Does Medicaid Make Private Health Insurance Cheaper?

Jonathan Holmes

April 6, 2021

Please click here for the most recent version of the paper.

Models of consumer choice and selection in health insurance markets typically focus on decisions within one market setting (e.g. the choice of plans offered by an employer). This paper focuses on selection and choice at the broader level of health care markets, comparing Medicaid, employer sponsored insurance, and insurance exchanges. I study cross-market adverse selection in the context of the Affordable Care Act, which simultaneously expanded Medicaid while creating health insurance exchanges. Using state decisions to expand Medicaid and variation in federal reinsurance payments as identification, I find that Medicaid expansion reduced prices on individual health insurance exchanges by an average of $504 (9.7%) per person, per year. These price reductions on the individual market represent an indirect benefit of Medicaid expansion that is equal to 9% of the total costs of expansion, and increased enrollment of higher income individuals (above 400% of the Federal Poverty Level) on the individual market by 10%. I then estimate a structural demand model with rich individual heterogeneity across consumers using a combination of public and administrative datasets. This model shows that, while the 2014 Medicaid expansion reduced prices in private markets, further expansions to Medicaid would have the opposite effect because they target a healthier population.

JEL: H51, I11, I13

In health insurance markets, adverse selection occurs when consumers with more costly medical conditions are more likely to purchase some insurance product. Existing models of adverse selection generally assume all insurance is sold in a single market (e.g. Einav, Finkelstein, and Cullen (2010); Handel, Hendel, and Whinston (2015); Tebaldi, Torgovitsky, and Yang (2019)).[1] But, in the United States, there exist multiple public and private health insurance markets, including Medicaid, Medicare, employer-sponsored insurance, and the health insurance exchanges. Average disease burden, the price of coverage, and plan characteristics vary widely across these different markets. This paper asks: Is adverse selection *across* health insurance markets relevant for U.S. policy?

In particular, I study how adverse selection affects our interpretation of the impacts of Medicaid expansion, enacted as a part of the Affordable Care Act (ACA). The ACA increased health insurance coverage in a piece-meal approach, by both expanding Medicaid eligibility to a larger group of low-income Americans, and by creating individual health insurance exchanges for those with higher incomes. If there is selection across these two markets, then Medicaid expansion can lower premiums on the individual market by reducing the number of high-risk enrollees who choose to purchase individual insurance. However, because the Affordable Care Act made so many changes to insurance markets simultaneously, it is difficult to identify the what changes were driven specifically by Medicaid expansion. Moreover, because health insurance data tends to be siloed, it is rare to have access to comprehensive data that covers all insurance markets simultaneously.

This paper addresses these challenges using an identification strategy based on arbitrary features of two major policy changes: (1) state-level differences in the timing of Medicaid expansion, and (2) large price variation on individual exchanges driven by federal reinsurance payments. These two sources of identification allow me to trace out a demand curve for individual insurance and measure the degree of substitution across the individual exchanges, Medicaid, and employer-sponsored coverage. My preferred specification is a discrete choice model of consumer demand which allows for rich heterogeneity across different demographic and income groups. To estimate this model, I combine national data from surveys and regulatory filings with detailed claims data from the Utah All-Payer Claims.

I find that cross-market adverse selection is very important to measuring the benefits of Medicaid expansion. In non-expansion states, I find that 34% of enrollees on the individual market would have enrolled in Medicaid if their state had expanded Medicaid. Because these enrollees were adversely selected, I find that expansion states had premiums that were $504 lower (a 9.7% reduction) in the individual market. These premium reductions represent indirect benefits of Medicaid expansion worth 9% of the cost of expanding Medicaid, whose incidence fell on individuals purchasing individual plans and the federal government that

[1] Other studies also look at adverse selection in who purchases health insurance at all (e.g. Hackmann, Kolstad, and Kowalski (2015)), or adverse selection on both of these margins (e.g. Azevedo and Gottlieb (2017); Geruso, Layton, McCormack, and Shepard (2019))

paid less in subsidies. Nevertheless, in policy simulations, I find that any further expansion of Medicaid (for example, by increasing the cutoff for Medicaid eligibility from 138% of the Federal Poverty Level to 200% of FPL) would actually increasing premiums in the individual market rather than decrease them, because this additional expansion targets a healthier population. Together, these facts suggest policymakers should account for adverse selection when designing health insurance policies, and that simulating the effects of these policies requires knowledge of the joint distribution of consumer demand across markets and health risk.

At the core of my empirical approach is a theoretical model which shows how adverse selection operates across markets. Existing models of adverse selection show how adverse selection is driven by variation in demand among consumers (e.g. Einav et al. (2010)). In my cross-market model, adverse selection is driven by variation in both consumer demand and *eligibility* for insurance. In the case of Medicaid expansion, new Medicaid enrollment crowds-out enrollment in private markets, such as the individual market. If those who switch from the individual market to Medicaid are adversely selected, then Medicaid expansion would reduce prices in the individual market. Lower prices on the individual market could then lead to an increase in coverage among those ineligible for Medicaid. The reduced-form estimate of an adversely selected Medicaid expansion on individual market thus combines two effects: Crowd-out from expanded Medicaid eligibility, and crowd-in from lower prices.

To measure the crowd-out effect of Medicaid, I first perform a descriptive analysis where I compare states which expanded Medicaid in 2014 to those that did not. While Medicaid expansion is an imperfect natural experiment,[2] this analysis clearly shows that Medicaid expansion crowded out private coverage, and that expansion states have a healthier population on their individual markets. In expansion states, I find that the total size of the private market decreased by 34% in the individual market and by 2% in the employer-sponsored market, relative to non-expansion states. Using Medical Loss Ratio data on insurance claims, I find that enrollees on individual insurance exchanges in non-expansion states incurred 10% more health insurance spending relative to non expansion states, which ultimately led to higher premiums relative to non-expansion states.

Next, to estimate how lower prices on the individual market affect demand, I estimate the aggregate demand for health insurance on aggregate demand in the individual market at the state level. To predict price changes on the individual market, I use arbitrary changes in premiums driven by the phase-out of federal reinsurance and risk corridor programs in 2017. Federal reinsurance and risk corridor programs payments varied from an average of $140 per enrollee in Rhode Island to $2200 per enrollee in South Dakota.

[2]Selection into Medicaid expansion was largely driven by partisan goals: All non-expansion states were either fully or partially controlled by Republicans in 2014

Using an event-study design estimated at the state level, I show that when these payments were phased-out in 2017, premiums rose one-for-one on individual exchanges. I then find that, for every $1000 increase in prices on the individual market, total enrolment on the individual market drops by 19%. Incorporating this demand elasticity, a back-of-the-envelope estimate suggests that for every 3 enrollees who are "crowded out" on the individual market in expansion states due to expanded eligibility the individual market to Medicaid, there is one who "crowds-in" in 2017 from lower prices.

While these state-level analyses give a picture of substitution across health insurance markets at the aggregate level, they mask substantial heterogeneity across individuals of different demographic and income groups. Estimating a model which accounts for this heterogeneity is necessary to simulate the effects of new policies. My choice of model is a standard discrete choice model with logit errors. Consumer demand for health insurance is estimated using a flexible functional form that is smooth in both income and demographics. Both prices and individual-level choice sets vary within a market.

This model is identified based on differences across individuals in their exposure to national policy changes. I estimate this model using data from 2015-2017, which is a period of relative stability in health policy, with two major exceptions. First, 5 states expanded Medicaid during this period, which led to large changes in enrollment across markets in these states. Second, prices on the individual market rose dramatically and predictably in 2017 as a result of the phase-out of the federal reinsurance and risk corridor programs. In this model, individual demand is identified by differences across individuals in their exposure to these large policy changes. In the case of Medicaid expansion, the affected population were those in a specific income band (mainly childless adults earning under 138% of the Federal Poverty Line), while other groups were unaffected. In the case of individual-market prices, many low and middle income consumers pay a much lower price than the sticker price, and are fully insulated against price rises. The model implicitly uses these unaffected groups to control for aggregate changes occurring in the economy or in insurance markets at the state level.

Using this model, I then consider the following question: Given Medicaid expansion had spillover benefits to the individual market, would an even larger expansion have had even larger benefits? To answer this question, I use the model to simulate other potential cutoffs for Medicaid eligibility. What I find is that the 138% of FPL cutoff established as a part of the ACA is very close to the optimal, when considering the spillover effects on other markets. Any additional expansion would end up increasing the price of coverage on individual markets. This is because the population earning 138-200% of FPL is much more positively selected than the population earning under 138% of FPL that got access to Medicaid under the ACA. Thus, understanding the role of adverse selection critically depends on identifying who is at the margin with respect to policy changes.

This paper is among a large literature since Akerlof (1970) that have applied a model of adverse selection to data from U.S. health insurance markets. This literature has looked at adverse selection in firms (Einav et al., 2010; Bundorf, Levin, & Mahoney, 2012), and in regulated health insurance exchanges (Handel et al., 2015; Tebaldi et al., 2019; Hackmann et al., 2015; Azevedo & Gottlieb, 2017; Geruso et al., 2019; Saltzman, 2017). These papers have made important advancements in understanding how adverse selection operates in a homogeneous market, where all consumers purchase insurance or remain uninsured. My contribution to this literature is to show that, in the U.S. context, the existence of multiple health insurance markets is very important for interpreting both welfare and for simulating the effects of policy.

This paper contributes to and reinterprets an older literature on the crowd-out effect of public insurance expansions, which began with Cutler and Gruber (1996) (see Gruber and Simon (2008) for an early review, and Frean, Gruber, and Sommers (2017) for a later review). This literature generally assumed that crowd-out makes insurance expansions less effective, because increases in public coverage are counter-acted by decreases in private coverage. My contribution to this literature is to show that crowd-out can actually strengthen private markets if it targets an adversely selected population. A more recent literature estimates the aggregate effect of crowd-out in the ACA, and finds that crowd-out was either non-existent or small (Frean et al., 2017). I show that focusing on aggregate crowd-out can mask substantial heterogeneity across different types of consumers. In the case of Medicaid expansion, substantial crowd-out among the newly Medicaid eligible was counteracted by crowd-in among those responding to lower prices on the individual market.

The closest antecedents to my work are that of Clemens (2015) and Peng (2017). These papers report the effect of Medicaid expansion on the price in certain private markets (Clemens (2015) shows this for community rated markets pre-ACA, and Peng (2017) shows this for Pennsylvania and Indiana post ACA). My paper using Medicaid expansion as a natural experiment confirms that these results hold when applied to the ACA as a whole. I then build on these results to show how these price changes on the individual markets affected how consumers re-sorted across markets, including interactions with employer-sponsored insurance. Finally, I integrate these findings into a general empirical model which allows me to estimate a variety of counterfactuals.

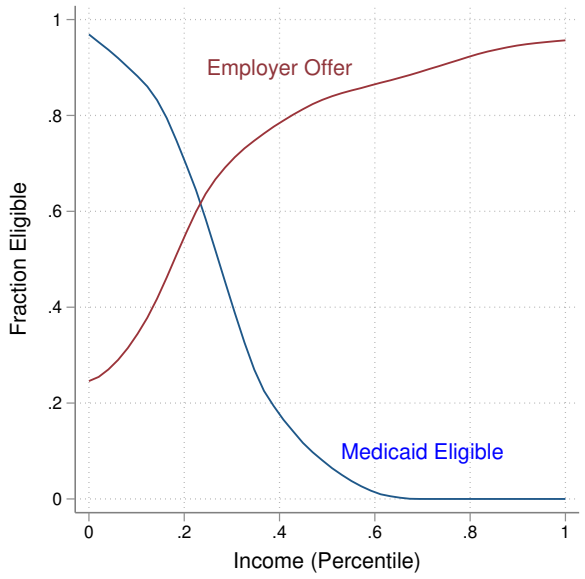
This paper proceeds as follows. In Section 1, I describe U.S. health insurance markets and explain how adverse selection leads to substantial variation in the cost of health insurance across Americans. In Section 2, I present my theoretical model which shows how adverse selection interacts with eligibility when considering how consumers sort across markets. In Section 3, I show reduced-form estimates using cross-state regression. In Section 4, I set-up and in section 5 I estimate the full discrete choice demand model. In Section 6, I conclude.

1 Insurance Markets in the United States

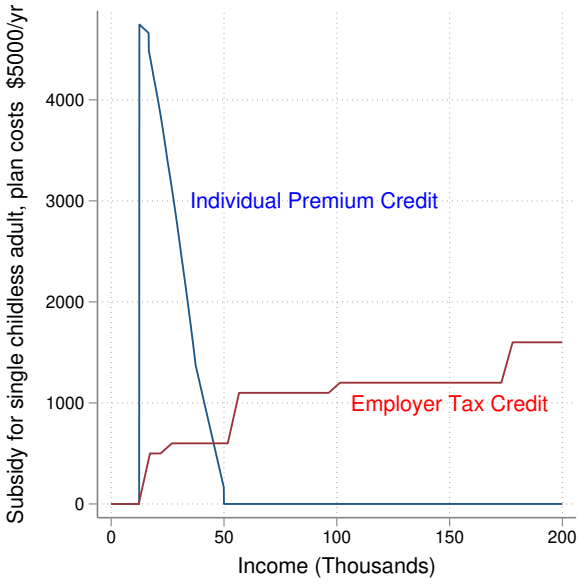
In this section, I show how US health insurance regulations lead to major differences in the price of health insurance across markets. Broadly speaking, Americans under age 65 get health insurance from the government (mostly Medicaid), through their place of work, or directly through a health insurance exchange.[3] These markets vary in three key dimensions: who has access, the sticker price of coverage, and the degree they are subsidized. Because of how regulations are structured, I will show that both rich and poor Americans generally have access to cheaper insurance with larger subsidies relative to the middle class.

Figure 1: Eligibility Rates and Subsidies by Market

(a) Fraction Eligible by Market



(b) Net Subsidies by Market



Notes: Employer offer rates are taken from state-level estimates of employer offer rates by firms of different wage quartiles in MEPS-IC for 2017, smoothed to cover the whole wage distribution and averaged across states using survey weights from the American Community Survey. Medicaid eligibility is simulated using state eligibility limits by eligibility group, gathered by the Kaiser Family Foundation. Employer tax credit calculated assuming existing marginal tax rate schedule. Incomes are ranked using the 2017 sample of the American Community Survey.

The first way that these markets differ is based on access, which as Figure 1a shows, is tightly linked to socioeconomic status. High wage firms are more likely to offer the kind of stable, full-time work that usually comes with health insurance: In 2018, 87% of employees that work for firms in the highest wage quartile were eligible for health insurance, compared to only 35% for firms in the lowest wage quartile.[4] By contrast,

[3] Other types of insurance not include insurance provided by the military (Tricare for active duty military and VA health care for veterans) and Medicare (which covers all Americans over 65, and certain people with a disability or who are diagnosed with end-stage renal disease). In the empirical portion of this paper, I will exclude from my analysis all products that do not qualify as health insurance (such as health sharing ministries or mini-med plans).

[4] Source: Medical Expenditure Panel Survey, Insurance Component. The fraction of employees offered health insurance is calculated as the percent of employees in establishments that offer health insurance times the percent of employees eligible for

as a means-tested program, Medicaid is available to only poor households. Since 2014, everyone has had access to the individual market.

The second way health insurance plans differ is subsidies, which are shown in Figure 1b. This figure shows the expected subsidy a person would receive if they purchased a plan with an annual premium of $5000. Insurance is subsidized in three main ways: Government provided insurance (like Medicaid) is generally provided in kind, with no (or very small) premiums, so is fully subsidized. On the individual market, households with incomes between 100% and 400% of the federal poverty level are given partial subsidies linked to the price of coverage. These subsidies cover most of the price of coverage for households earning 138% of the federal poverty level, and this subsidy decreases for higher income households. Finally, because employer-sponsored health coverage is not taxed as income, employer-sponsored insurance has a large implied subsidy for those in higher tax brackets.

Lastly, the price of similar coverage can vary across markets. From a regulatory perspective, insurers can sell insurance in the individual, small group, or large group markets, where the small group market is available to firms with under 50 employees,[5] and the large group market is available to firms with over 50 employees.[6] While there are strong regulations governing how premiums can vary *within* each of these markets, no regulation restricts how premiums vary *across* these markets. Since the same insurance companies often operate in all three of these markets simultaneously, nearly identical plans are commonly offered in each market, but with very different premiums.

To demonstrate how prices can differ across markets, I show in Figure 2 that the sticker prices (which is the price before subsidies) for plans sold in the individual market are substantially higher than for identical plans sold in the small group market. To construct this figure, I identify all gold plans sold on the individual market where I can find a corresponding plan in the small group market with nearly identical characteristics. I then report the median difference in price across the individual and small group plans by state, and plot it against the average price of purchasing a gold plan as a small group. This figure shows that the price of an individual market plan is higher in nearly all states, except New Jersey. In states like Iowa, Alabama, and Virginia, an individual plan costs over 50% more than its identical counterpart in the small group market.

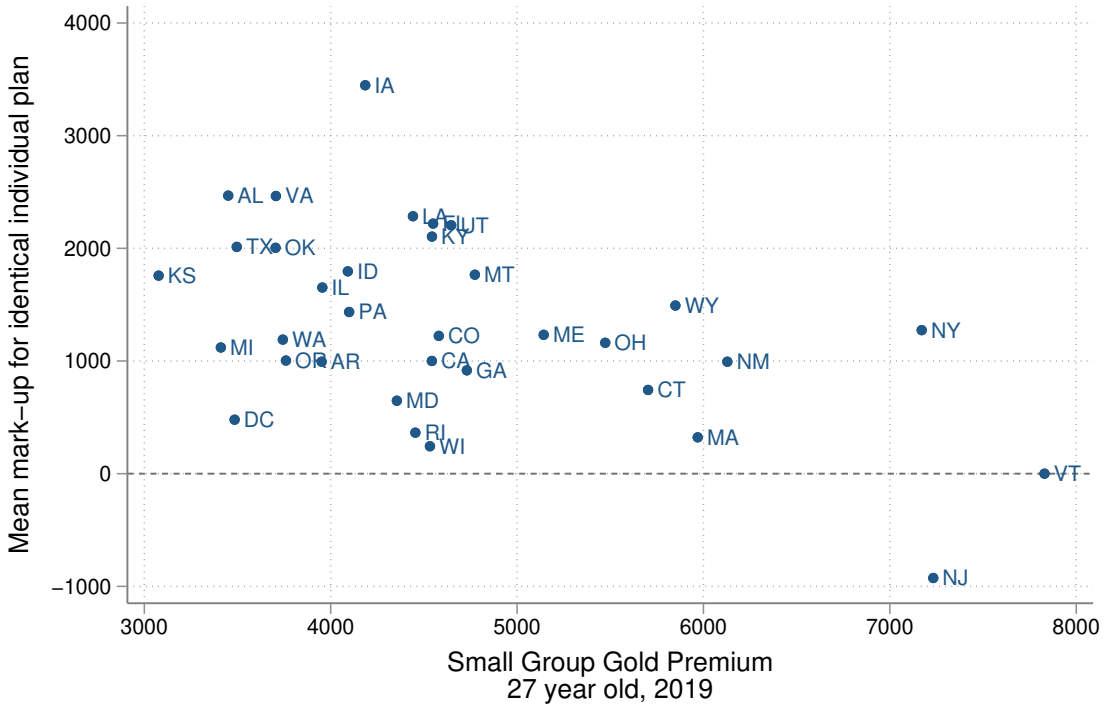
To augment this data, in Appendix B I do a more general analysis of how prices across all health insurance markets. In Appendix B, I show that large groups are also advantageously selected relative to both the small group and the individual markets, but the difference between the large group market and the small group market is much smaller than the difference between the small group market and the individual market. In sum, I find that in 19 states an average large firm would pay over 50% more to purchase their existing

health insurance in establishments that offer health insurance.

[5] The small group market includes firms with up to 100 full-time employees in California, Colorado, New York, and Vermont.

[6] For a full description of how these markets differ from a regulatory standpoint, see Appendix A

Figure 2: Relative Premium of Identical Plans in the Small Group and Individual Markets

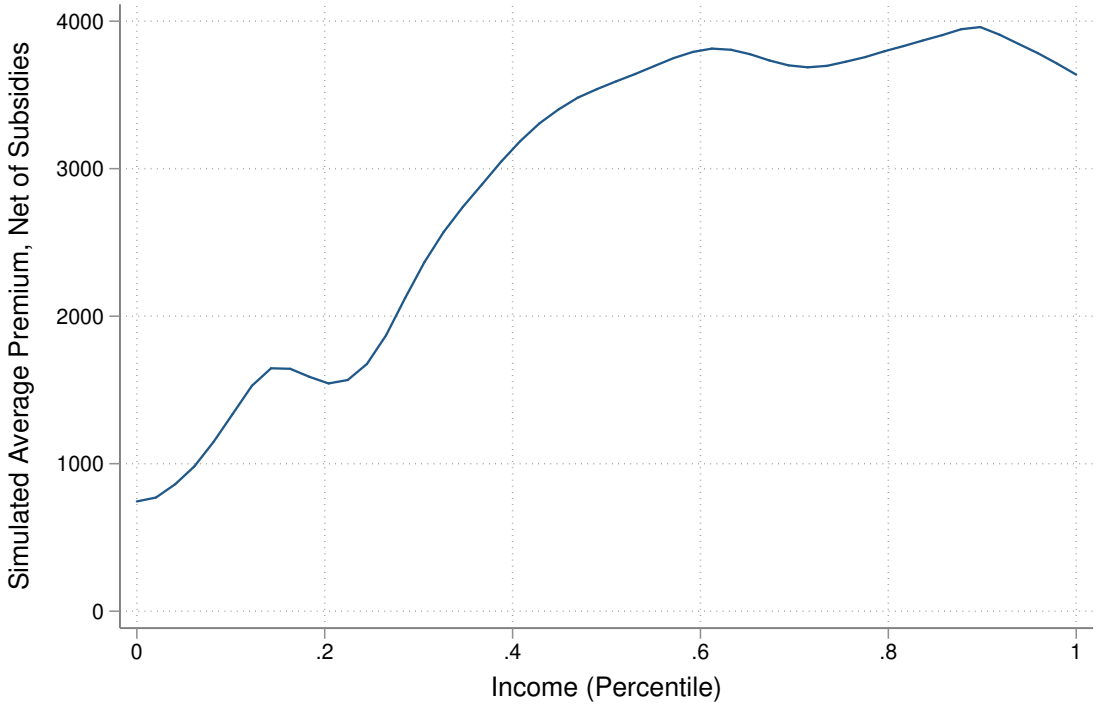


Notes: Figure uses premium data from data compiled by HIX compare (https://www.hixcompare.org). Premiums are measured for a 27-year old with no subsidies. The x-axis reports the average premium of all gold plans available in the ACA compliant market both on and off exchange. The y-axis reports the median difference in price for near-identical individual and small group plans. Two plans are near-identical if they share an insurer, network, plan type (HMP/PPO, etc), metal tier, and geographic rating area. If an insurer offers multiple plans with these identical characteristics on the same market, I take the average premium across all of them. Plans are weighted such that each region has equal weight within a state.

insurance plan if they had to purchase it at the price listed on the individual market. In these states, independent contractors and part-time workers have to pay thousands of dollars more for coverage relative to the price paid by peers in both small and large firms.

Finally, Figure 3 combines the insights from all parts of this section to report how the price of coverage varies across income groups in the United States. In this figure, I report the average amount that it would cost to purchase a standard Gold plan or Medicaid, net of subsidies. The inverse-U shaped graph shows that Americans at the middle of the income distribution generally pay a much higher rate for health insurance relative to either low-income or high-income individuals. The inverse-U shape is caused both by subsidies and market access. Middle income individual have access to the lowest subsidies, and also often only have access to the expensive insurance market, the individual market. This is particularly true for self-employed individuals at just above the median of the income distribution.

Figure 3: Predicted Average Premium for Gold Plan, Net of Subsidies



Notes: This figure reports the estimated average yearly paid by all Americans by income group, net of subsidies, in 2017. To construct this figure, I find the insurance market for each individual in the American Community Survey, and assign the average predicted premium (net of subsidies) for a Gold insurance plan in that market. I set a price of zero for Medicaid enrollees, and I impute the price of an individual market plan for those who are uninsured.

2 Theoretical Model

In this section, I develop a theoretical model of how consumers select across multiple insurance markets. This model is similar in spirit to Einav et al. (2010), except that consumers select an insurance *market* rather than an insurance *product*. Selecting a market is similar to choosing products except that many consumers are not eligible to purchase insurance in all markets. For example, Medicaid is available only to those who pass a means test, and employer-sponsored insurance is available to only those are household members of an employees of firms that offer coverage. As a result, adverse selection will be driven by both demand and eligibility, and policies impacting eligibility can impact adverse selection.

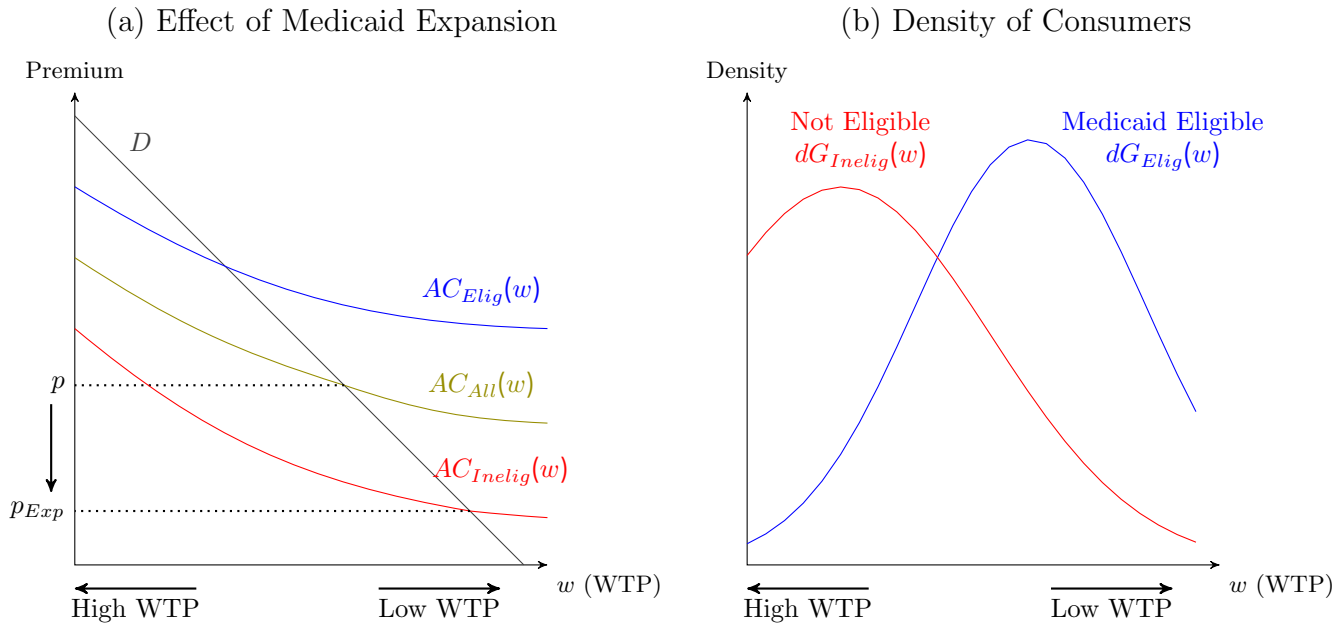
2.1 Graphical Model

In Section 4, I will present a full model including all of the different types of insurance in the United States. For now, I present a simplified version of the model which shows how a public health insurance expansion can affect a private market. To correspond to the natural experiment described in future sections, I will

denote the public market as "Medicaid," and the private market as "individual market."

Consider a population that, initially, is able to purchase insurance on only the individual market. Suppose there exists only a single, standardized health insurance product on the individual market, sold at a premium p. This plan is guaranteed-issue, meaning insurers cannot deny coverage to any consumer, and community-rated, meaning prices do not vary across consumers according to their health status. Suppose each consumer i has a willingness to pay for private insurance w_i, and the cumulative distribution of willingness to pay in the population is $G(w)$. When insured, each consumer generates c_i medical costs, and the average cost to insure all consumers with willingness to pay under p is $AC_{All}(w)$. Finally, assume that the insurance market is competitive and there exists a single equilibrium where $p = AC_{All}(p)$.

Figure 4: Market Equilibrium in the Individual Market

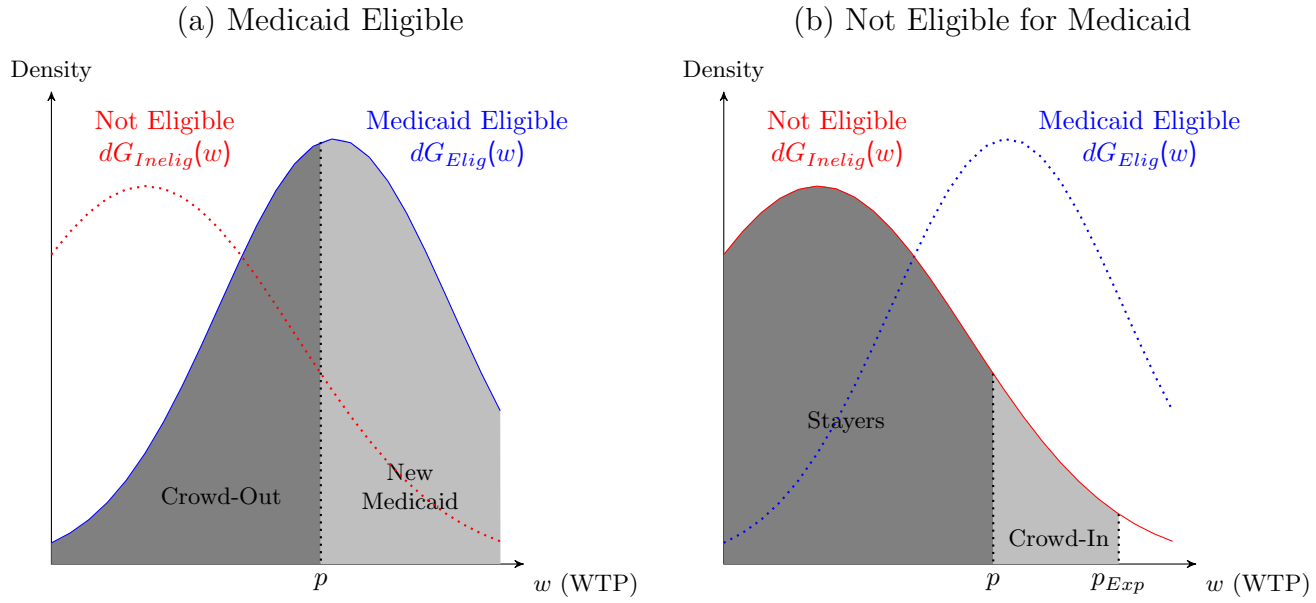


The baseline equilibrium is shown in Figure 4a, where $AC_{All}(w)$ crosses the demand curve. This figure replicates the core graph from Einav et al. (2010), except that I scale the x-axis by willingness to pay. Scaling by williness to pay means that demand will always be a 45° line, and the underlying distribution of people is not uniform. An example of how individuals can be distributed is shown in Figure 4b.

Now suppose the government is considering an expansion of public health insurance, Medicaid. Suppose further that there is a 100% take-up of Medicaid (i.e. all those who are eligible for Medicaid take it up). I can then define new functions G_{Elig} and G_{Inelig} where are the cumulative distributions of willingness to pay for those who are eligible for Medicaid expansion and those who are not. Similarly, $AC_{Elig}(w)$ and $AC_{Inelig}(w)$ are average costs for each of these groups.

As shown in Figure 4a, the effect of Medicaid expansion on the private market depends on the relative position of the two curves $AC_{Inelig}(w)$ and $AC_{All}(w)$. The figure shows an example where $AC_{Inelig}(w)$ is below $AC_{All}(w)$, which means the population that is ineligible for Medicaid is advantageously selected relative to the population that is eligible for Medicaid, conditional on willingness to pay. In this case, after Medicaid expansion, prices in the individual market would drop from p to p_{Exp}.

Figure 5: Impact of Medicaid Expansion on Individual Types



To decompose the effect of Medicaid expansion it is helpful to divide the population into groups. In Figure 5a, I show the effect of the policy on those who gain access to Medicaid. Among this group, those who have a high willingness to pay for insurance are "crowded-out" of private coverage by Medicaid. Those with a low willingness to pay for insurance go from having no insurance to having Medicaid. Among those who remain ineligible for Medicaid, which I show in Figure 5b, there are the "stayers" who are willing to purchase insurance at price p. In addition, if prices drop as a result of Medicaid expansion, there will be additional "crowd-in" of individuals who are ineligible for Medicaid but are willing to purchase insurance at the lower price p_{Exp}. To calculate the welfare effects of Medicaid expansion, governments should consider both the direct effects to new Medicaid enrollees, as well as the indirect effects to individual market participants, both to stayers (who pay lower prices) and to those who crowd-in.

The main takeaway from this section is that, to calculate the effect of government health insurance expansions on private markets, it is critically important to understand how average costs conditional on willingness to pay vary across groups that are eligible and groups that are ineligible for insurance expansions.

Since willingness to pay can vary across groups, it is not sufficient to simply measure average health status across eligible and ineligible groups. What matters is how adversely selected the group of eligibles are who would also purchase insurance in the private market.

2.2 Estimating Demand at the State Level

While the overall health insurance market is much more complicated, the lessons of the graphical model can still be applied to analyzing certain policies, given an appropriate natural experiment. In this section, I map out a framework to analyze the effect of an expansion of Medicaid on the individual market at the state level. To do so requires a slight generalization of the concepts in the graphical model.

This approach requires some experimental or quasi-experimental variation in prices on the individual market. In this paper, I will use a set of federal subsidies which I describe in Section 3 to measure price differences. Since I am not modelling plan selection conditional on market, I collapse the full menu of prices on the individual market into a single price. In my baseline specification, this price is the price of a benchmark silver plan, which is defined in legislation as the price of the second-lowest cost silver plan on an individual insurance exchange. Measuring only a single price per market means I ignore some potentially important within-market price variation, which Geruso et al. (2019) discuss at length. Showing how within-market and cross-market adverse selection can interact is a subject for future research.

Assuming quasi-experimental variation in prices, it is straightforward to plot aggregate demand in the individual market with respect to prices, $D(p|\Omega)$. Here, I define Ω to be the current policy environment in all other health insurance markets. Ω includes, for example, the price and availability of employer-sponsored health insurance, and whether an individual health insurance mandate exists. I include Ω in this equation to emphasize that these curves are valid in only the specific policy context where they are studied, and that changes in that context could drastically affect how prices affect demand.

Now, define Ω_{Exp} to be the policy environment in expansion states and Ω_{Nexp} to be the policy environment in non-expansion states. Combining quasi-experimental variation in prices with Medicaid expansion allows me to identify four objects: $D(p|\Omega_{Exp})$, $D(p|\Omega_{Nexp})$, $AC(p|\Omega_{Exp})$, and $AC(p|\Omega_{Nexp})$. If, in addition, the only difference between expansion and non-expansion states is the availability of Medicaid, then the entire difference between $D(p|\Omega_{Exp})$ and $D(p|\Omega_{Nexp})$ can be attributed to the existence of Medicaid.

This approach enables a direct comparison of prices and enrollment in the individual market between expansion and non-expansion states. Let p_{Exp} be average prices in expansion states and p_{Nexp} be average prices in non-expansion states. Then, by comparing enrollment in expansion and non-expansion states, it is straightforward to decompose differences in demand into price effects and the crowd-out effect of Medicaid.

$$\overbrace{D(p_{Nexp}|\Omega_{Nexp}) - D(p_{Exp}|\Omega_{Exp})}^{\text{Difference in Demand}} = \overbrace{[D(p_{Exp}|\Omega_{Nexp}) - D(p_{Exp}|\Omega_{Exp})]}^{\text{Crowd-out from Medicaid expansion}} + \underbrace{[D(p_{Nexp}|\Omega_{Nexp}) - D(p_{Exp}|\Omega_{Nexp}]}_{\text{Crowd-in from lower prices}} \tag{1}$$

Further, using data on premiums and average costs, I can document whether premium differences across states were driven by adverse selection. Average premiums in a market p are just equal to average costs times a mark-up. In the following sections, I will show that mark-ups are small and do not vary much across markets, meaning the majority of premium differences can be attributed to differences in average costs across markets.

3 Reduced Form Estimates

In this section, I use arbitrary features of the Affordable Care Act to estimate how Medicaid expansion affected demand for insurance in the individual market, and the rate of substitution across insurance markets. The goal is to estimate all of the pieces of Equation (1), to estimate both the crowd-out effect of Medicaid and whether it affected prices in the individual market. Since no dataset has all the pieces I need, this analysis makes use of a combination of several datasets. I will begin by describing the natural experiments that I will use, and then I will present the estimates.

3.1 The Affordable Care Act

A central goal of the Affordable Care Act (ACA) was to increase the number of people covered by health insurance. For low-income individuals (those earning under 138% of the federal poverty level) the ACA did so by expanding Medicaid eligibility. The ACA also created individual health insurance exchanges, which were meant for those who earn above 138% of the federal poverty level.

While Medicaid expansion was designed to be enacted nationally, the planned Medicaid expansion was not actually enacted in all states. In the case *National Federation of Independent Business v. Sebelius*, the Supreme Court decided in June of 2012 that the federal government could not compel states to expand their Medicaid programs. As a result, 23 states decided to not expand Medicaid, of which four subsequently expanded in 2015 and 2016. In non-expansion states, no provision in the Affordable Care Act provides subsidized health insurance for those earning less than 100% of the federal poverty level. More details on Medicaid expansion are provided in Appendix A.

When it was passed, the Affordable Care Act also included two temporary programs to help stabilize prices in the new individual health insurance exchanges: reinsurance and risk corridors.[7] The reinsurance program provided additional funding to reimburse plans for individuals with outlier medical costs (above $45,000 per year). The risk corridor program acted as a transfer policy from more profitable to less profitable plans. Plans with claims-to-premiums ratios of under 77.6% pay into a fund which reimburses plans with claims-to-premiums ratios of over 82.4%.

As a result, the passage of the Affordable Care Act actually acted as a two-stage experiment. Medicaid expansion and the individual health insurance exchanges were enacted in 2014. In 2014-2016, the the risk profile of populations in health insurance exchanges differed between expansion and non-expansion states, but prices were relatively similar. Once the risk adjustment and risk corridor subsidies were removed, then a substantial price difference became apparent between expansion and non-expansion states. In Figure 6, I show the role of reinsurance and risk corridor subsidies by showing how the price of coverage varied between the individual market and small group market in each state. Between 2014-2016, plans sold on the individual market were actually cheaper than equivalent small group plans even though the individual market enrolled a sicker population, because the reinsurance and risk corridor programs allowed prices to remain low. Prices increased drastically once these programs were fully phased out in 2017, and the difference in prices between expansion and non-expansion states only appeared in 2017.

The full effect of Medicaid expansion is not felt until 2017, after all reinsurance and risk corridor subsidies were removed. To tranlate this into the variables from the model, my goal is to compare insurance demand in expansion states in 2017 ($D(p^{2017}_{Exp}|\Omega_{Exp})$) to insurance demand in non-expansion states in the same year ($D(p^{2017}_{NExp}|\Omega_{Nexp})$) and determine which portion of this difference in demand is due to crowd-out and subsequent crowd-in.

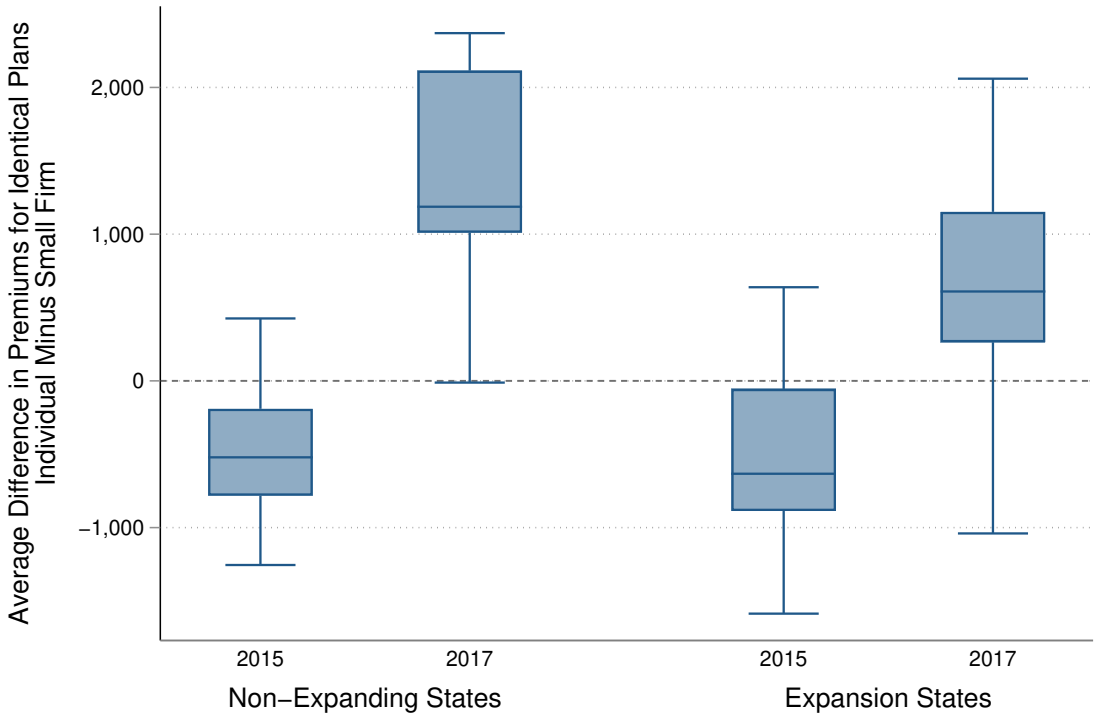
The advantage of this natural experiment is that we have a natural counterfactual (since p^{2015}_{NExp} is similar to p^{2017}_{Exp}, we can use that to predict the counterfactual $D(p^{2017}_{Exp}|\Omega_{Nexp})$. As a result, I basically use the early years of the ACA to predict the effect of crowd-out without any price effects $D(p^{2014}_{Nexp}|\Omega_{Nexp})-D(p^{2014}_{Exp}|\Omega_{Exp}) \approx D(p^{2014}_{Nexp}|\Omega_{Nexp}) - D(p^{2014}_{Nexp}|\Omega_{Nexp})$. I then use the price shocks in 2017 to trace out demand curves.

3.2 Measuring Crowd-Out

To measure the degree to which Medicaid crowded-out private coverage, I estimate the event study in Equation (2):

[7] The ACA also contained a provision for risk adjustment, which unlike reinsurance and risk corridors was designed to be permanent. The risk adjustment program is revenue neutral and only transfers funds across plans within a given state.

Figure 6: Adverse Selection in the Individual Market, by State



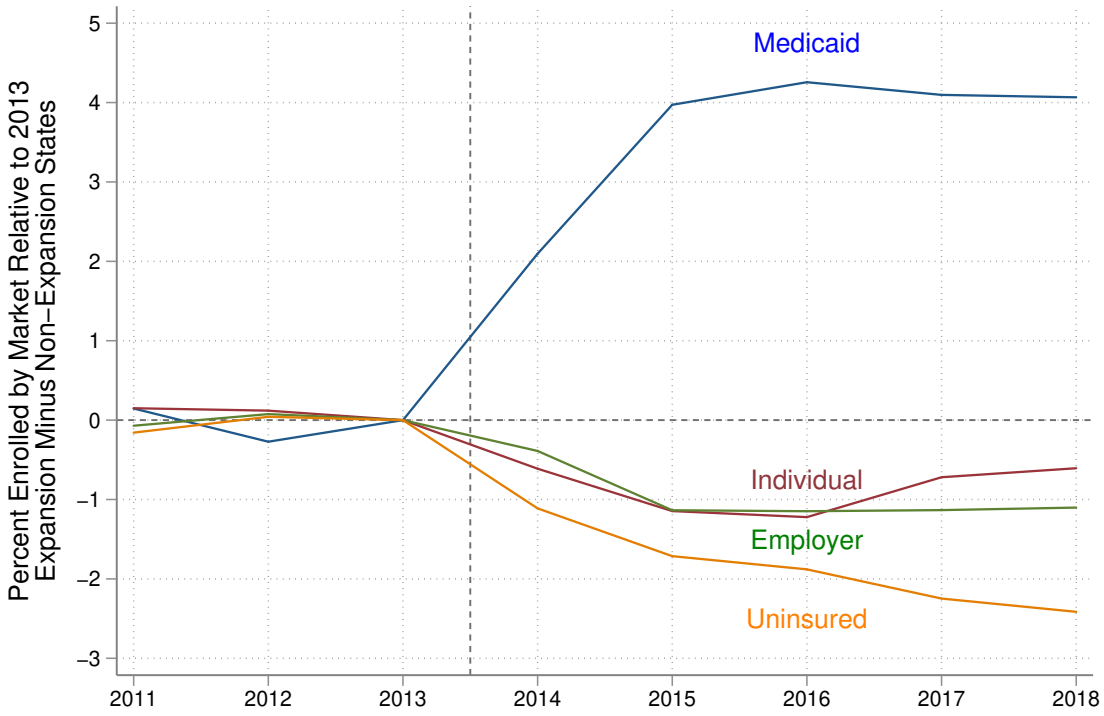
Notes: This figure measures the mean mark-up for an individual market plan relative to an identical small group plan, which is a way of measuring the degree of adverse selection in the individual market. Using data complied by hixcompare (www.hixcompare.org), I identify all Gold plans that are sold in both the individual and small group markets simultaneously. I then find the average difference in prices for identical plans within a state (premiums are measured for a 40-year-old person). The Box Plot reports the 25th, Median, and 75th percentile states, and whiskers report adjacent values (I exclude outside values). I exclude states which expanded Medicaid in 2015 or 2016 (Pennsylvania, Indiana, Alaska, Montana, and Louisiana). See Appendix C for full details on my approach for estimating the effect of Medicaid on prices in the individual market using plan-specific controls.

$$Y_{it} = \beta_t E_{s(i,t)} + \alpha_{s(i,t)} + \gamma_t + \epsilon_{it} \tag{2}$$

In this regression, Y_{it} is an indicator for whether person i participates in a given insurance market at time t. The key dependent variables are participation in Medicaid, the employer market, the individual market, and non-insurance. E_s is an indicator equal to one if the state is a Medicaid expansion state, α_s are state fixed effects, and γ_t are time fixed effects.

The identifying assumption in this regression is that expansion and non-expansion states are subject to common trends. In figure 7, it appears that all markets are trending similarly before 2013, which supports the assumption of common trends. To attribute changes in enrollment across markets to Medicaid expansion, this paper also assumes that all other elements of the Affordable Care Act affected expansion and non-

Figure 7: Enrollment by Market, Expansion Minus Non-Expansion States



Notes: This figure plots difference in enrollment between expansion and non-expansion states by market, normalized to 2013, using data from the American Community Survey. The sample includes all individuals aged under 65, and is weighted using ACS population weights. I exclude states which expanded Medicaid in 2015 (Pennsylvania, Indiana, and Alaska).

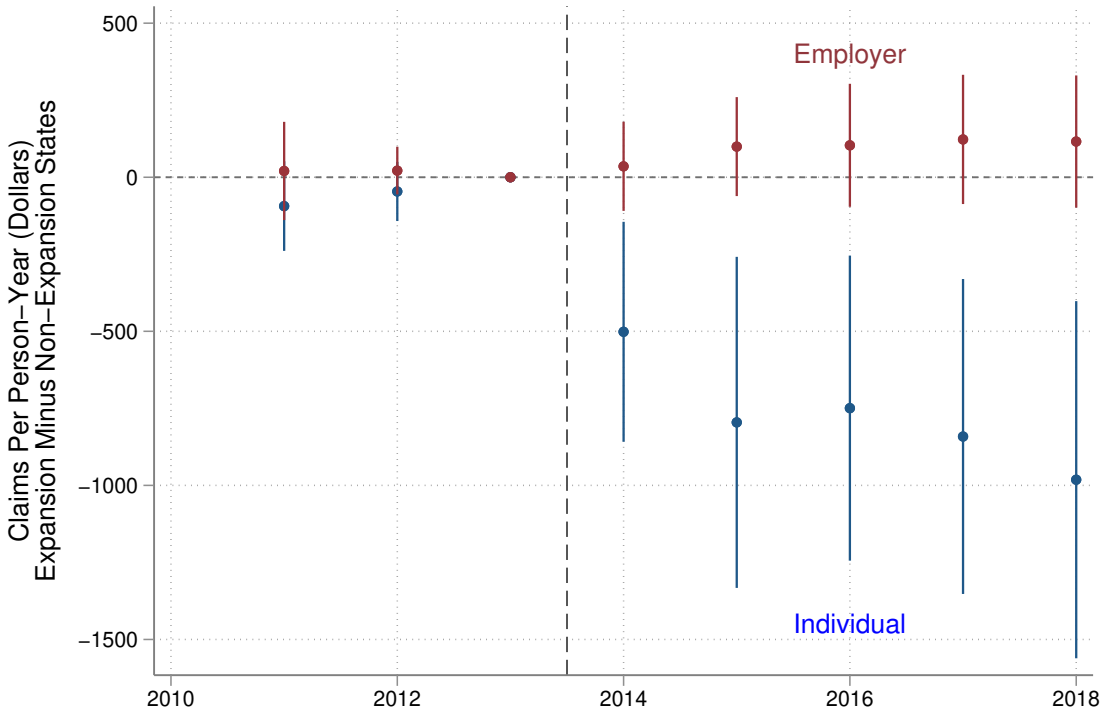
expansion states equally.[8] Finally, note that in the language of heterogeneous treatment effects, this event-study analysis estimates the effect of Medicaid for Medicaid expansion states (average treatment effect on the treated). Since states which expanded Medicaid expansion had more generous Medicaid programs and a lower uninsured population prior to the enactment of the Affordable Care Act, the measured effect of Medicaid expansion is almost certainly a lower bound for the predicted effect of Medicaid expansion in non-expansion states.

To estimate the effect of Medicaid expansion on enrollment, I estimate Equation (2) using the American Community Survey Public Use Microdata Sample from 2011-2018. I restrict my sample to adults aged 21-64. To get as clean a comparison as possible, I drop people living in Pennsylvania, Indiana and Alaska which expanded Medicaid in 2015. I use sampling weights and cluster standard errors at the state level.

I report event-time coefficients β_t from two instances of Equation (2) in Figure 7, where the dependent

[8] While there were many policy changes enacted as a part of the Affordable Care Act, most of these were applied equally in all states. However, two policies which differed between expansion and non-expansion states were the fraction of plans covered by "grandmothering" and the fraction of states which ran their own health insurance exchange rather than use the federal health insurance exchange. In a separate analysis (available on request) I show that neither of these factors played a big role in the price of on-exchange plans or enrollment in the individual market.

Figure 8: Impact of Medicaid Expansion on Claims Per Person



Notes: This figure reports event-time coefficients from the regression Equation (2), estimated at the state level, and using total medical claims per enrollee-year as the dependent variable. I exclude states which expanded Medicaid in 2015 (Pennsylvania, Indiana, and Alaska). Data taken from Medical Loss Ratio data, published by The Center for Consumer Information and Insurance Oversight of the Centers for Medicare and Medicaid Services.

variables are Medicaid enrollment and non-insurance. Prior to 2014, both the fraction of the population receiving Medicaid and the fraction that is uninsured trended similarly in expansion and non-expansion states. Starting in 2014, Mediciad enrollment rose for 2-3 years in expansion states relative to non-expansion states until it plateaued at 5-6% higher between 2016 and 2018. The fraction of the population not covered by any insurance fell at the same time, but only by an average of 2-3%. This figure implies that approximately half of new Medicaid enrollees in expansion states would still have been insured by some other source had their state not expanded Medicaid.

The next question is, were these new Medicaid enrollees adverseley selected? Using the same event-study approach, I test whether expansion states have a sicker or healthier population in their individual and employer-sponsored markets. I find that selection into Medicaid plays an important role in setting prices on the individual health insurance exchanges, but does not substantially affect prices in the employer-sponsored market.

To test for adverse selection, I track how insurance claims vary across market segments and over time. I

measure total claims per enrolled member, which is the amount paid by insurance companies for all medical services (including from hospitals and physicians) and prescription drugs. The idea is that, if a market enrolls an adversely selected population, total claims in that market will increase. Adverse selection is not the only reason claims could vary across markets. For example, markets with more generous insurance or that insure patients in areas with more expensive hospitals, could have higher claims per member. However, in Appendices B and C, I show that the results from this section hold after controlling for both geography and plan characteristics.

I measure claims using the Medical Loss Ratio data, which includes aggregate premiums, enrollment, and health spending by state, insurance company, and market segment. I aggregate total claims per insured person by market to the state level, and report results in Figure 8. This figure shows that spending in all markets was trending similarly prior to 2014 in (with only one slightly statistically significant coefficient). Starting in 2014, there is a dramatic trend break in the individual market. Spending per member drops considerably in expansion relative to non-expansion states. In particular, spending in expansion states is approximately $800-$1000 lower per enrollee in 2015-2017 relative to non-expansion states.

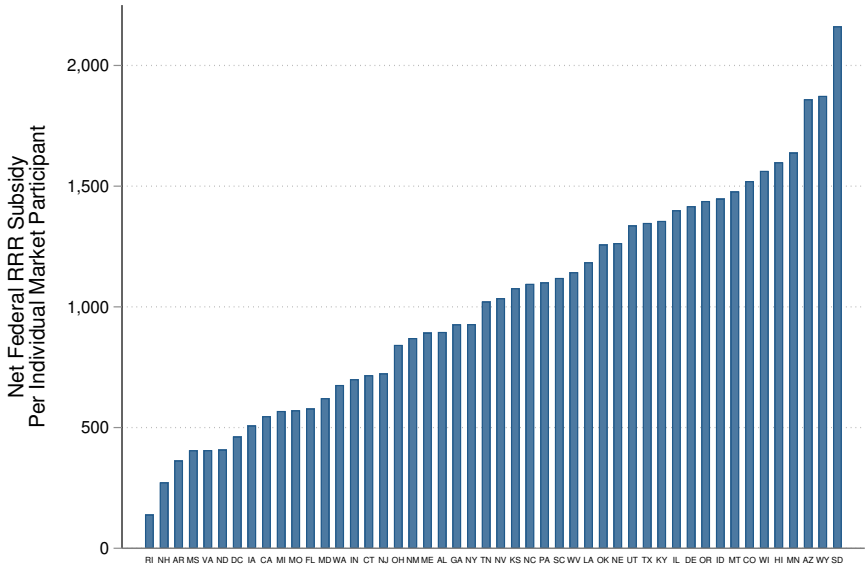
The reason that total claims per person in the individual market diverges so drastically between expansion and non-expansion states comes down to the group of "marginal" consumers who take up Medicaid in expansion states but who would purchase insurance in the individual market in non-expansion states. The adverse selection described here also lowers premiums per person in the individual market in expansion states relative to non-expansion states. In Appendix C, I show that after controlling for plan characteristics, plans sold in expansion states are on average 20 log points cheaper than in non-expansion states by 2018, which is a similar order of magnitude as the results I estimate using the ACS.

3.3 Identifying Demand for Individual Insurance Exchanges

Since the enactment of the Affordable Care Act in 2014, average prices on individual health insurance exchanges have varied widely both over time and across different counties in the U.S. Much of these price changes are due to arbitrary features of U.S. law, and have nothing to do with the supply of health insurance. This arbitrariness allows us to use certain price variation on these exchanges as a natural experiment.

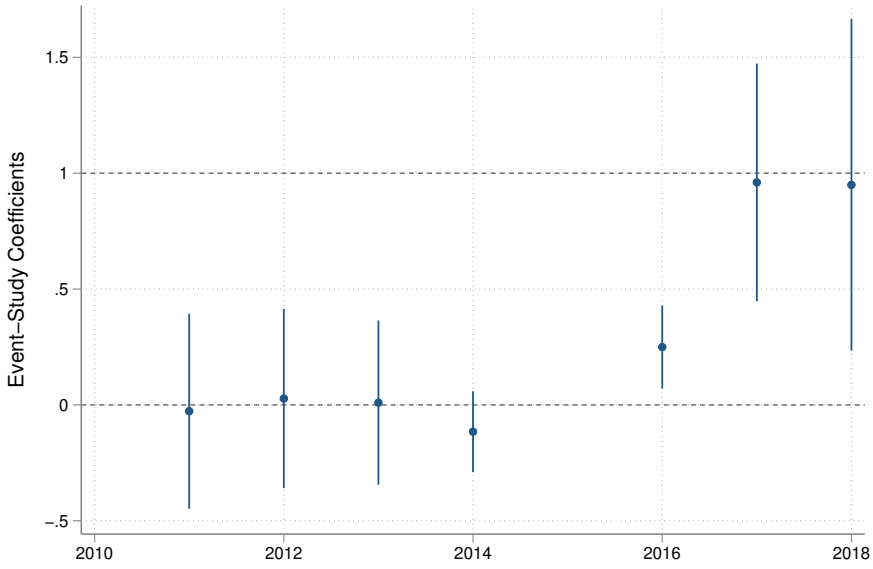
As it happens, certain states benefited much more than other states from these subsidies. In figure 9, I show that in 2015, the average combined subsidy per individual market enrollee varied between $140 to $2200 per individual market participant. Both the reinsurance and risk corridor programs were temporary, operating only in 2014-2016. After these transfers were removed, prices increased, particularly in states which had large subsidies before. I use the removal of these subsidies as an instrument for on-exchange

Figure 9: Reinsurance and Risk Corridor Subsidies by State



Notes: This figure reports net Reinsurance and Risk Corridor payments per enrollee-year by state in 2015. I exclude Massachusetts and Vermont (which have a combined individual/small group market) and Alaska (which is a huge outlier). Data taken from Medical Loss Ratio data, published by The Center for Consumer Information and Insurance Oversight of the Centers for Medicare and Medicaid Services.

Figure 10: Individual Market Premium vs. 2015 Reinsurance and Risk Corridor Subsidies



Notes: This figure reports event-time coefficients from the regression equation (3), estimated at the state level, and using total medical premiums per enrollee-year in the individual market as the dependent variable. I exclude Massachusetts and Vermont (which have a combined individual/small group market) and Alaska (which is a huge outlier). Data taken from Medical Loss Ratio data, published by The Center for Consumer Information and Insurance Oversight of the Centers for Medicare and Medicaid Services.

prices in these states.

To measure the effect of removing risk corridor and reinsurance subsidies, I use a parametric event study:

$$Y_{it} = \beta_t R_{s(i,t)} + \alpha_{s(i,t)} + \gamma_t + \epsilon_{it} \tag{3}$$

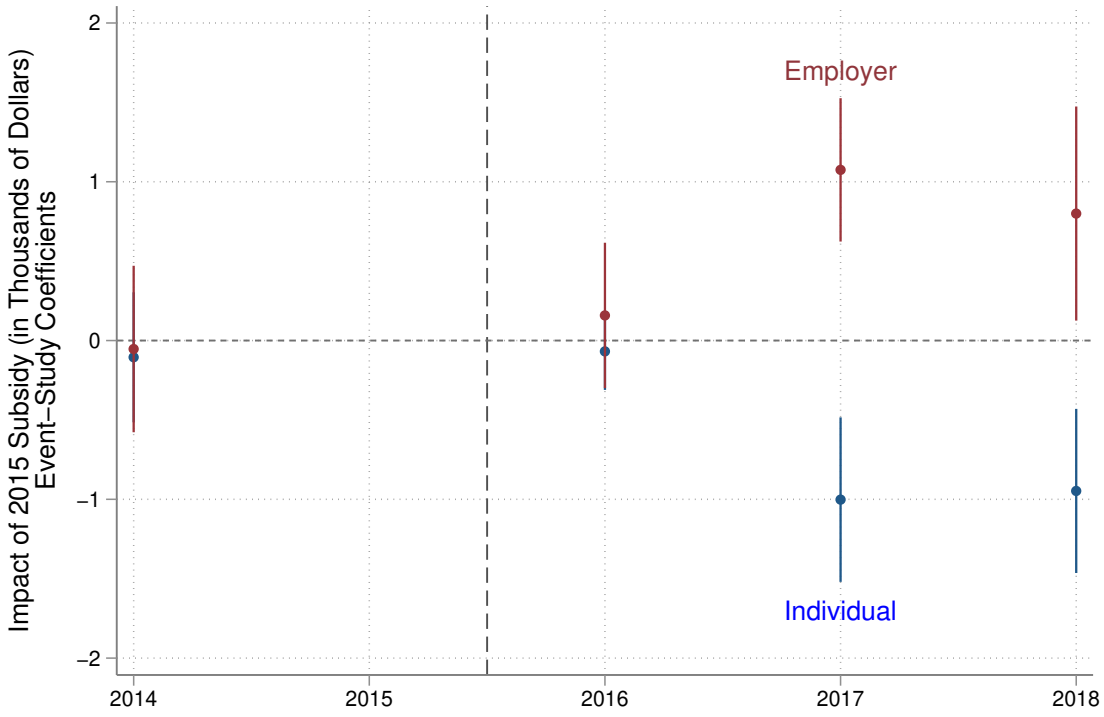
In this equation, R_s is the average combined reinsurance and risk corridor subsidy for state s in 2015, α_s are state fixed effects, and γ_t are year fixed effects. Dependent variables will include a measure of average premiums by state (for the first-stage regression) and indicators for participation in each insurance market by state (for the second-stage regression). The identification condition for this regression is that counterfactual changes in enrollment across markets between 2015 and 2017 (what would have occurred without the removal of subsidies) are uncorrelated to average subsidies in 2015. I will provide evidence this is a reasonable assumption by showing the absence of a major pre-trend. Moreover, other known factors that influence prices on ACA exchanges are not strongly correlated with promised reinsurance and risk corridor subsidies. For example, both expansion and non-expansion states are well represented in the states which receive a large subsidy and those what receive a small subsidy.

I first show the first-stage results from equation (3) on premiums in the individual market. The dependent variable in this regression is average premiums per enrolled member in the individual market by state, which I get from the Medical Loss Ratio data. The dependent variable is total promised reinsurance and risk corridor subsidies on the individual market in 2015, which I also calculate using the Medical Loss Ratio data. This analysis excludes Massachusetts and Vermont, that each have a joint individual/small group market and thus do not have separate data for reinsurance/risk corridor payments for their individual market. I also exclude Alaska which is a huge outlier, with both higher premiums and dramatically higher reinsurance/risk corridor subsidies.[9]

Results are shown in Figure 9. The fact that the coefficient on $\widehat{\beta_{2017}}$ is approximately 1 means that, on average, prices rose one-for-one when subsidies were removed in 2017. What is potentially more surprising is that there is not another trend break in 2014, since the reinsurance and risk corridor subsidies were enacted in 2014. I take this to indicate that the differences across states in adverse selection onto their individual markets were not a factor before the enactment of the Affordable Care Act. Prior to 2014, insurance companies in most states were allowed to deny individuals with pre-existing conditions from getting coverage, and this means that individual market enrollees were relatively healthy in all states. Between 2014 and 2016, states differed widely in terms of the average health of their enrollee populations, but premiums were kept fairly similar across states because states with a sicker individual markets benefited more from reinsurance and

[9] Results are similar with Alaska included, but point estimates are much noisier.

Figure 11: Event-Study Coefficients of Effect of Reinsurance and Risk Corridor Subsidies on Enrollment



risk corridor subsidies. It was only in 2017 that the full extent of differences in adverse selection across states was fully incorporated into premiums on the individual market.

3.4 Impact of Subsidy Removal on Enrollment

To measure the effects of reinsurance and risk corridor subsidies, I estimate equation (3) in the American Community Survey Public-Use Microdata Sample. My sample is adults aged 20-64. Moreover, since individuals earning under 400% of the federal poverty level have access to subsidies (and therefore are not fully subject to price increases on the individual exchange), I restrict the sample to individuals earning over 400% of the federal poverty level. I use sampling weights and cluster standard errors by state.

Results from this parametric event study are listed in Figure 11. In this figure, all coefficients are scaled to measure the fraction of the population that responds to a $1000 change in the overall subsidy. After subsidies are removed, markets with higher subsidies have a larger drop in enrollment on the individual market in 2017-2018, equal to approximatly .8-.9% of the population per $1000 in subsidy removed. In 2017, it appears that enrollment in the employer-sponsored insurance market rose one-for-one with a drop

in enrollment on the individual exchanges, suggesting that price increases on the individual market have led to substitution towards employer-sponsored coverage.

4 Econometric Demand Model

Consider now a more general demand model, which allows for multiple types of public and private health insurance markets. Let each consumer choose to participate in a single insurance market $j \in \mathscr{J}$.[10] The set of insurance markets $\mathscr{J}$ includes the individual market (I), the market for employer-sponsored coverage (F), and Medicaid (MC), and consumers may also choose to stay uninsured (U). In this model, three elements are important for consumer choice: demand, eligibility, and prices.

First, demand among consumers can vary across different markets. Demand across markets can vary because markets differ in the quality of insurance plans available for purchase. In particular, Medicaid plans tend to have inferior hospital networks but lower cost-sharing relative to similar plans available in private markets. Second, even if consumers have access to the same insurance plans in multiple markets, demand may vary across these settings because the hassle cost of purchasing a plan may vary between Medicaid, individual coverage, and employer-sponsored coverage. I model these preferences as differences in the willingness to pay for insurance across the different markets for health insurance. Formally, let individuals indexed by i have a vector of characteristics ζ_i, which includes all individual-level characteristics which may affect demand for insurance. Then, define willingness to pay for each market j as $W^j(\zeta_i)$. In this case, willingness to pay is measured as the price to purchase a "typical" plan in each market, which I will define below.

Markets also differ based on eligibility requirements. Let ψ_i be a sub-vector of ζ_i, where ψ_i are the elements of ζ_i which can be used in order to determine eligibility in different insurance markets. Examples of variables in ψ_i are income (to determine eligibility for Medicaid) and an employer identifier (to determine eligibility for employer-sponsored health insurance). Then, each person has a person-specific choice set $J(\psi_i) \in \mathscr{J}$.

Finally, markets can differ in the price to purchase health insurance, due to adverse selection. For example, as I showed in Section 1, it is common for identical plans to be sold in the individual and small firm markets, but for the individual market plans to be sold at a higher price. To account for these price differences, let P^j be the price for purchasing a standardized plan in each market.[11] The standard plan I use is a benchmark plan, which is defined in the Affordable Care Act as the second-lowest cost silver plan

[10] I take notation from the discrete choice literature. In a typical discrete choice model, consumers select "products" indexed by j in "markets" indexed by t. In the current paper, "insurance type" will take the place of the product and is indexed by j.

[11] Characterizing price differences in this way is a simplification, because it collapses a menu of prices in a each market to a single price. See Geruso et al. (2019) for a discussion of how having multiple products in each market can affect equilibrium outcomes.

on the health insurance exchanges. I measure prices on the employer-sponsored market using plans with similar characteristics to the benchmark on-exchange plans,[12] and Medicaid plans have a cost of zero. The price an individual pays P_i^j may differ from the overall sticker price in a region because of age adjustment and government subsidies.

$$j_i^* = argmax_{k \in J(\Psi_i)}\{W^k(\zeta_i) - P_i^k\} \tag{4}$$

Consumers will select the plan which maximizes utility among those that are in their choice set, as in equation (4). Consumers select the outside option of non-insurance if no option in their choice-set $J(\psi_i)$ has a utility above zero.

4.1 Assumptions

I make the following assumptions about the data generating process:

Random utility model with product-by-region shock

I assume that utility W_i^j follows the following process:

$$W_t^j = \delta_r^j(d_i) + \xi_{tr}^j + \sigma\epsilon_{itr}^j \tag{5}$$

In this equation, $\delta_r^j(d_i)$ is a flexible function of demographics d_i which is continuous in age and income. I assume two different demand shocks: ξ_{tr}^j is a product-market specific demand shock, and ϵ_{itr}^j is an iid error term distributed as a type-1 extreme value.

This characterization of demand is similar to Tebaldi (2020), but with an additional restriction for how demand varies over time in the same region. I assume that, within a region r, the distribution of demand is constant over time except for a single product-by-region demand shock ξ_{tr}^j that is shared by all consumers in that region. Assuming a single unobservable at the product-by-region level is standard in many models of differentiated product demand (Berry & Haile, 2020). What is unique in my case is that I assume a single market-level unobservable in a model where prices vary within that market. This assumption will imply that the elasticity of demand with respect to prices will be estimated using variation within a market both in the cross section and over time.

Pre-determined choice sets

[12] Typically, employees pay only a small part of the premiums for employer-sponsored insurance, with the employer covering the rest of the premium out of general revenues. However, it is likely that at least a portion of the portion paid by employers is passed-through to employees in the form of lower wages. Since it is unclear what portion of premiums employees are responsible for, in robustness checks (available on request) I conduct robustness where employees are responsible for between 0% and 100% of their total health insurance premiums.

I assume that individuals take their existing choice sets J_i as given. This assumption means that consumers do not respond to changes in either the price or quality of health insurance products by making choices that alter their eligibility to purchase insurance any particular market. In particular, this assumption means that when prices change on the individual insurance markets, consumers do not change jobs to gain access to employer-sponsored coverage, nor do they intentionally lower their incomes to gain access to Medicaid. Assuming static choice sets substantially reduces the complexity of the model, because it does not require a full model of labor market participation and employment choices.[13]

I feel comfortable making this assumption in this setting because a growing literature shows that the Affordable Care Act had either small or no impact on the labor market (Duggan, Goda, & Jackson, 2019; Frean et al., 2017; Kucko, Rinz, & Solow, 2018; Heim, Hunter, Isen, Lurie, & Ramnath, 2019). I confirm the findings of this literature in a supplementary analysis of both Medicaid expansion and price changes on individual markets using data form the American Community Survey and the Medical Expenditure Panel Survey - Insurer's Component (available on request). Comparing Medicaid expansion and non-expansion states, I find no statistically significant changes in aggregate labor market participation or wages after the enactment of the Affordable Care Act. I also find no effect of Medicaid expansion on the number of employees who are eligible for insurance and the number of firms that offer health insurance. Similarly, when prices rose on individual markets in 2017 in response to the removal of reinsurance and risk corridor subsidies, I find no evidence that labor market participation or firm offer rates changed. In response to both natural experiments, the take-up rate of insurance in firms conditional on insurance being offered was the only outcome affected in employer-sponsored markets: Take-up of offered insurance rose when prices on the individual market fell, and it fell in response to Medicaid expansion.

Cost model

Once covered by health insurance, individuals accrue health insurance costs for doctor's visits, hospital admissions, and prescription drugs. Variation in these costs across individuals can be explained, broadly, by factors related to the individual (which drive adverse selection), and factors related to the chosen insurance plan. In measuring and predicting adverse selection across markets, my goal is to abstract from the differences in costs attributable to insurance characteristics and moral hazard, to focus on how costs vary across markets holding constant plan characteristics. To do so, I model costs as in equation (6).

$$c_{ip} = h(d_i) + \epsilon_p + u_{ip} \tag{6}$$

[13] To see models which incorporate both health insurance and labor market choices, see for example Aizawa and Fang (2013); Aizawa and Fu (2020)

In this equation, anticipated insurance costs c_{ip} are total predicted health spending for consumer i purchasing plan p. This equation disaggregates costs into a portion explained by individual characteristics ($h(d_i)$), a portion explained by plan-specific component (ϵ_p), and a portion potentially explained by both (u_{ip}). The portion explained by individual characteristics varies flexibly in the same demographic variables I use in my demand model. Insurance characteristics include the price negotiated between insurers and providers: For example, Medicaid plans tend to contract with cheaper hospitals and clinics, and they pay lower prices at these clinics, relative to private plans. Also, insurance plans both within and across markets can vary in their cost-sharing (like deductibles and coinsurance rates), and consumers respond to more generous plan characteristics by increasing their health insurance utilization.

In estimating this equation, I make the assumption that $E[u_{ip}|d_i, p] = 0$. What this assumes is that, when predicting medical spending, more generous plans (those which lead to higher medical costs) are more generous for consumers of all demographic groups equally.

4.2 Identification

Differences in prices on the individual ACA market

The first source of identification in this model comes from arbitrariness in the price of health insurance on the individual health insurance exchanges. Equation (7) shows how the price of the benchmark health insurance plan varies across individuals on the ACA exchange:

$$P_{ij}^I = min\{P^{MAX}(inc_{it}), \alpha(age_{it})P_r^I\} \tag{7}$$

In each region, insurers set prices for each plan in their menu, and the price of the benchmark plan is simply the second-lowest cost silver plan in that market. To come up with an individual price, the initial price P_r^I is multiplied by an age-adjustment factor $\alpha()$ which is a schedule set in law. The default price schedule in the Affordable Care Act varies within a ratio of 3:1 (a 64-year-old pays 3 times what a 20-year-old would pay) and some states have state-specific price curves. To arrive at the net price, the government calculates an affordable P^{MAX} which a household of income inc_{it} is assumed to be able to afford. If the price of the benchmark plan rises above P^{MAX}, then subsidies are provided to offset price increases one-for-one.

These arbitrary features lead to two key sources of identification in the model. First, the functions P^{MAX} and inc_{it} are non-linear functions with both kinks and notches. Assuming demand varies smoothly across age groups and income levels, one can use nonlinearities in these formulas to identify demand parameters. Finkelstein, Hendren, and Shepard (2019) make use of the notches in this price schedule to identify demand

in the individual market among low-income adults, and Tebaldi (2020); Tebaldi et al. (2019) make use of the arbitrariness of the full price-schedule.

Arbitrariness in these price schedules also allows me to identify the effect of price changes over time. In all states, there exists a group of individuals (who tend to be older, have low incomes, or both) who are completely insulated against price increases on the individual health insurance exchanges: regardless of changes in the overall price level P_r^I, they will still pay exactly $P^{MAX}(inc_{it})$. According to Equation (5), changes in this group's overall demand for insurance will come from changes in ξ_{tr}^I, which enables me to identify this parameter. Then, after controlling for ξ_{tr}^I, I can use changes in the price of insurance for those who do face changes in prices over time to identify the price elasticity of demand on the individual exchange.

Variation in choice sets

The second source of identification is variation in individual-level choice sets. In the cross-section, identification comes from differences in choice sets across otherwise similar individuals in the same region and year. Across two people of similar ages and incomes, one person may work for a firm which offers health insurance, and the other may not. I can use differences in choices across these individuals to estimate the demand for employer-sponsored coverage. Moreover, legislation sets specific income cutoffs for Medicaid eligibility. Comparing how demand for various types of insurance changes across these discrete cutoffs is a further source of identifying variation. Choice sets can also vary over time. In particular, since many states expanded Medicaid after 2015 (see Appendix A.1 for a list), I take advantage of these "later expanders" to measure demand changes after a chance in individual-level choice-sets.

5 Estimation

Consumer demand for market j is listed in equation (8)

$$u_{irt}^j = \alpha^j p_{irt}^j + \delta_r^j(d_i) + \xi_{tr}^j + e_{itr}^j \tag{8}$$

In this specification, demand by demographic group $\delta^j(d_i)$ is estimated using a flexible function of 10 dummy variables for deciles of income, and 10 dummy variables for the deciles of age. I include fixed effects at the state-market-year level, ξ_{tr}^j, and e_{itr}^j is a type-1 extreme value. I normalize coefficients against the outside option of no insurance. The price of coverage in each market is equal to the cost of purchasing a benchmark silver plan.[14] On the individual market, prices also vary because of differences in the rate of subsidies across markets.

[14] The benchmark plan on health insurance exchanges is defined as the second lowest cost Silver plan.

To finish constructing the likelihood function, I then construct individual-level choice sets J_i for each person in the sample. I calculate Medicaid eligibility for each individual according to their income, demographics, and family composition, and state-level rules for eligibility.[15] Then, given person i's choice set J_i, the implied probability of purchasing coverage in market j, $s_{ijrt}^{J_i}$, is listed in equation (9).

$$s_{ijrt}^{J_i} = \begin{cases} \frac{exp(\alpha^i p_{irt}^j + \delta^j(d_i) + \xi_{tr}^j)}{\sum_{j' in J_i} exp(\alpha^i p_{irt}^j + \delta^{j'}(d_i) + \xi_{tr}^{j'})} & j \in J_i \\ 0 & j \notin J_i \end{cases} \tag{9}$$

Because neither the American Community Survey nor the Utah All Payer Claims data reports whether individuals were offered health insurance by their employer, I do not perfectly observe each individual's actual choice set J_i. Instead, I take employer offer rates from survey data, and I estimate a likelihood function which integrates over potential choice sets including and excluding employer-sponsored insurance. I get firm-level offer rates from the Medical Expenditure Panel Survey - Insurers Component, given employment status (full/part time) and average firm wage. I then calculate p_i^E, the probability that individual i is offered employer-sponsored coverage, as the probability that at least one of the two household heads is offered health insurance by their employer.

$$P(a_i^j) = p_i^E s_{ij}^{J_i^E} + (1 - p_i^E) s_{ij}^{J_i^{\neg E}} \tag{10}$$

The final likelihood function is listed in equation (10). In this equation, a_i^j is a an indicator that person i is observed purchasing insurance in market j. J_i^E is person i's choice set including employer-sponsored coverage, and $J_i^{\neg E}$ is their choice set excluding employer-sponsored coverage. In my baseline demand model, I estimate this equation using maximum likelihood.

5.1 Data

I estimate the demand model using two individual-level datasets. The Utah All-Payer Claims Dataset (APCD) includes administrative quality data on health insurance status, income, and health status of all individuals purchasing insurance in Utah between 2013-2015. The key limitations of this data are that it only includes data from one state, and is missing data on uninsured individuals. In contrast, the American Community Survey (ACS) includes data on insurance choice, income, and age across a representative sample of Americans across all U.S. states. To make use of the relative strengths of each dataset, I estimate

[15] I find state-level rules for Medicaid eligibility from the Kaiser Family Foundation's Annual Survey of Eligibility Rules, Enrollment and Renewal Procedures, and Cost-Sharing Practices in Medicaid and CHIP https://www.kff.org/medicaid/report/annual-updates-on-eligibility-rules-enrollment-and-renewal-procedures-and-cost-sharing-practices-in-medicaid-and-chip/

equation (10) using maximum likelihood using the American Community Survey Public Use Microdata Sample, which allows me to use the full extent of the large policy-driven variation in both prices and choice sets to identify demand. I then use the Utah APCD to estimate a cost model to show how costs vary across people and markets using detailed controls.

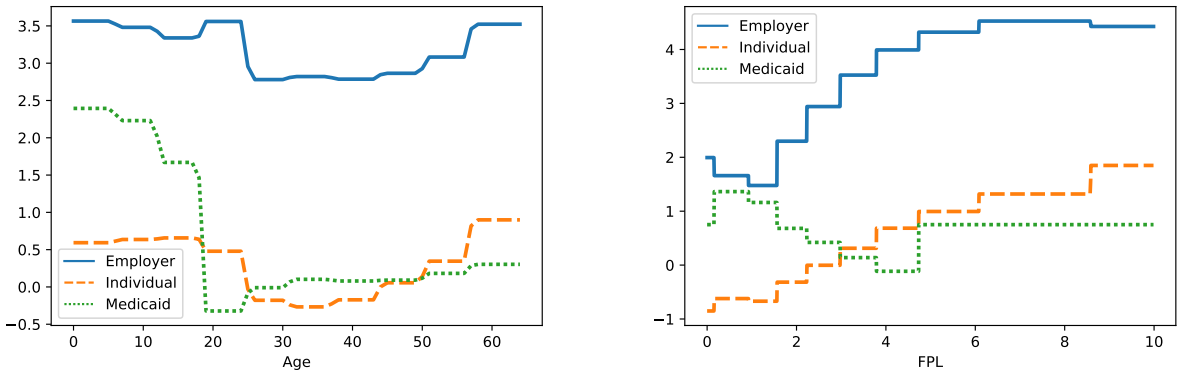
To estimate the demand model, I use the full ACS sample between 2015-2017, which was a period of relative stability in in insurance market policy. During this period, the two most important policy changes which affected the price of insurance and enrollment were the phase-out of reinsurance subsidies and the late expansion of Medicaid in 5 states, both of which I use for identification. I do not include data from 2014 when enrollment in Medicaid and the individual exchanges was ramping up but administrative delays meant that total enrollment may have been lower than ultimate demand in these markets. I also do not use data starting in 2018 in order to avoid modelling policies subsequently made by the Trump administration, like the elimination of cost-sharing subsidies and the elimination of the individual mandate, which affected health insurance markets starting in 2018 and 2019. My baseline specification includes all Americans under 65 years old who do not qualify for Medicare.

5.2 Demand Model Results

Figure 12: Model Parameters

(a) Age Coefficients

(b) Income Parameters



This figure reports coefficients from functions δ^j in equation (8). Each function is modelled with 10 dummy variables for age and income (as a fraction of the Federal Poverty Level) respectively. The y-axis is measured in the units of the discrete choice model, and is normalized against the outside option of non-insurance. Age coefficients are plotted for someone of average income, and income parameters are polotted for someone of average age.

My baseline demand model estimates equation (10) using maximum likelihood. I report estimates from the demand model in figure 12. I estimate the model in the American Community Survey. All of the

dependent variables are normalized to have a mean of zero and a standard deviation of 1. These figures report the coefficients from flexible demand functions with respect to both age and income (measured as a fraction of the Federal Poverty Line).

Looking first at the results with respect to age, I find that demand for insurance is high among children, low and stable for young adults, with increasing demand for insurance between age 40 and age 65 in all markets. There is stronger demand for both employer-sponsored and individual coverage among those aged 20-26 relative to those aged 26-50, likely because these individuals still have access to their parent's insurance policies under the Affordable Care Act.

Looking next at the results for income, I find that demand for employer-sponsored coverage and individual coverage is increasing in income, demand for Medicaid is decreasing in income. These results are consistent with the fact that, for low-income individuals individuals, the characteristics of Medicaid - low cost sharing but generally worse hospital and physician networks - are more attractive to lower income individuals.

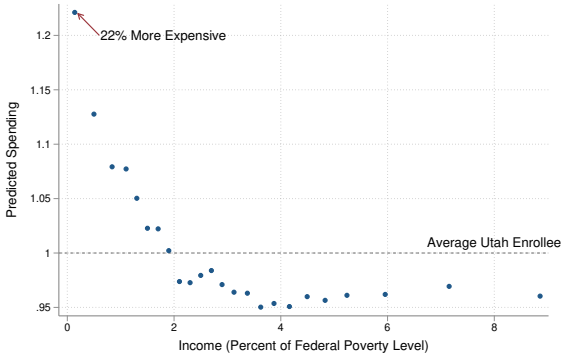
5.3 Cost Model Results

Figure 13: Cost Model Parameters

(a) Age/Gender Coefficients



(b) Income Parameters



Notes: This figure graphs coefficient estimates from equation (6) using data from the Utah All-Payer Claims Dataset on costs by demographic group. The dependent variable is total predicted health costs using the ACG model. The dependent variable is normalized such that spending for the mean person in Utah is equal to 1. I include dummy variables for each year of age interacted with gender, and 25 dummy variables for individuals of different income levels. All specifications include plan-specific fixed effects.

I estimate the cost model in equation (6) using data from the Utah All-Payer Claims Dataset. The dependent variable in this regression is predicted health insurance claims using ACG scores measured using the John Hopkins ACG model. Using ACG scores instead of actual claims is useful because ACG scores are based only on diagnosis codes and are not affected by negotiated prices which can vary across insurers and plans. As a result, I am able to ignore the potential effect of negotiated premiums or plan characteristics.

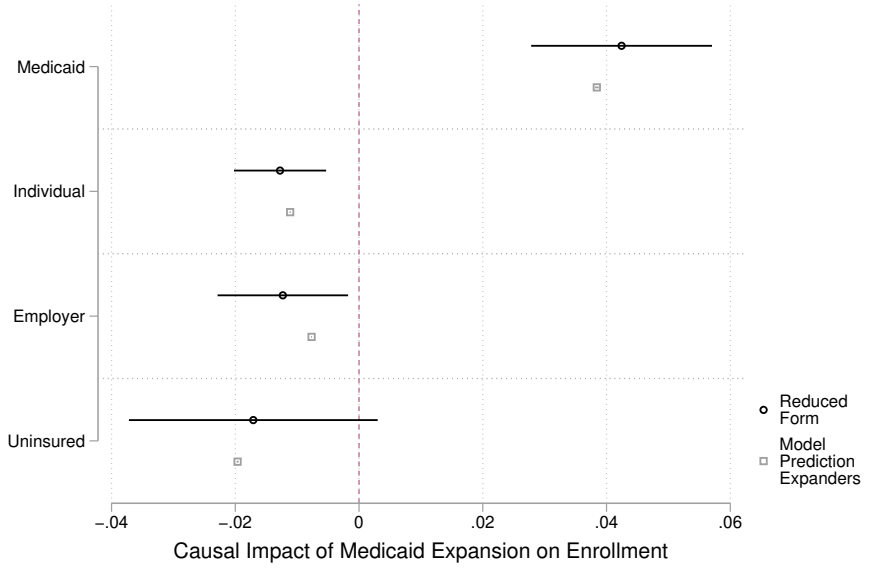
Second, I estimate these models including plan-level fixed effects, so differences in costs are estimated relative to individuals purchasing exactly the same insuruance plan.

Results from this exercise are shown in figure 13, which report coefficients from the model with respect to both income and demographics. These figures show that predicted health spending rises with age, and is higher for women relative to men for individuals of working age. For example, a 56-year old women is predicted to have 2.2 times higher spending than the average person in Utah. These figures also show that health status is regressive, with lower-income people having higher predicted health spending after controlling for age, gender and plan characteristics. The negative relationship between predicted health spending and income and is particularly apparent for those earning under 200% of the Federal Poverty Level.

In a supplementary analysis (available on request), I show that the results from this section generalize well to the overall population in the United States using data from the Medical Expenditure Panel Survey. Using this data, I show that both self-reported health status and the prevalence of major health conditions are generally regressive in income - that is, higher-income individuals are healthier than lower-income individuals after controlling for age and gender. I also replicate the general finding that spending rises with age, and is generally higher for women relative to men for individuals of working age.

5.4 Counterfactuals

Figure 14: Causal Impact of Medicaid Expansion, 2015, Model VS Reduced Form Moments

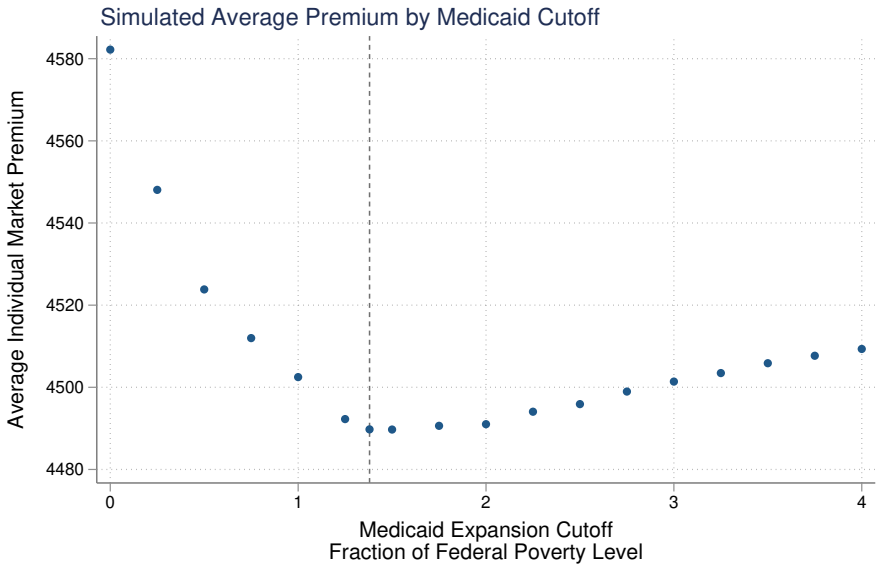


As a first counterfactual, I compare the predicted impact of Medicaid expansion according to the model

to the reduced-form moments I estimated comparing expansion to non-expansion states in Figure 14. In this figure, the reduced form moments are taken from Equation (2), comparing outcomes in 2012 (to measure pre-ACA levels) and 2015 (to measure post-ACA levels) of coverage by market. I use 2012 and 2015 to exclude the transition period from 2013-2014 where Medicaid enrollment was ramping up.I compare these reduced form moments to simulated moments from the model. The simulated moments are equal to average predicted enrollment in 2015 with full Medicaid expansion minus average enrollment in 2015 assuming no expansion, measured in expansion states (to measure treatment on the treated). Note that non-expansion states tend to have higher predicted effects of Medicaid expansion because they started from a lower base than expansion states.

Figure 14 shows that the reduced-form and structural moments are very similar, suggesting that the empirical model correctly predicts the degree of substitution between Medicaid and other markets. Since the empirical model was estimated using data from 2015, 2016 and 2017, and the reduced form experiment compares outcomes between 2012 and 2015, these findings are not mechanically related.

Figure 15: Different simulated Medicaid expansion



As a second conterfactual exercise, I ask the following question: What would the effect of expanding Medicaid to a different threshold than the one set in the Affordable Care Act? In the Affordable Care Act, policymakers chose 138% of FPL as the threshold for Medicaid eligibility, and assumed that those with a higher income would purchase private insurance on the individual insurance exchanges. In figure 15, I show 18 counterfactual simulations of different levels of Medicaid expansion, from 0% (where Medicaid eligibility

is determined by pre-2014 state-level rules) to 400% of FPL. On the y-axis, I plot the average premiums for a plan purchased on the individual market. This figure shows that, from the perspective of those operating on the individual market, prices are lowest when Medicaid expansion is approximately at 138%, which is what was set out in legislation.

6 Conclusion

In this study, I present a framework which shows how different eligibility conditions across health insurance markets can lead to differences in the price of insurance in these markets. I show that these price differences can be very large in practice: For example, in Tennessee, purchasing a plan in the individual market is over 70% more expensive than an equivalent plan in the small firm market. Then, I show that policies which affect one market can spill over to other markets through risk sorting. For example, I find that the Affordable Care Act's Medicaid expansion reduced premiums by 9.7% on the individual markets, which increased the number of people purchasing individual insurance. Finally, I propose a structural model which estimates how these policies operate at the individual level.

These findings are very important to consider as politicians debate different policies to expand health insurance access in the United States. As of 2019, 9.2% of the the United States population still does not have health insurance.[16] However, there is strong support in Congress for preserving existing health insurance options, because many who are currently insured report they are satisfied with their existing plan.[17] As a result, it seems likely that even if the rules governing insurance markets change, the existing system of multiple competing markets will stay for the foreseeable future. Understanding how this system affects the distribution of what people pay for coverage and how policy affects this distribution is thus important.

[16] Source: American Community Survey

[17] Medicaid: https://jamanetwork.com/journals/jamainternalmedicine/fullarticle/2643347 Employer-sponsored: http://files.kff.org/attachment/Report-KFF-LA-Times-Survey-of-Adults-with-Employer-Sponsored-Health-Insurance

References

Aizawa, N., & Fang, H. (2013, January). *Equilibrium labor market search and health insurance reform* (Working Paper No. 18698). National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w18698 doi: 10.3386/w18698

Aizawa, N., & Fu, C. (2020, January). *Interaction of the labor market and the health insurance system: Employer-sponsored, individual, and public insurance* (Working Paper).

Akerlof, G. (1970). The market for 'lemons': Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, *84*(3), 488-500.

Azevedo, E. M., & Gottlieb, D. (2017). Perfect competition in markets with adverse selection. *Econometrica*, *85*(1), 67-105. Retrieved from https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA13434 doi: 10.3982/ECTA13434

Berry, S. T., & Haile, P. A. (2020). *Nonparametric identification of differentiated products demand using micro data* (Working Paper).

Bundorf, M. K., Levin, J., & Mahoney, N. (2012, December). Pricing and welfare in health plan choice. *American Economic Review*, *102*(7), 3214-48. Retrieved from http://www.aeaweb.org/articles?id=10.1257/aer.102.7.3214 doi: 10.1257/aer.102.7.3214

Clemens, J. (2015, April). Regulatory redistribution in the market for health insurance. *American Economic Journal: Applied Economics*, *7*(2), 109-34. Retrieved from http://www.aeaweb.org/articles?id=10.1257/app.20130169 doi: 10.1257/app.20130169

Cutler, D. M., & Gruber, J. (1996, 05). Does Public Insurance Crowd out Private Insurance? *The Quarterly Journal of Economics*, *111*(2), 391-430. Retrieved from https://doi.org/10.2307/2946683 doi: 10.2307/2946683

Duggan, M., Goda, G. S., & Jackson, E. (2019). The effects of the affordable care act on health insurance coverage and labor market outcomes. *National Tax Journal*, *72*(2), 261-322. doi: dx.doi.org/10.17310/ntj.2019.2.01

Einav, L., Finkelstein, A., & Cullen, M. R. (2010, 08). Estimating welfare in insurance markets using variation in prices. *The Quarterly Journal of Economics*, *125*(3), 877-921. Retrieved from https://doi.org/10.1162/qjec.2010.125.3.877 doi: 10.1162/qjec.2010.125.3.877

Finkelstein, A., Hendren, N., & Shepard, M. (2019, April). Subsidizing health insurance for low-income adults: Evidence from massachusetts. *American Economic Review*, *109*(4), 1530-67. Retrieved from http://www.aeaweb.org/articles?id=10.1257/aer.20171455 doi: 10.1257/aer.20171455

Frean, M., Gruber, J., & Sommers, B. D. (2017). Premium subsidies, the mandate, and medicaid expansion:

Coverage effects of the affordable care act. *Journal of Health Economics*, *53*, 72 - 86. Retrieved from http://www.sciencedirect.com/science/article/pii/S0167629616302272 doi: https://doi.org/10.1016/j.jhealeco.2017.02.004

Geruso, M., Layton, T. J., McCormack, G., & Shepard, M. (2019, September). *The two margin problem in insurance markets* (Working Paper No. 26288). National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w26288 doi: 10.3386/w26288

Gruber, J., & Simon, K. (2008). Crowd-out 10 years later: Have recent public insurance expansions crowded out private health insurance? *Journal of Health Economics*, *27*(2), 201 - 217. Retrieved from http://www.sciencedirect.com/science/article/pii/S0167629607000963 doi: https://doi.org/10.1016/j.jhealeco.2007.11.004

Hackmann, M. B., Kolstad, J. T., & Kowalski, A. E. (2015, March). Adverse selection and an individual mandate: When theory meets practice. *American Economic Review*, *105*(3), 1030-66. Retrieved from http://www.aeaweb.org/articles?id=10.1257/aer.20130758 doi: 10.1257/aer.20130758

Handel, B., Hendel, I., & Whinston, M. D. (2015). Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica*, *83*(4), 1261-1313. Retrieved from https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA12480 doi: 10.3982/ECTA12480

Heim, B. T., Hunter, G., Isen, A., Lurie, I. Z., & Ramnath, S. P. (2019, November). *Income responses to the affordable care act: Evidence from a premium tax credit notch* (Tech. Rep.). Retrieved from https://www.dropbox.com/s/t39nss2eh4f40r0/Income%20Responses%20to%20the%20Affordable%20Care%20Act.pdf?dl=0

Kucko, K., Rinz, K., & Solow, B. (2018, March). *Labor market effects of the affordable care act: Evidence from a tax notch* (Working Paper). SSRN. Retrieved from https://ssrn.com/abstract=3161753 doi: 10.2139/ssrn.3161753

Peng, L. (2017). How does medicaid expansion affect premiums in the health insurance marketplaces? new evidence from late adoption in pennsylvania and indiana. *American Journal of Health Economics*, *3*(4), 550-576. Retrieved from https://doi.org/10.1162/ajhe_a_00088 doi: 10.1162/ajhe_a_00088

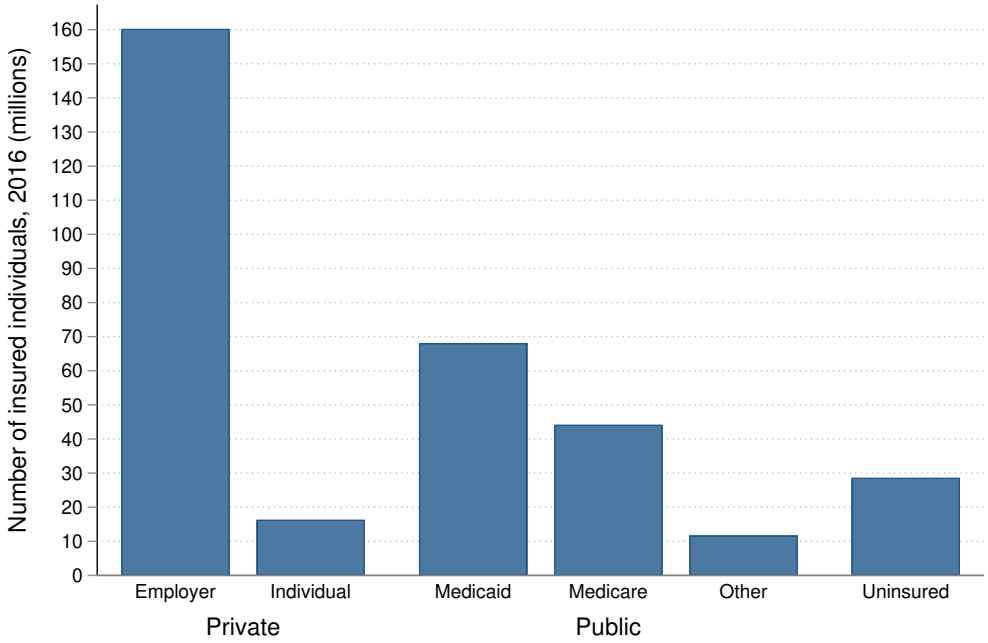
Saltzman, E. (2017). *The welfare implications of risk adjustment in imperfectly competitive markets* (Working Paper). University of Pennsylvania.

Tebaldi, P. (2020). *Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca* (Working Paper).

Tebaldi, P., Torgovitsky, A., & Yang, H. (2019, May). *Nonparametric estimates of demand in the california health insurance exchange* (Working Paper No. 25827). National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w25827 doi: 10.3386/w25827

A United States Health Insurance Regulations

Appendix Figure 1: Enrollment by Health Insurance Market Segment in the United States, 2016



Sources: Individual and small group totals taken from the Medical Loss Ratio data, published by The Center for Consumer Information & Insurance Oversight. The other categories are taken from the American Community Survey 1-year PUMS file. "Other" public insurance includes Tricare and the VA. Since individuals with government insurance often erroneously report non-group insurance, the difference between reported individual plans in the ACS and reported individual plans in the Medical Loss Ratio data is classified as "other" public plans. Totals measured as of 2016 because the Medical Loss Ratio data does not separately report ACA and non-ACA market totals after the completion of the risk corridor program in 2016.

The market for U.S. health insurance is divided into segments, each governed by separate legislation. The private market has three main segments: Firms with 2-50 employees purchase insurance in the small group market,[18] firms with over 50 employees purchase insurance in the large group market, and individuals who don't purchase insurance through their employer participate in the individual market. Public health insurance is provided primarily through Medicaid and Medicare. Medicaid is a state-level program that covers individuals who meet a means test and those who qualify for Supplemental Security Income. Eligibility threshholds vary by state and applicant type, with parents, children, pregnant women, and seniors generally having higher eligibility threshholds. Medicare is an entitlement program offered to all those who are over 65 years of age, as well as those who receive Social Security Disability Income and individuals diagnosed with End Stage Renal Disease. Figure 1 shows total population by market segment.

[18] The small group market includes firms with up to 100 full-time employees in California, Colorado, New York, and Vermont.

The large firm market is regulated federally. The regulatory framework is defined in the Employee Retirement Income Security Act of 1974 (ERISA), which was subsequently amended by the Consolidated Omnibus Budget Reconciliation Act of 1986 (COBRA), the Health Insurance Portability and Accountability Act of 1996 (HIPAA), and the Patient Protection and Affordable Care Act (ACA). At the employer-level, the price of healthcare is experience-rated, and there are few restrictions on what insurance companies can charge employers. However, employers are barred under HIPAA from charging employees more for "health factors." HIPAA also prevents firms from denying insurance on the basis of pre-existing conditions (as long as the prospective employee had coverage already), and the ACA expanded this rule to cover all individuals with pre-existing conditions. As a result, I model the large group insurance market as being community-rated within a firm, where each employer operates its own separate risk pool.

Prior to the 2010 passage of the Patient Protection and Affordable Care Act (ACA), prices in the individual and small group markets were regulated at the state level. While exceptions exist, these markets were experience-rated in most states prior to 2014, meaning insurers could vary premiums based on health status.[19] After the ACA was enacted in 2014, all new plans sold in the individual and small group markets were community rated, meaning prices did not depend on health status. However, existing "grandfathered" plans that had existed prior to 2010 were allowed to continue to operate under pre-existing state-level pricing regulations.

A.1 The Affordable Care Act

The Patient Protection and Affordable Care Act (ACA) was passed in 2010, with the goal of expanding overall health insurance coverage while minimizing disruption to existing markets. When he was selling this plan to the public, President Barak Obama repeatedly said "If you like your health care plan, you can keep it."[20] As a result of this promise, the ACA preserved the distinctions between the large group, small group, and individual markets, and allowed most existing plans to continue under "grandfathered" status. As I will show, preserving these market definitions has lead to wide variation in prices for coverage based on which market a consumer uses to purchase health insurance.

Starting in 2014, all new plans sold in the US are guaranteed issue, meaning insurance can no longer be denied to individuals with pre-existing conditions, and premiums cannot vary according to an individual's health status. Because the ACA worked within the confines of existing market definitions, these provisions were enacted separately in each market. The individual and small group markets were each community

[19] States with a community-rated individual market prior to 2014 include: Maine, Massachusetts, New Jersey, New York, Oregon, Vermont, and Washington. For a detailed description of the adoption of community-rating prior to 2014, see Clemens (2015))

[20] See https://www.politifact.com/truth-o-meter/article/2013/dec/12/lie-year-if-you-like-your-health-care-plan-keep-it/

rated, meaning each had listed prices which varied only by age (using a common 3:1 age rating curve) and geography.[21] In the large group market, the ACA strengthened regulations outlawing price differences at a given employer based on health status. While all new plans had a constant price within a market (or within a given firm), there were no restrictions on the relative cost of health insurance across market segments. In practice, a single health insurance product could have wide ranging sticker prices based only on whether it was purchased in the individual or small group markets, or by a healthier or sicker large firm.

Finally, to cover full-time workers at large firms, the ACA included an employer mandate which required all large firms to offer affordable coverage to their full-time workforce.

Appendix Table 1: Medicaid expansion decisions by state

Date	**State**
January, 2014	Arizona, Arkansas, California, Colorado, Connecticut, Delaware, District of Columbia, Hawaii, Illinois, Iowa, Kentucky, Maryland, Massachusetts, Minnesota, Nevada, New Jersey, New Mexico, New York, North Dakota, Ohio, Oregon, Rhode Island, Vermont, Washington, West Virginia
April, 2014	Michigan
August, 2014	New Hampshire
January, 2015	Pennsylvania
February, 2015	Indiana
September, 2015	Alaska
January, 2016	Montana
July, 2016	Louisiana
January, 2019	Maine, Virginia
January, 2020	Utah, Idaho
Date TBD	Nebraska, Missouri, Oklahoma
No Expansion	Alabama, Florida, Georgia, Kansas, Mississippi, North Carolina, South Carolina, South Dakota, Tennessee, Texas, Wisconsin, Wyoming

Source: Kaiser Family Foundation (https://www.kff.org/medicaid/issue-brief/status-of-state-medicaid-expansion-decisions-interactive-map/)

A.2 Price Setting Under the Affordable Care Act

Starting in 2014, the Affordable Care Act mandated that all new health insurance products sold in the individual market be qualified health plans. A qualified health plan must be community-rated, meaning that prices are posted and all Americans may purchase a plan at the market rate with no restrictions based on health status. Qualified health plans are sold in geographic health insurance markets, called rating areas. Each rating area is generally a set of one or more counties, and states vary in the number of rating areas allowed. Each insurer then offers a variety of types of plans within each rating area, which can vary based on things like network reach or cost-sharing characteristics. Each of these plans is assigned a base price b_j.

[21] Note: Prices were also allowed to vary by whether a person smoked tobacco, but in practice few insurers charged different prices to smokers and non-smokers, and price differences between these plans tend to be small

To calculate the final price paid by consumers for a plan, the base price b_j is then modified based on an age-rating curve, and potentially by cost-sharing subsidies. In most states, age adjustment follows a 3:1 curve, meaning that a 64-year old pays 3 times what a 20-year old would to purchase the same health insurance plan. Let α_i be the age adjustment, with $\alpha_i = 3$ for a 64-year old and $\alpha_i = 1$ for a 21-year old.[22] The Affordable Care Act also allows plans to charge a higher rate to tobacco users, though this provision is only available in certain states (and many plans do not charge a higher rate even though they can). In this paper, I ignore the tobacco provision.

For low to middle income individuals, the ACA also provided subsidized access to the individual market for all those who between 100%- 400% of the Federal Poverty Line. These subsidies were benchmarked to the cost of the second lowest priced silver plan in each region, which is also called the benchmark plan. Specifically, each family received a subsidy such that the total cost of the benchmark plan would not exceed that cap as a fraction of the federal poverty line. In 2020, these caps varied from 2.06% of the Federal Poverty Line (for a family earning between 100%-133% of the FPL) to 9.76% (for a family earning between 300-400% of the Federal Poverty Line). In practice, this means that any family receiving subsidies is fully insolated against increases in the price of a benchmark plan, since subsidies increase 1-for-1 with increases in the premium.

In sum, on-exchange plans are priced according to the following equation:

$$P_{ij} = max(\alpha_i b_j - s_i, 0) \tag{11}$$

In this equation b_j is the base price of plan j in consumer i's market, and α_i is the age adjustment. s_i is the premium subsidy received by person i. Premium subsidies are determined by the following equation:

$$s_i = max(P_{ij^*} - c(p_i), 0) \tag{12}$$

Where P_{ij^*} is the price of the benchmark plan j^* for person i, and $c(p_i)$ is the premium cap for person i based on their income.

A.3 Changes after the Affordable Care Act

The history of the risk corridor program is not complete without a discussion of its political context. In the Affordable Care Act, there was no requirement that risk corridor payments net to zero. In practice, the number of insurance companies with claims-to-premiums ratio of over 82.4% was much higher than

[22] Note that Washington DC, Massachusetts, Minnisota, Utah, New York, and Vermont each maintain their own separate age-rating curves.

the number of insurance companies with claims to premiums ratio under 77.6%, meaning that the bill required additional funds. However, the 2015 and 2016 appropriations bills included a rider that required risk corridor payments to net to zero, which in practice led to much lower payments than insurance companies were promised. Most firms operating on the individual insurance exchange in 2014 and 2015 lost money, largely because they did not receive the risk corridor payments they were promised. The rider was challenged in the courts, and in *Maine Community Health Options v. United States* the U.S. Supreme Court ruled in April 2020 that this particular rider was not legal, and ordered the government to pay back the promised payments. While insurance companies could (and did) update their prices in 2016 in response to this policy change, the rider was added too late for them to update their prices in 2015, and it is my understanding that this particular policy was unexpected from the insurers' perspective. In this section, I use promised risk corridor payments (as opposed to actual payments after taking the rider into account) from 2015 (before the rider was announced).

Prior to 2014, Medicaid eligibility was determined at the state level, and states varied widely in who was eligible. Under the ACA, eligibility for Medicaid was standardized to include (at a minimum) all those earning under 133% of the Federal Poverty Line. This represented an expansion of Medicaid eligibility in most states, especially among men and adults without children.

B Measuring Premiums

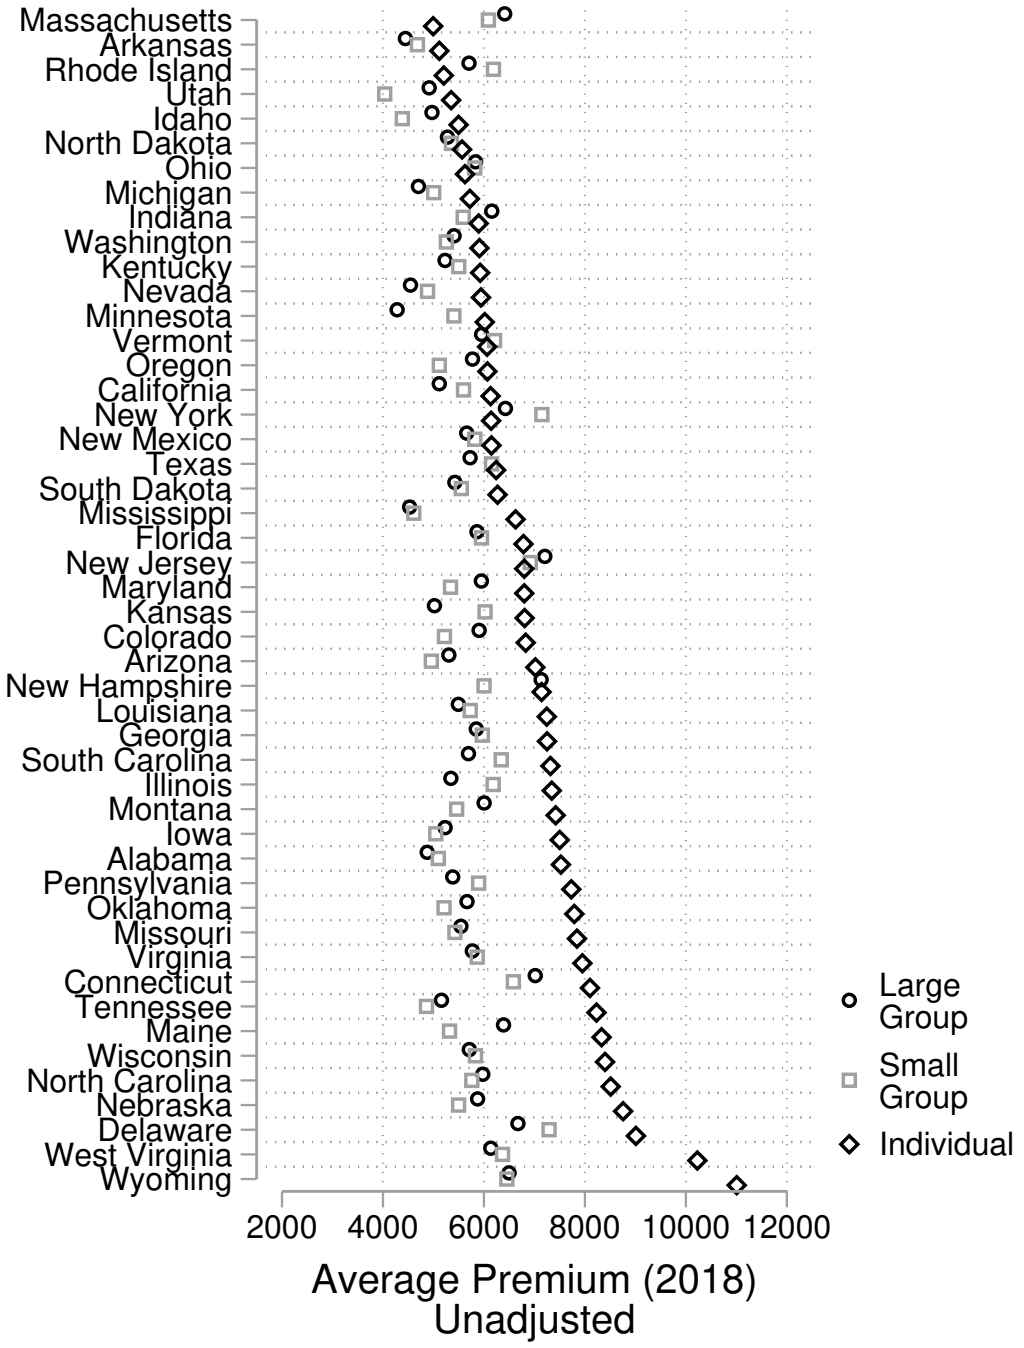
Existing research has not been able to document price differences for a given product across market segments in the United States. This is principally because premiums in both the grandfathered and large group health insurance markets are negotiated separately for each contract, and these contract details are typically treated as trade secrets. As a result, premiums for health insurance products do not appear in most administrative data sources accessible to researchers, and the most of our knowledge about health insurance premiums comes from surveys, which often lack data on the characteristics of the enrolled population. This paper will suffer from the same issues: I do not also do not observe representative contract-level data from the large-firm market linked to plan and employee characteristics. What I will provide is the best possible approximation of these price differences using publicly available data. I will show that prices across these markets vary widely, both within and across states.

Approximating price differences between the ACA compliant and non-compliant markets is more difficult because prices in the non-compliant markets are treated as industry secrets. I will document aggregate price differences across markets using the the Medical Loss Ratio data. This data includes aggregate premiums and enrollment by state, insurance company and market segment. Because this data is based on administrative records, it avoids the possibility of bias that might exist in survey measures. The downside is that these are market aggregates without a breakdown by plan-type and health status. As a result, I will also present a version where I have controlled for an index of plan generosity and demographic characteristics by state and market.

Figure 2 reports raw average premiums from the Medical Loss Ratio data in the individual, small group and large group markets. This figure shows wide variation between states in premiums paid on the individual market. In Massachusetts and Rhode Island, the average policy purchased on the individual market was actually cheaper than the average policy purchased by a small or large firm. In Wyoming, the average policy on the individual market was close to double the price of the average policy in the large group market. Since these figures are raw means, prices may vary across markets for many reasons, including both adverse selection and also differences in the average plan quality across markets.

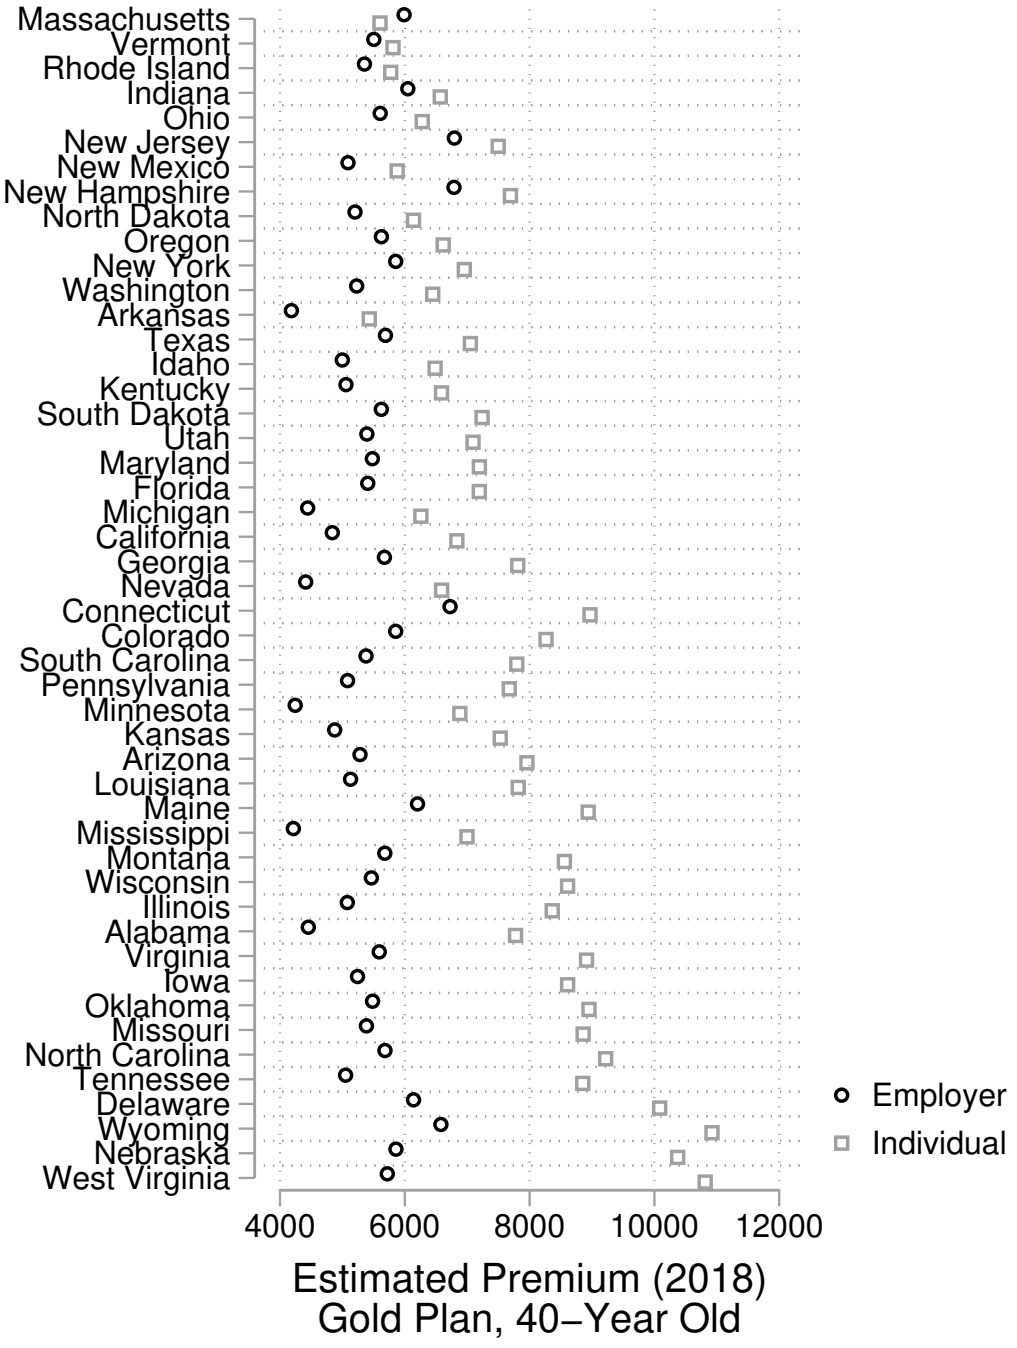
To account for differences across market in plan quality, I also provide adjusted figures in Figure 3. This figure takes the raw premiums in Figure 2 and divides them by both an demographic index and a plan generosity index. The result should be read as a back-of-the envelope estimate of the average difference in prices for an identical product. Relative to the unadjusted figures, the gap between the individual and large group market is higher. The main reason is that plans sold on the individual market are less generous, with higher deductibles and coinsurance rates on average.

Appendix Figure 2: Average Premium Paid per Person and Year, by State and Market in 2018.



This figure reports average premium paid by person, market and state in 2018. Small group and individual markets include both grandfathered and non-grandfathered plans (note: Very few individual grandfathered plans exist in 2018, but some states continue to have a large grandfathered small group market). The large group market includes only fully-insured large firms. No adjustments are made for individual demographics or plan quality. Source: Author calculations using the Medical Loss Ratio data, published by The Center for Consumer Information & Insurance Oversight

Appendix Figure 3: Average Premiums for Individual and Large Group Markets by State.



This figure reports average premium paid by person, market and state in 2018. The individual market include both grandfathered and non-grandfathered plans, but the grandfathered individual market is of neglibile size in most states in 2018. The large group market includes only fully-insured large firms. Average premiums are adjusted for average plan characteristics and individual demographics, to approximate the price of a Gold plan for a 40-year old individual (see text for details). Source: Author calculations using the Medical Loss Ratio data, published by The Center for Consumer Information & Insurance Oversight

To construct the age index, I calculate predicted spending by age and gender (calculated using the John Hopkins ACG model) and apply it to the distribution of age and gender in the American Community Survey, for all individuals participating in the individual and large group markets by state. I then normalize each state's age index so that the US average is 1, with younger states having an index below 1, and older states having an index above 1. To construct the plan quality index, I assume a linear relationship between premiums and deductibles at the plan level. This allows me to estimate an index of plan quality by sector using only data on average deductibles in that sector. Using an linear model of premiums using HIX Compare, I estimate that a 1000$ increase in deductible is associated with a 7% decrease in average premiums. I then apply these coefficients to average deductible by state using data from the MEPS-IC (insurer component) survey of health plans. For on-exchange plans, I calculate average plan quality using purchase data by metal tier using the OEP public use files. I normalize each index to be relative to the price of a generic gold plan.

While each of these datasets have their weaknesses, both consistently show that the small group and large group markets are unambiguously cheaper than the individual market in most states. By contrast, the small and large group markets appear to have similar average premiums. In the next section, I will look at how Medicaid expansion affected relative premiums across the individual and small group markets. Because the large group data is much worse, I will not be able to perform a similar analysis across the individual and large group markets. Nevertheless, because the small and large group markets appear pretty similar, I expect premium differences between the individual and large group markets to be very similar to premium differences between the individual and small group markets.

C Event-Study with Plan-Specific Controls

To estimate the degree of adverse selection in each individual market, I take advantage of the fact that the Affordable Care Act created separate individual and small group markets for health insurance. The individual and small group markets operate in separate risk pools, but share most regulations on plan quality and pricing. The same insurers tend to participate in both markets, and similar (often identical) plans are offered. As a result, the difference in premiums is a useful proxy for how adversely selected the individual market is in each state.

Let $P_{j,t,g,m}$ be the premium for plan j in year t at location s and in market m, where m is either "Individual ACA" or "Small Group ACA." I then estimate the regression:

$$\begin{aligned} log(P_{j,t,g,m}) = \beta_t 1\{m = \text{Individual ACA}, g \in \text{Non-Expansion State}\} \\ + \delta_t 1\{m = \text{Individual ACA}, g \in \text{Expansion State}\} \\ + \psi_{t,g,m} + \epsilon_{j,t,g,m} \end{aligned} \tag{13}$$

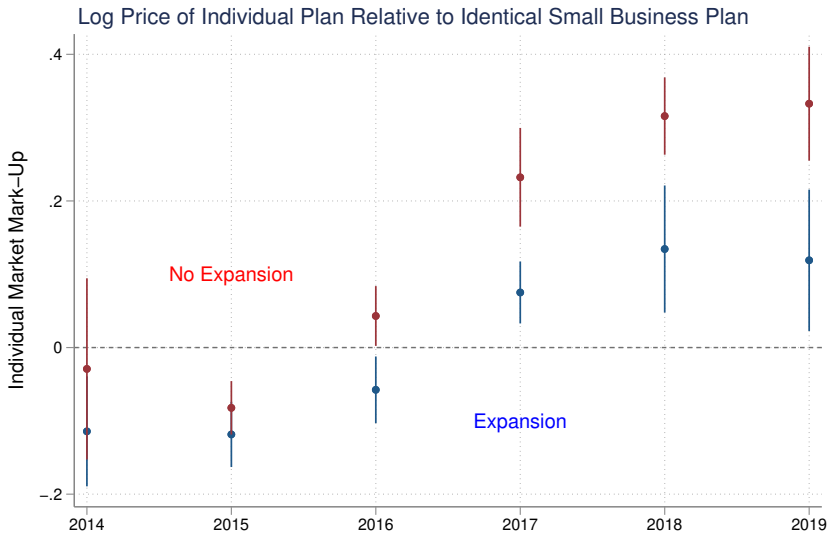
In this equation, $\psi_{t,g,m}$ is a separate fixed effect for the interaction of plan,[23] geographic rating area, and year. β_t is thus an estimate of the average difference in prices between a individual and small group plan in a non-expansion state, and δ_t is an estimate from an expansion state. This regression is weighted so each state has equal weight, and standard errors clustered at the state level. Results are reported in Figures 4 and 5

In Figure 4, I report coefficient estimates from Equation (13) using the sticker price of insurance purchased by a 40-year old individual. This figure shows clearly that prices in the individual market were actually lower than in the small group market in 2014 and 2015. The reason prices were lower was that the individual market received large risk adjustment and risk corridor subsidies relative to the small group market in 2014-2016 (and especially 2014-2015). To get a better measure of adverse selection, I also calculate a measure of adjusted premiums by adding back anticipated subsidies, which I take from the Medical Loss Ratio data. This measure of adjusted premiums represents the total price paid for each health insurance product, including both premiums paid by individuals and subsidies paid by the government. Results for adjusted premiums are listed in Figure 5.

These regressions suggest that the individual market has always been adversely selected relative to the

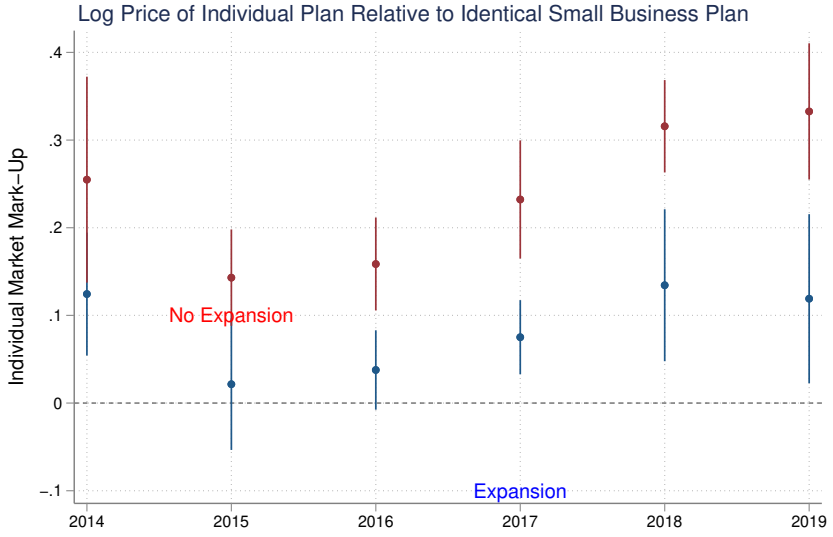
[23] Note that, while insurance companies often offer near-identical plans in both the individual and small group markets, there are often minor differences in the plans offered. In the primary specification, all products of a given metal tier, insurer, and product type (HMO/PPO/POS) are lumped together as the same plan. I have run these results using more stringent definitions of plan (eg: Shares all characteristics in the data as identical). When I do this, the sample size drops, but the qualitative and quantitative results are robust.

Appendix Figure 4: Log Price of Individual Plan Relative to Small Group Plan (Raw)



This table plots regression coefficients from Equation (13). The dependent variable is the log price of a plan purchased by a 40-year old. The unit of observation is at the plan/rating area level, and the sample includes all plans of all metal tiers except catastrophic plans. The sample only includes plans which are sold in both the individual and small group markets. Two different products are the same "plan" if they share an insurer, network, network type (PPO or HMO), and metal tier. Regressions are weighted so each state has equal weight in the regression and standard errors are clustered at the state level.

Appendix Figure 5: Log Price of Individual Plan Relative to Small Group Plan (Adjusted for Subsidies)



Coefficients are from Equation (13). All notes are the same except as Figure 4 except that the dependent variable is adjusted prices. Adjusted prices are calculated by multiplying the sticker price by the an inflation index, equal to $\frac{1}{1-subsidy}$, where the subsidy is calculated as the net combined expected risk corridor and reinsurance subsidies as a fraction of premiums reported in the Medical Loss Ratio files by market and state in 2014-2016.

small firm market. In expansion states, premiums plus subsidies in the individual market have fluctuated between 5 and 15 log points higher than in the small group market. In non-expansion states, premiums plus subsidies have been between 15 and 35 log points higher. Average premiums paid by consumers were held artificially low between 2014-2016 because of temporary risk corridor and reinsurance subsidies, and the price paid by consumers was similar across expansion and non-expansion markets. Since 2016, prices rose in all markets, and rose more drastically in non-expansion states.