeed, simulations reveal that repealing the mandate would reduce coverage and welfare; roughly a third of the total effect of repealing the mandate operates through prices rather than directly through the mandate. Removing both the mandate and the premium tax credit has a very large effect on coverage, reducing it by 20 percentage points. This is much larger than the observed coverage increase of 9 percentage points from 2013 to 2014. This difference suggests that, absent the mandate and premium tax credits, guaranteed issue and community rating could have caused coverage to fall, because of worsened adverse selection.

This simulation shows that repealing without replacement can have large adverse effect on the individual insurance market. I also simulate the effect of replacing the mandate and premium tax credit with a set of age based tax credits, proposed in the American Health Care Act. For the whole population of people at risk to use the individual insurance market, these credits are more generous than the premium tax credit, and so repeal and replace has only a small effect on prices, coverage and welfare. Thus, the mandate by itself is not a necessary condition for the operation of the individual insurance market: sufficiently generous subsidies can replace it. The key is a mechanism to induce healthy people to purchase insurance. Of course, different mechanisms have very different distributional implications. Fully tracing out the incidence of replacing the income-based premium tax credit with an age-based tax credit would be a valuable question for future work.

An important caveat is that the counterfactual simulations of the American Health Care Act’s age credits hold fixed many aspects of the Affordable Care Act, particularly guaranteed issue and community rating. This may seem implausible, since the AHCA would also eliminate community rating and cost sharing reductions, and rollback Medicaid expansions, among other changes. However, changes to the financial side of the ACA—the subsidies—can be passed much more easily than can changes to the regulatory aspect, because Senate rules allow budget reconciliation bills to be passed without the threat of a filibuster. Thus changes to the mandate or tax credits can be passed without bipartisan support. The counterfactuals in this paper may be informative for current policy debate.

References

Abaluck, Jason and Jonathan Gruber, "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," *American Economic Review*, 2011, *101* (4), 1180–1210.

Aizawa, Naoki and You Suk Kim, "Advertising and Risk Selection in Health Insurance Markets," 2015. unpublished working paper.

Akerlof, George A., "The Market for "Lemons:" Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 1970, *84* (3), 488–500.

Azevedo, Eduardo M. and Daniel Gottlieb, "Perfect Competition in Markets with Adverse Selection," *Econometrica*, 2017, (1), 67–105.

Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston, "How does risk selection respond to risk adjustment? New evidence from the Medicare Advantage Program," *American Economic Review*, 2014, *104* (10).

Buchmueller, Thomas and John DiNardo, "Did Community Rating Induce and Adverse Selection Death Spiral? Evidence from New York, Pennsylvania and Connecticut," *American Economic Review*, 2002, *92* (1), 280–294.

Bundorf, M. Kate, Jonathan Levin, and Neale Mahoney, "Pricing and Welfare in Health Plan Choice," *American Economic Review*, 2012, *102* (7), 3214–3248.

Carlin, Caroline and Robert Town, "Adverse Selection, Welfare, and Optimal Pricing of Employer-Sponsored Health Plans," 2008. Unpublished working paper.

Cox, Cynthia, Ashley Semanskee, Gary Claxton, and Larry Levitt, "Explaining Health Care Reform: Risk adjustment, Reinsurance, and Risk Corridors," August 2016. Kaiser Family Foundation Issue Brief.

Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya, "Can Health Insurance Competition Work? Evidence from Medicare Advantage," 2015. NBER Working Paper Number 20818.

Cutler, David M. and Sarah J. Reber, "Paying for Health Insurance: The Trade-Off Between Competition and Adverse Selection," *Quarterly Journal of Economics*, 1998, *113* (2), 1–41.

de Meza, David and David C. Webb, "Advantageous Selection in Insurance Markets," *RAND Journal of Economics*, 2001, *32* (2), 249–262.

DeLeire, Thomas, Andre Chappel, Ken Finegold, and Emily Gee, "Are Low-Income Individuals Aware of and Responsive to Cost-Sharing Subsidies in the Afffordable Care Act's Health Insurance Marketplaces?," 2016. Unpublished working paper.

Eibner, Christine and Evan Saltzman, "Charging Older Adults Higher Premiums Could Cost Taxpayers," September 2015. The Commonwealth Fund.

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

Ericson, Keith M. Marzilli and Amanda Starc, "Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange," *Review of Economics and Statistics*, 2015, *97* (3), 667–682.

Fang, Hanming, Michael P. Keane, and Dan Silverman, "Sources of Advantageous Selection: Evidence from the Medigap Insurance Market," *Journal of Political Economy*, 2008, *116* (2), 303–350.

Finkelstein, Amy and Kathleen McGarry, "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market," *American Economic Review*, 2006, *96* (4), 938–958.

Flood, Sarah, Miriam King, Steven Ruggles, and J. Robert Warren, "Integrated Public Use Microdata Series, Current Population Survey, Version 4.0," 2015. cps.ipums.org.

Frean, Molly, Jonathan Gruber, and Benjamin D. Sommers, "Premium subsidies, the mandate, and Medicaid expansion: Coverage Effects of the Affordable Care Act," *Journal of Health Economics*, 2017, *53*, 72–86.

Geruso, Micahel, Timothy J. Layton, and Daniel Prinz, "Screening in Contract Design: Evidence from the ACA Health Insurance Exchanges," 2016. Available at SSRN: `https://ssrn.com/abstract=2867542`.

Giovannelli, Justin, Kevin W. Lucia, and Sabrina Corlette, "Implementing the Affordable Care Act: State Regulations of Marketplace Plan Provider Networks," May 2015. The Commonwealth Fund.

Gruber, Jonathan, "Health Care Reform Is a "Three Legged Stool:" The Costs of Partially Repealing the Affordable Care Act," August 2010. Center For American Progress, available at `https://www.americanprogress.org/issues/healthcare/reports/2010/08/05/8226/health-care-reform-is-a-three-legged-stool/`.

Hackman, Martin B., Jonathan T. Kolstad, and Amanda E. Kowalski, "Adverse Selection and an Individual Mandate: When Theory Meets Practice," *America Economic Review*, 2015, *105* (3), 1030–1066.

Handel, Benjamin, Igal Hendel, and Michael D. Whinston, "Equilibria in Health Exchanges: Adverse Selection vs. Reclassification Risk," *Econometrica*, 2015, *83* (3), 1261–1313.

Handel, Benjamin R., "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 2013, *103* (7), 2643–2682.

_ and Jonathan T. Kolstad, "Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare," *American Economic Review*, 2015, *105* (8), 2449–2500.

Heim, Bradley T., Gillian Hunter, Adam Isen, Ithai Z. Lurie, and Shanthi P. Ramnath, "Income Responses to the Affordable Care Act: Evidence from the Premium Tax Credit," 2016. Unpublished working paper.

Hendren, Nathan, "Private Information and Insurance Rejections," *Econometrica*, 2013, *81* (5), 1713–1762.

Jaffe, Sonia and Mark Shepard, "Price-Linked Subsidies and Health Insurance Markups," January 2017. `http://www.soniajaffe.com/articles/Jaffe_Shepard_Subsidies.pdf`.

Kowalski, Amanda E., "The Early Impact of the Affordable Care Act State-by-State," *Brookings Papers on Economic Activity*, 2014, pp. 277–333.

Levitt, Larry, Cynthia Cox, and Gary Claxton, "Data Note: How has the Individual Insurance Market Grown Under the Affordable Care Act," May 2015. Kaiser Family Foundation.

Rothschild, Michael and Joseph Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 1976, *90* (4), 629–649.

Royalty, Anne Beeson and Neil Solomon, "Health Plan Choice: Elasticities in a Managed Competition Setting," *Journal of Human Resources*, 1999, *34* (1), 1–41.

Sasso, Anthony T. Lo and Ithai Z. Lurie, "Community rating and the market for private non-group health insurance," *Journal of Public Economics*, 2009, *93*, 264–279.

Shepard, Mark, "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange," 2016. Unpublished wokring paper.

Simon, Kosali Ilayperuma, "Adverse Selection in health insurance markets? Evidence from state small-group health insurance reforms," *Journal of Public Economics*, 2005, *89*, 1865–1877.

Spinnewijn, Johannes, "Heterogeneity, Demand for Insurance, and Adverse Selection," *American Economic Journal: Economic Policy*, 2017, *9* (1), 308–343.

Sydnor, Justin, "(Over)insuring Modest Risks," *American Economic Journal: Applied Economics*, 2010, *2* (4), 177–199.

Tebaldi, Pietro, "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA," 2016. Unpublished working paper.

Uberoi, Namrata, Ken Finegold, and Emily Gee, "Health Insurance Coverage and the Affordable Care Act, 2010-2016," 2016. ASPE Issue brief, available at `https://aspe.hhs.gov/sites/default/files/pdf/187551/ACA2010-2016.pdf`.

Table 1: Summary statistics, by year and insurance status

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Year	2012		2013		2014		DD
Insured?	Yes (59%)	No (41%)	Yes (58%)	No (42%)	Yes (67%)	No (33%)	$\Delta_{2014} - \Delta_{2012-2013}$
Gross premium	3.79	3.52	3.83	3.57	3.83	3.45	0.11
	(1.80)	(1.55)	(1.82)	(1.59)	(1.82)	(1.58)	[0.10]
Net premium	3.79	3.52	3.83	3.57	2.42	1.76	0.39
	(1.80)	(1.55)	(1.82)	(1.59)	(1.68)	(1.33)	[0.09]
Income/FPL*100	445	315	452	309	452	311	4.44
	(320)	(241)	(339)	(219)	(338)	(215)	[16.5]
Age	38.9	37.7	38.5	38.2	39.3	37.2	1.37
	(18.4)	(15.1)	(18.5)	(14.9)	(17.9)	(15.1)	[0.88]
Female	49.6	42.5	48.4	41.9	50.4	38.7	4.9
	(50.0)	(49.4)	(50.0)	(49.3)	(50.0)	(48.7)	[1.9]
Any chronic condition	50.1	39.0	49.4	39.8	48.6	37.4	0.9
	(50.0)	(48.8)	(50.0)	(49.0)	(50.0)	(48.4)	[2.4]
Expected insurer cost	4.07	3.36	3.83	2.57	3.71	2.44	0.29
	(4.04)	(3.35)	(5.22)	(3.67)	(4.83)	(3.67)	[0.20]

Notes: Sample is 2012-2014 MEPS, restricted to people aged 0-64 with income above 138% of FPL, not covered by Medicaid, and not offered employer sponsored health insurance. The full sample consists of 15,636 person-years in 8,656 families. Table reports mean of the indicated variable for the indicated groups. Standard deviations in parentheses. The final column gives the "difference-in-differences" $(E[Y|ins, 2014] - E[Y|unins, 2014]) - (E[Y|ins, pre - 2014] - E[Y|unins, pre - 2014])$ and its standard error (heteroskedasticity robust, clustered on family, in brackets). Expected insurer cost, gross premium, and net premium are measured in thousands of dollars. Female and any chronic conditions are dummy variables equal to zero or 100, so differences are in percentage points.

Table 2: Validity checks for instrumental variables

	(1)	(2)	(3)
Dependent variable	Bad health	Chronic condition	Expected cost
Price coefficient	2.40	3.48	0.19
	(1.67)	(2.55)	(0.17)
Outcome mean	9.69	46.39	3.57
# Observations	12,298	12,320	12,320
# Families	7,109	7,122	7,122

Notes: The sample is described in Table 1, and further limited to people with family income at least 20% from their kink or discontinuity. Table shows the (non-)effect of premiums on the indicated predetermined variables, estimated via 2SLS, for the indicated dependent variable. The excluded instruments are post×aboveKink and post×aboveDisc. Additional controls include post, aboveKink, aboveDisc, discSample, and discSample×post; dummy variables for family composition; quadratics in age, income, and prices; and a quadratic in income interacted with post. Bad health is a dummy variable for self-reported health status fair, poor, or bad, measured in percentage points; the sample in column (1) excludes people with missing self-reported health. Any chronic is a dummy variable for any chronic conditions, also measured in percentage points. Expected insurer cost is predicted insurer expenditures given private information; see text for details. Prices and costs are measured in thousands of dollars. Robust standard errors, clustered on households, in parentheses.

Table 3: Instrumental variables estimates of the demand for insurance

Specification	(1) Base	(2) Dist>10%	(3) Dist>30%	(4) State-Year Fixed Effects	(5) +CSR	(6) Placebo
	Panel A: First stage					
Above Kink × Post	0.83	0.75	0.87	0.76	0.72	0.92
	(0.12)	(0.10)	(0.16)	(0.12)	(0.09)	(0.09)
Above Disc × Post	2.45	2.37	2.49	2.37	2.34	2.56
	(0.16)	(0.13)	(0.20)	(0.15)	(0.12)	(0.14)
F-statistic	378.5	427.8	321	396.2	459.1	405.2
	Panel B: Instrumental variables estimates					
Price	-6.18	-4.70	-8.91	-6.33	-5.78	-0.18
	(2.91)	(2.69)	(3.49)	(2.98)	(2.60)	(3.18)
% Insured	60.1	61.0	58.5	60.1	60.1	56.8
Mean price	3.17	3.21	3.15	3.17	3.17	2.82
Elasticity	-0.33	-0.25	-0.48	-0.33	-0.31	-0.01
# Observations	12,320	14,018	10,604	12,320	12,320	8,654
# Families	7,122	7,909	6,208	7,122	7,122	5,209

Notes: The sample is described in Table 1, and further limited to people with family income at least 20% from their kink or discontinuity, except in columns (2) and (3), where the distance limit is 10% or 30%. The table shows instrumental variables estimates of the demand for insurance. Price is measured in $1000s and insurance coverage is 0 or 100. The excluded instruments are aboveKink×post and aboveDisc×post. Additional controls always include post, aboveKink, aboveDisc, discSample, and discSample×post; dummy variables for family composition; quadratics in age, income, and prices; and a quadratic in income interacted with post. In column (4), I add a full set of state-year fixed effects. In column (5), I include a dummy variable for income below 200% of FPL, and its interaction with post; these proxy for eligibility for CSR subsidies. The elasticity is calculated as $\beta\bar{p}/\bar{y}$, where $\bar{p}$ and $\bar{y}$ are average premium and insurance coverage in sample. Robust standard errors, clustered on households, in parentheses.

Table 4: The effect of price on expected insurer costs

Specification	(1) Base	(2) Dist>10%	(3) Dist>30%	(4) State-Year Fixed Effects	(5) +CSR	(6) Placebo
	Panel A: First stage estimates					
Above Kink X Post	0.79	0.73	0.73	0.98	0.66	0.93
	(0.10)	(0.09)	(0.14)	(0.06)	(0.09)	(0.10)
Above Disc X Post	2.51	2.44	2.39	3.09	2.66	2.74
	(0.15)	(0.13)	(0.19)	(0.12)	(0.13)	(0.15)
F-statistic	326.8	331.5	305.2	289.5	1315.2	637.5
	Panel B: Instrumental variables estimates					
Price	0.44	0.33	0.65	0.20	0.36	-0.05
	(0.20)	(0.17)	(0.25)	(0.08)	(0.18)	(0.20)
Mean expected ins. cost.	3.79	3.78	3.84	3.79	3.79	3.83
Mean price	3.28	3.31	3.26	3.28	3.28	3.09
# People	6,017	6,990	5,012	6,017	6,017	3,938
# Families	3,675	4,155	3,104	3,675	3,675	2,482

Notes: The dependent variable is expected insurer costs of health care. The sample is described in Table 1, and further limited to people with insurance and with family income at least 20% from their kink or discontinuity, except in columns (2) and (3), where the distance limit is 10% or 30%. The table shows instrumental variables estimates of the price sensitivity of demand for insurance. Price is measured in \$1000s and insurance coverage is 0 or 100. The excluded instruments are aboveKink×post and aboveDisc×post. Additional controls always include post, aboveKink, aboveDisc, discSample, and discSample×post; dummy variables for family composition; quadratics in age, income, and prices; and a quadratic in income interacted with post. In column (4), I add a full set of state-year fixed effects. In column (5), I include a dummy variable for income below 200% of FPL, and its interaction with post; these proxy for eligibility for CSR subsidies. Robust standard errors, clustered on households, in parentheses.

Table 5: Simulated effect of repealing the mandate and premium tax credit

Scenario	(1) 2014	(2) 2014 Policy	(3) 2017 Policy	(4) Repeal Mandate	(5) Repeal PTC	(6) Repeal Both
	Data	Equilibrium		Difference vs. 2017 policy		
Mandate penalty ($1000s)	0.27	0.27	0.90	−0.90	0.00	−0.90
Subsidy offered ($1000s)	1.23	1.15	1.28	0.31	−1.28	−1.28
Gross price ($1000s)	3.70	3.58	3.83	0.54	0.57	1.35
Effective price ($1000s)	2.20	2.16	1.65	1.13	1.85	3.53
Coverage rate	0.67	0.65	0.68	−0.06	−0.11	−0.21
Welfare per capita ($1000s)	NA	2.41	2.48	−0.23	−0.10	−0.42

Notes: Prices, tax credits, and welfare are measured in thousands of dollars. Column (1) shows the average of the indicated variables for the 2014 data. Columns (2) and (3) simulate equilibrium under the policy as legislated in 2014 and 2017. Equilibrium changes because of changes in the mandate and reinsurance. Columns (4)-(6) show the difference relative to 2017 of alternative policies. Column (4) simulates equilibrium with the mandate eliminated, but otherwise retaining 2017 policies, column (5) simulates equilibrium with the premium tax credit (PTC) eliminated, and column (6) simulates equilibrium with both the mandate and the PTC eliminated. The equilibrium calculations use demand and cost curves are from column (1) of Tables 3 and 4, and the load is equal to $480 per enrollee per year.

Table 6: Simulated effect of AHCA Tax Credits and Age Curve

Scenario	(1) 2017 Policy	(2) Repeal PTC & Mandate	(3) AHCA Credits	(4) AHCA Age Curve	(5) AHCA Credits & AHCA Curve
	Equilibrium	Difference vs. 2017 policy			
Mandate penalty ($1000s)	0.90	−0.90	−0.90	0.00	−0.90
Subsidy offered ($1000s)	1.28	−1.28	1.10	−0.01	1.10
Unsubsidized price ($1000s)	3.83	1.35	−0.22	−0.19	−0.21
Effective price ($1000s)	1.65	3.53	−0.42	−0.18	−0.41
Coverage rate	0.68	−0.21	0.02	0.01	0.03
Welfare per capita ($1000s)	2.48	−0.42	0.09	−0.01	−0.01

Notes: Prices, tax credits, and welfare are measured in thousands of dollars. Columns (1) simulates equilibrium under the policy as legislated in 2017. Columns (2)-(5) show the difference relative to 2017 of alternative policies. Column (2) simulates equilibrium with both the mandate and the PTC eliminated. Column (3) replaces the mandate and PTC with the age-based tax credits of the AHCA. Column (4) keeps the mandate and PTC but uses the 5:1 age-curve of the AHCA instead of the 3:1 age curve of the ACA, and column (5) replaces the mandate and PTC with the AHCA tax credits, and moves to a 5:1 age curve. The equilibrium calculations use demand and cost curves are from column (1) of Tables 3 and 4, and the load is equal to $480 per enrollee per year.

Figure 1: Expected family contribution, given income

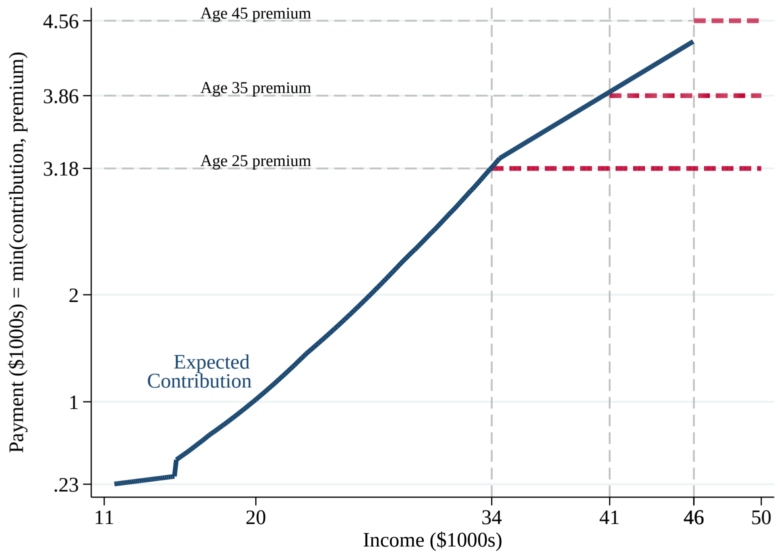

Notes: Figure shows the expected family contribution for a single person, as a function of adjusted gross income. The actual price owed is the minimum of the market price and the expected family contribution. The figure shows market prices for a 25 year-old, 35 year-old, and 45 year-old in Monroe County, Indiana. For 25 and 35 year-olds, prices net of subsidies are kinked in income, but for a 45 year-old, it is discontinuous

Figure 2: Equilibrium with adverse selection

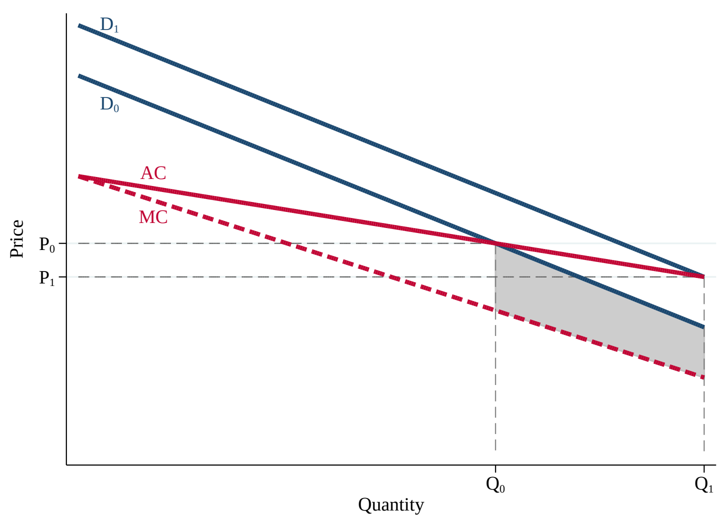

Notes: Figure depicts equilibrium with adverse selection, given by the intersection of AC and D_0. The demand curve always lies above the MC curve, so full coverage is efficient. Equilibrium occurs when $D_0 = AC$, however, and at that price, not everyone purchases insurance. A mandate or subsidy that shifts the demand curve up from D_0 to D_1 results in lower prices and greater coverage when the average cost curve is downward sloping, because it draws healthy people into the market. The shaded area represents the welfare loss from adverse selection and the welfare gain from fully correcting it.

Online Appendices

A Non-measurement of pre-period prices does not lead to bias

In this appendix, I present a simple model to show why unobserved base prices in the pre-ACA period are not a problem for identification. As in the main text, let $\tilde{p}_{it}$ be the per-person price of insurance net of the subsidy and mandate, and let p_{it} be the gross price, which may be unobserved. By definition,

$$\tilde{p}_{it} = p_{it} + s_{it}, \tag{16}$$

where s_{it} is the sum of the PTC payment and mandate penalty for non-coverage. After 2014, I observe $\tilde{p}_{it}, p_{it}$ and s_{it}. Before 2014, I do not observe p_{it} or $\tilde{p}_{it}$, but s_{it} is known to be zero.

Substituting this of $\tilde{p}_{it}$ into the estimating equations, Equation 12 and Equation 13, and rearranging, we obtain

$$\begin{aligned} y_{it} =& \beta_0 + \beta_1 s_{it} & (17)\\ &+ \beta_2 aboveKink_{it} + \beta_3 aboveDisc_{it} + \beta_4 discSample_i \\ &+ \beta_5 post_{it} + \beta_6 discSample_{it} \times post_{it} + [\beta_1 p_{it} + \varepsilon_{it}] \\ s_{it} =& \alpha_0 + \alpha_1 aboveKink_{it} \times post_{it} + \alpha_2 aboveDisc_{it} \times post_{it} & (18)\\ &+ \alpha_3 aboveKink_{it} + \alpha_4 aboveDisc_{it} + \alpha_5 discSample_i \\ &+ \alpha_6 post_{it} + \alpha_7 discSample_{it} \times post_{it} + [-p_{it} + \eta_{it}] \end{aligned}$$

These regressions show that the actual estimating equations, which relate demand or cost to net prices, are equivalent to an alternative set of estimating equations equation, which relate demand or cost to the *subsidy*, except with a new error term, in brackets. The new error term contains the price gross of the subsidy. Instrumental variables estimates of Equation 17 yield consistent estimates of price sensitivity β_1 under the identification assumption that the instruments ($post \times aboveKink$ and $post \times aboveDisc$) are uncorrelated with p_{it} and ε, conditional on the controls.

It may seem that the requirement that the instruments be uncorrelated with p_{it} adds an identification assumption. However, this requirement is automatically satisfied. Recall that p_{it} varies across areas because of differences in benchmark premia across areas, but it does not vary over time. The only way p_{it} could be correlated with $post \times aboveKink$ or $post \times aboveDisc$ is if p_{it} is correlated with $aboveKink$ or $aboveDisc$. But I control for these this coefficient is equal to the price sensitivity. Thus by construction, the residual variation in the instruments is uncorrelated with p_{it}. In sum, it is not a source of bias that pre-period *prices* are unobserved, because pre-period *subsidies* are observed (they are equal to zero), and identification relies on changes in subsidies.

B Description of premium data

I assembled a dataset with the benchmark premium of each county in 2014, that is the premium of the second lowest cost silver plan (SLCSP) for each county. I began with the public use files from `data.healthcare.gov`, which provide information about every plan offered in every county among states with Federally Facilitated Marketplaces or partnerships with HHS. This information includes the plan's metal level and premium, so it is straightforward to find the benchmark premium for each county.

For the 14 states and DC that did not use `healthcare.gov`, I collected data in a variety of ways. I began by visiting state exchange websites in the fall of 2016. Some, like California, provide easily accessibly information about the prices of all plans in all counties. Others, like Washington state, provide booklets with information about all plans, from which it is possible to recover premiums. Still others, like Washington, D.C., do not directly provide historical premium data at all. However, because the benchmark premium is necessary for calculating tax credits, all state exchanges enable taxpayers to look up their own benchmark premium. When this was the only option for determining benchmark premiums, I looked up the benchmark premium for each rating area in the state, and assigned this premium to each county in the rating area. I was unable to determine benchmark premiums for Kentucky, because the state website required a Social Security number of a Kentucky resident before it would disclose premiums. Phone calls to the exchange operator were unhelpful. I omit Kentucky from the analysis sample.

C Replication with CPS data

I replicate all analyses reported in this paper using data from the Current Population Survey (as prepared by Flood et al. (2015)). Relative to the MEPS, the CPS suffers from three major disadvantages: it lacks any measure of total healthcare spending, it has incomplete geographic identifiers, and the survey question underlying the main insurance variable changed in 2013.[30] Because of these disadvantages, I focus on the MEPS results. Nonetheless I also analyze the CPS data because the CPS has much larger sample sizes and because it is a second, independent sample, so replication in the CPS provides reassurance that any statistically significant results in the MEPS are unlikely to be due to data mining or specification searching.

The replication keeps the sample, specification, and variable definition as similar as possible, but differences between the CPS and MEPS makes exact replication impossible. Here I highlight the important differences. First, the CPS sample is limited in the same way as the MEPS sample (Marketplace eligible, far from kink/discontinuity), but instead of excluding families with offers of employer sponsored insurance (which is not ascertained in the CPS), I exclude families in which anyone reports having employer sponsored insurance. I also exclude families with missing county identifiers; these families generally live in smaller counties, which are suppressed to preserve confidentiality, and I exclude insurance year 2012 because of the change in insurance definition in 2013. Second, the CPS does not measure health care spending. It does, however, ask about health status. The question wording is the same as in MEPS, but in the CPS, the respondent answers for everyone in the household. I impute expected insurer costs in the CPS using the regression estimates from the MEPS reported in column (1) of Table C.1.

Table C.5 shows summary statistics for the Marketplace sample in the CPS. Like the MEPS, the CPS data show a large increase in the insurance rate in 2014, but no obvious worsening of the risk pool. Older people were not especially likely to gain insurance, and the expected costs of the insured did not rise, relative to the uninsured expected costs. The differences are narrower in the CPS than in the MEPS, consistent with the fact that the measure of expected spending allows for less private information in the CPS.

Table C.6 shows the validity checks for the instruments in the CPS. Price has no effect on bad health, or expected insurer costs. The instruments appear valid in the CPS as well. Table C.7 shows the first stage and IV estimates of the effect of price on insurance demand. The point estimates and elasticities are somewhat smaller than in the MEPS, with elasticities around -0.2. The differences

[30] In particular, the question changed in survey year 2014, which refers to insurance coverage year 2013. Before the change, the CPS asked respondents whether they had coverage in the past calendar year. The CPS insurance module takes place in March, and the "past calendar year" may have been misinterpreted as "the past 12 months." To avoid ambiguity, beginning in the 2014 survey, CPS enumerators began asking about coverage in each month of the past calendar year (December, November, etc.).

between the samples are generally not statistically significant, although the sampling frames are different (the CPS sample excludes low population counties, in which premiums cannot be imputed, because the county identifiers are suppressed).

Table C.8 shows the effect of prices on expected insurer costs. A $1000 increase in prices raises the expected costs of the insured by about $190, also smaller than but not significantly different from the MEPS estimate. Interestingly, the CPS estimates imply that a price change that reduces the coverage rate by one percentage point increases the average cost of the insured by about $60, quite similar to the MEPS estimate.

Table C.9 reports the results of the counterfactual experiments on repealing the mandate or the premium tax credit, performed now in the CPS using the CPS estimates from column (1). Repealing the mandate reduces the insurance rate by about 3 percentage points, driving up costs and prices as relatively healthy people exit the market, so the gross price increases by about 5%. Repealing the mandate as well as the premium tax credit would have a larger effect, reducing coverage by 9 percentage points, raising prices by 13 percent, and reducing welfare per capita by about $130. These effects are similar to the estimates from the MEPS, but generally smaller because the CPS estimated price sensitivity is lower. This means that there are smaller direct effects and smaller equilibrium effects. Table C.10 shows the simulated effect of replacing the mandate and premium tax credit with the AHCA age credits and age curve. As in the MEPS, I find that in the CPS, replacing the PTC and the mandate with the age credits undoes essentially all of the effects of repeal. Coverage, prices, and welfare change very little.

Table C.1: Regression for imputing expected insurer spending

Outcome = insurer spending	Coefficient	Standard Error
Income/FPL * 100	-0.87	(0.80)
High school	-1331	(1165)
Some college	1612	(1107)
College or more	35	(1133)
Race = black	284	(855)
Age 6-17	125	(1682)
Age 18-30	-744	(7410)
Age 31-40	-738	(7418)
Age 41-50	-43	(7406)
Age 51-60	2490	(7398)
Age 60-64	2754	(7383)
Female	566	(498)
Employed	-2465	(619)
Health very good	-272	(613)
Health good	1835	(701)
Health Fair	5509	(1100)
Health poor	12367	(1841)
Constant	3958	(7459)
R^2	0.08	
Sample size	2,648	

Notes: Sample is as described in Table 1, but further limited to insured people in 2014. Table reports estimates from a regression of insurer medical spending on the indicated variables. Additional regressors, not shown, include: dummy variables for missing education, race, health status for missing education, missing. Robust standard errors clustered on household in parentheses

Table C.2: Summary statistics, various samples

Sample	Full MEPS		Analysis sample		Estimation sample	
	Mean	(SD)	Mean	(SD)	Mean	(SD)
Insured	0.86	(0.35)	0.61	(0.49)	0.60	(0.49)
Age	37.96	(22.82)	38.45	(17.11)	38.87	(17.09)
Female	0.51	(0.50)	0.46	(0.50)	0.47	(0.50)
Child	0.28	(0.45)	0.15	(0.36)	0.15	(0.36)
Income/FPL	360	(331)	396	(304)	407	(344)
# Observations	104,743		15,636		12,320	
# Families	38,102		8,656		7, 122	

Notes: Sample is 2012-2014 MEPS. The analysis sample is further restricted to people aged 0-64 with income above 138% of FPL, not covered by Medicaid, and not offered employer sponsored health insurance. The estimation sample is further restricted to people with household incomes at least 25% away from their kink or discontinuity in premium.

Table C.3: Robustness of simulated effects of repealing the mandate or PTC

Scenario	(1) 2014 Policy	(2) 2017 Policy	(3) Repeal Mandate	(4) Repeal PTC	(5) Repeal Both
	Panel A: Baseline parameterization				
	Equilibrium		Difference vs. 2017 policy		
Mandate penalty ($1000s)	0.27	0.90	−0.90	0.00	−0.90
Subsidy offered ($1000s)	1.15	1.28	0.31	−1.28	−1.28
Gross price ($1000s)	3.58	3.83	0.54	0.57	1.35
Effective price ($1000s)	2.16	1.65	1.13	1.85	3.53
Coverage rate	0.65	0.68	−0.06	−0.11	−0.21
Welfare per capita ($1000s)	2.41	2.48	−0.23	−0.10	−0.42
	Panel B: Set load to $1000				
	Equilibrium		Difference vs. 2017 policy		
Mandate penalty ($1000s)	0.27	0.90	−0.90	0.00	−0.90
Subsidy offered ($1000s)	1.50	1.64	0.31	−1.64	−1.64
Gross price ($1000s)	4.22	4.47	0.53	0.78	1.59
Effective price ($1000s)	2.44	1.92	1.13	2.42	4.14
Coverage rate	0.64	0.67	−0.07	−0.15	−0.26
Welfare per capita ($1000s)	2.01	2.07	−0.21	−0.09	−0.39
	Panel C: Monopolist insurer				
	Equilibrium		Difference vs. 2017 policy		
Mandate penalty ($1000s)	0.27	0.90	−0.90	0.00	−0.90
Subsidy offered ($1000s)	1.40	1.59	0.40	−1.59	−1.59
Gross price ($1000s)	4.04	4.37	0.69	0.99	2.09
Effective price ($1000s)	2.36	1.88	1.19	2.58	4.58
Coverage rate	0.64	0.67	−0.07	−0.16	−0.28
Welfare per capita ($1000s)	2.36	2.43	−0.27	−0.2	−0.61

Notes: Table shows simulated equilibrium under the policy as legislated in 2014 and 2017 in columns (1) and (2), and the change relative to 2017, for the indicated policy, in columns (3)-(5). In column (3), the mandate is eliminated but not the premium tax credit (PTC), in column(4) PTC is eliminated but not the mandate, and in column (5) both are eliminated. Panel A shows the baseline: demand and cost curves are from column (4) of Tables 3 and 4, and the load is equal to $480 per enrollee per year. Panel B changes the baseline by calibrating the load to an extreme value of $1000. Panel C simulates equilibrium at baseline parameters but assuming a monopolist insurer, rather than a competitive equilibrium.

Table C.4: Robustness of simulated effects of AHCA age credits and age curve

Scenario	(1) 2017 Policy	(2) Repeal PTC & Mandate	(3) AHCA Credits	(4) AHCA Age Curve	(5) AHCA Credits & Curve
	Panel A: Baseline parameterization				
	Equilibrium	Difference vs. 2017 policy			
Mandate penalty ($1000s)	0.90	−0.90	−0.90	0.00	−0.90
Subsidy offered ($1000s)	1.28	−1.28	1.10	−0.01	1.10
Gross price ($1000s)	3.83	1.35	−0.22	−0.19	−0.21
Effective price ($1000s)	1.65	3.53	−0.42	−0.18	−0.41
Coverage rate	0.68	−0.21	0.02	0.01	0.03
Welfare per capita ($1000s)	2.49	−0.43	0.10	0.05	0.20
	Panel B: Set load to $1000				
	Equilibrium	Difference vs. 2017 policy			
Mandate penalty ($1000s)	0.90	−0.90	−0.90	0.00	−0.90
Subsidy offered ($1000s)	1.64	−1.64	0.74	−0.03	0.74
Gross price ($1000s)	4.47	1.59	0.00	−0.13	0.56
Effective price ($1000s)	1.92	4.14	0.17	−0.10	0.73
Coverage rate	0.67	−0.26	−0.01	0.00	−0.05
Welfare per capita ($1000s)	2.07	−0.39	0.08	0.01	0.07
	Panel C: Monopolist insurer				
	Equilibrium		Difference vs. 2017 policy		
Mandate penalty ($1000s)	0.90	−0.90	−0.90	0.00	−0.90
Subsidy offered ($1000s)	1.59	−1.59	0.79	−0.09	0.79
Gross price ($1000s)	4.37	2.09	−0.07	−0.26	−0.09
Effective price ($1000s)	1.88	4.58	0.03	−0.17	0.01
Coverage rate	0.67	−0.28	0.00	0.01	0.00
Welfare per capita ($1000s)	2.43	−0.61	0.08	0.04	0.16

Notes: Table shows simulated equilibrium under the policy as legislated in 2017 in columns (1), and the change relative to 2017, for the indicated policy, in columns (2)-(5). In column (2) both the mandate and PTC are eliminated. In column (3), they are replaced with the AHCA age credits. In column (4), I keep the mandate and PTC but change the age curve from 3:1 to 5:1. In column (5) I eliminate the mandate and PTC, replace them with the AHCA age credits, and change the age curve. Panel A shows the baseline: demand and cost curves are from column (4) of Tables 3 and 4, and the load is equal to $480 per enrollee per year. Panel B changes the baseline by calibrating the load to an extreme value of $1000. Panel C simulates equilibrium at baseline parameters but assuming a monopolist insurer, rather than a competitive equilibrium.

Table C.5: Summary statistics, by year and insurance status, CPS

	(1)	(2)	(3)	(4)	(5)
Year	2013		2014		DD
Insured?	Yes (51%)	No (49%)	Yes (63%)	No (37%)	$\Delta_{2014}-\Delta_{2012-2013}$
Gross premium	3.4	3.4	3.5	3.5	0.05
	(1.6)	(1.5)	(1.7)	(1.5)	[0.07]
Net premium	3.4	3.4	2.1	1.8	0.22
	(1.6)	(1.5)	(1.5)	(1.3)	[0.07]
Income/FPL∗100	486.7	342.6	448.1	364.2	-60.20
	(612.3)	(434.8)	(480.9)	(413.9)	[24.91]
Age	34.4	36.5	34.6	36.1	0.54
	(18.1)	(15.5)	(17.9)	(15.2)	[0.67]
Female	48.5	45.1	49.0	43.9	1.67
	(0.5)	(0.5)	(0.5)	(0.5)	[1.47]
Expected insurer cost	3.8	3.3	3.6	3.2	-0.06
	(4.1)	(3.4)	(3.7)	(3.2)	[0.14]

Notes: Sample is 2013-2014 CPS, restricted to people aged 0-64 with income above 138% of FPL, not covered by Medicaid, with no one in the family receiving employer sponsored health insurance. The full sample consists of 19,803 person-years in 12,547 families. Table reports mean of the indicated variable for the indicated groups. Standard deviations in parentheses. The final column gives the "difference-in-differences" $(E[Y|ins, 2014] - E[Y|unins, 2014]) - (E[Y|ins, 2013] - E[Y|unins, 2013])$ and its standard error (heteroskedasticity robust, clustered on family, in brackets). Expected insurer cost, gross premium, and net premium are measured in thousands of dollars. Female is a dummy variable equal to zero or 100, so differences are in percentage points.

Table C.6: Validity checks for instrumental variables, CPS

	(1)	(2)
Dependent variable	Bad health (0-100)	Expected insurer cost ($1000s)
Price coefficient	0.49	0.07
	(0.77)	(0.09)
Outcome mean	8.15	3.63
# Observations	15,260	15,260
# Families	9,243	9,243

Notes: The sample is described in Table C.5, and further limited to people with family income at least 20% from their kink or discontinuity. Table shows the (non-)effect of premiums on the indicated predetermined variables, estimated via 2SLS, for the indicated dependent variable. The excluded instruments are aboveKink×post and aboveDisc×post. Additional controls always include post, aboveKink, aboveDisc, discSample, and discSample×post; dummy variables for family composition; quadratics in age, income, and prices; and a quadratic in income interacted with post. Bad health is a dummy variable for self-reported health status fair, poor, or bad, measured in percentage points. Expected spending is predicted insurer spending given private information; see text for details. It is measured in thousands of dollars. Robust standard errors, clustered on households, in parentheses.

Table C.7: Instrumental variables estimates of the demand for insurance, CPS

Specification	(1) Base	(2) Dist>10%	(3) Dist>30%	(4) State-Year Fixed Effects	(5) +CSR	(6) Placebo
	Panel A: First stage					
Above Kink ×Post	1.33	1.17	1.46	1.31	0.93	1.26
	(0.03)	(0.03)	(0.04)	(0.03)	(0.04)	(0.04)
Above Disc × Post	2.87	2.65	3.02	2.86	2.67	2.67
	(0.06)	(0.06)	(0.07)	(0.06)	(0.06)	(0.08)
F-statistic	1320.7	1190.1	1278	1304	1079.1	659.0
	Panel B: Instrumental variables estimates					
Price	-3.34	-2.51	-4.75	-3.13	-3.22	0.62
	(1.74)	(1.70)	(1.97)	(1.75)	(1.94)	(1.97)
Elasticity	-0.18	-0.136	-0.256	-0.169	-0.174	0.029
# Observations	15,260	17,624	13,076	15,260	15,260	15,614
# Families	9,243	10,694	7,886	9,243	9,243	9,594

Notes: The sample is described in Table C.5, and further limited to people with family income at least 20% from their kink or discontinuity, except in columns (2) and (3), where the distance limit is 10% or 30%. The table shows instrumental variables estimates of the price sensitivity of demand for insurance. Price is measured in $1000s and insurance coverage is 0 or 100. The excluded instruments are aboveKink×post and aboveDisc×post. Additional controls always include post, aboveKink, aboveDisc, discSample, and discSample×post; dummy variables for family composition; quadratics in age, income, and prices; and a quadratic in income interacted with post. In column (4), I add a full set of state-year fixed effects. In column (5), I include a dummy variable for income below 200% of FPL, and its interaction with post; these proxy for eligibility for CSR subsidies. The placebo specification in column (6) treats 2013 as the post period and 2012 as the pre period. The elasticity is calculated as $\beta\bar{p}/\bar{y}$, where $\bar{p}$ and $\bar{y}$ are average premium and insurance coverage in sample. Robust standard errors, clustered on households, in parentheses.

Table C.8: The effect of price on expected insurer costs, CPS

Specification	(1) Base	(2) Dist>10%	(3) Dist>30%	(4) State-Year Fixed Effects	(5) +CSR	(6) Placebo
	Panel A: First stage estimates					
Above Kink × Post	1.34	1.16	1.49	1.33	0.90	-0.19
	(0.04)	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)
Above Disc × Post	2.97	2.74	3.14	2.95	2.73	1.56
	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.11)
F-statistic	843.80	724.10	843.10	822.90	646.80	134.9
	Panel B: Instrumental variables estimates					
Price	0.19	0.17	0.13	0.20	0.13	-0.01
	(0.11)	(0.10)	(0.12)	(0.12)	(0.13)	(0.13)
# Observations	8,587	9,963	7,310	8,587	8,587	7,450
# Families	5,150	5,962	4,371	5,150	5,150	4,561

Notes: The dependent variable is expected insurer costs of health care. The sample is described in Table C.5, and further limited to people with insurance and with family income at least 20% from their kink or discontinuity, except in columns (2) and (3), where the distance limit is 10% or 30%. The table shows instrumental variables estimates of the price sensitivity of demand for insurance. Price is measured in $1000s and insurance coverage is 0 or 100. The excluded instruments are aboveKink×post and aboveDisc×post. Additional controls always include post, aboveKink, aboveDisc, discSample, and discSample×post; dummy variables for family composition; quadratics in age, income, and prices; and a quadratic in income interacted with post. In column (4), I add a full set of state-year fixed effects. In column (5), I include a dummy variable for income below 200% of FPL, and its interaction with post; these proxy for eligibility for CSR subsidies. The placebo specification in column (6) treats 2013 as the post period and 2012 as the pre period. Robust standard errors, clustered on households, in parentheses.

Table C.9: Simulated effect of repealing the mandate and premium tax credit (CPS)

Scenario	(1) 2014 Policy	(2) 2017 Policy	(3) Repeal Mandate	(4) Repeal PTC	(5) Repeal Both
	Equilibrium		Difference vs. 2017 policy		
Mandate penalty ($1000s)	0.27	0.88	−0.88	0.00	−0.88
Subsidy offered ($1000s)	1.31	1.56	0.13	−1.56	−1.56
Unsubsidized price ($1000s)	3.64	4.04	0.22	0.26	0.52
Effective price ($1000s)	2.06	1.61	0.96	1.81	2.95
Coverage rate	0.61	0.62	−0.03	−0.06	−0.09
Welfare per capita ($1000s)	4.51	4.53	−0.08	−0.01	−0.13

Notes: Prices, tax credits, and welfare are measured in thousands of dollars. Column (1) shows the average of the indicated variables for the 2014 data. Columns (2) and (3) simulate equilibrium under the policy as legislated in 2014 and 2017. Equilibrium changes because of changes in the mandate and reinsurance. Columns (4)-(6) show the difference relative to 2017 of alternative policies. Column (4) simulates equilibrium with the mandate eliminated, but otherwise retaining 2017 policies, column (5) simulates equilibrium with the premium tax credit (PTC) eliminated, and column (6) simulates equilibrium with both the mandate and the PTC eliminated. The equilibrium calculations use demand and cost curves are from column (1) of Tables C.7 and C.8, and the load is equal to $480 per enrollee per year.

Table C.10: Simulated effect of AHCA Tax Credits and Age Curve (CPS)

Scenario	(1) 2017 Policy	(2) Repeal PTC & Mandate	(3) AHCA Credits	(4) AHCA Age Curve	(5) AHCA Credits & AHCA Curve
	Equilibrium	Difference vs. 2017 policy			
Mandate penalty ($1000s)	0.88	−0.88	−0.88	0.00	−0.88
Subsidy offered ($1000s)	1.56	−1.56	0.75	0.13	0.75
Gross price ($1000s)	4.04	0.52	0.01	−0.04	0.33
Effective price ($1000s)	1.61	2.95	0.13	−0.17	0.45
Coverage rate	0.62	−0.09	0.00	0.01	−0.01
Welfare per capita ($1000s)	4.53	−0.13	0.03	−0.02	0.00

Notes: Prices, tax credits, and welfare are measured in thousands of dollars. Columns (1) simulates equilibrium under the policy as legislated in 2017. Columns (2)-(5) show the difference relative to 2017 of alternative policies. Column (2) simulates equilibrium with both the mandate and the PTC eliminated. Column (3) replaces the mandate and PTC with the age-based tax credits of the AHCA. Column (4) keeps the mandate and PTC but uses the 5:1 age-curve of the AHCA instead of the 3:1 age curve of the ACA, and column (5) replaces the mandate and PTC with the AHCA tax credits, and moves to a 5:1 age curve. The equilibrium calculations use demand and cost curves are from column (1) of Tables C.7 and C.8, and the load is equal to $480 per enrollee per year.