Market Segmentation and Competition in Health Insurance *

Michael J. Dickstein, Kate Ho, Nathaniel Mark

October 18, 2021

Abstract

In the United States, households obtain health insurance through distinct market segments. We explore the economics of this segmentation by comparing coverage provided through small employers versus the individual marketplace. Using data from Oregon, we find households with group coverage spend 26% less on covered health care than households with individual coverage yet face higher markups. We develop a model of plan choice and health spending to estimate preferences in both markets and evaluate integration policies. In our setting, pooling can both mitigate adverse selection in the individual market and benefit small group households without raising taxpayer costs.

*We thank conference and seminar participants at New York University and the University of Pennsylvania for helpful feedback. We also received valuable input from Eduardo Azevedo, Gautam Gowrisankaran, Robin Lee, Victoria Marone, and Pietro Tebaldi. We thank Quan Le, Gabriel Lesnick, and Nicole Ng for excellent research assistance.

Dickstein: New York University and NBER, michael.dickstein@nyu.edu; Ho: Princeton University, NBER, and CEPR, kate.ho@princeton.edu; Mark: U.S. Department of Justice, nathaniel.mark@usdoj.gov. The views expressed herein do not represent the views of the U.S. Department of Justice. Throughout this paper, we use the term "market" in ways that do not necessarily reflect the product and geographic boundaries of antitrust markets.

1 Introduction

In the United States, households obtain health insurance through payers in distinct market segments. In 2017, for example, 56% of the population obtained coverage through an employer, 43% used a government plan, and 16% purchased a plan individually. Roughly 9% lacked coverage.[1] This heterogeneity in coverage sources contrasts with the insurance market design in other developed countries. Single-payer public insurance systems operate in the United Kingdom, Canada, and Australia, for example. In Germany, France, and Israel, households receive coverage through large health plans, often private and non-profit, that compete with each other to cover consumers across a range of ages and employment groups (Tikkanen et al., 2020).

We explore how the segmented system in the U.S. affects the degree of adverse selection, the level of insurance premiums, and ultimately consumer surplus. We focus on the potential effects of combining the small group market, where households obtain coverage through small employers, with the individual market, where households purchase insurance directly from marketplaces or brokers. These two markets offer an ideal laboratory to study the economics of segmentation. Following the implementation of the Affordable Care Act (ACA) in 2014, insurers must standardize plan designs in the two markets in similar ways. Moreover, recent policy changes have moved the markets closer to pooling. Massachusetts, Vermont, and the District of Columbia already require common premiums in the two segments (Hall and McCue, 2018). Further, after a federal regulatory change in 2020, employers have the option to shift their employees' insurance coverage to the individual market. Using new individual-coverage health reimbursement arrangements (ICHRAs), employers fund tax-exempt accounts through which employees purchase individual coverage. If adopted widely, ICHRAs could transform the U.S. insurance system, eliminating the employer's role in plan choice.

In pooling the markets, two economic mechanisms operate that may generate opposing welfare effects. First, pooling can alter the extent of adverse selection. The change in selection depends on how the distribution of household characteristics—including family size, income, underlying health

[1] These shares sum to over 100%, as households often receive coverage from two or three sources simultaneously or sequentially in a year (U.S. Census Bureau, 2018).

needs, risk aversion, and the propensity for moral hazard—differ between the two populations being pooled. Second, under small group coverage, employers and brokers act as agents for eligible employees, typically choosing one or two plans to offer to households (AHRQ, 2016). The plan choice itself may act as a recruiting tool for the employer or may be skewed by a broker's incentives to recommend costly plans that pay higher sales commissions. Households insured in the small group market also do not face the full cost of premiums, both because of employer subsidies and because insurance is a tax-advantaged employee benefit. Insurer markups may be higher as a result of these frictions.

We quantify the welfare consequences of market segmentation using data from Oregon's insurance markets in years 2014 through 2016. Our all-payers claims data from Oregon includes information on premiums, household enrollment status by year and market segment, and health spending. We also observe household transitions between different types of private insurance, including between individual coverage and group coverage.

We begin by examining health spending and premium levels in Oregon's small group and individual insurance markets. In our sample, we observe that small group enrollees spend 26% less on covered health services, conditional on plan generosity. However, these enrollees face 7% higher age-specific gross premiums than they would in the individual market.[2] Taken together, these facts indicate that insurer markups over health care costs are higher in the small group market.

When comparing plan choice, we find important differences in the two populations. In the individual market, plans in the more generous gold tier appear highly adversely selected, with high enrollee spending leading to high premiums and low market share for these plans (Cutler and Reber, 1998). Only 14% of insured households choose gold plans; these consumers incur an average of $932 per month in covered health spending. The equilibrium outcome is quite different in the small group market, where nearly 36% of enrollees choose gold plans. Enrollees in these plans also spend less on health care than their counterparts in the individual market, averaging $586 per month in spending per household.

[2]We compute differences in health spending and premiums between the two market segments for plans in the same metal tier–i.e. with the same actuarial value. The overall statistic is a weighted average of these differences, weighted by the market share of each metal tier in the sample.

These descriptive analyses suggest that both economic mechanisms in our setting—the extent of adverse selection and the role of employer agency—may point to welfare gains from pooling. Individual market enrollees could benefit from a reduction in adverse selection after the entry of healthier small group households. Small group enrollees may also benefit from more competitively priced individual market plans. The overall welfare change, however, depends critically on the equilibrium premiums that would result in the pooled market. The equilibrium itself depends on households' characteristics and preferences for enrollment in particular plans. Given this theoretical ambiguity, we develop a model of demand and supply to recover these preferences, predict equilibrium premiums, and quantify changes in consumer welfare.

Our demand model predicts households' plan choices and subsequent medical care utilization (Einav et al., 2013). From the estimated model, we quantify the unobserved health needs of each household; their risk aversion; and also the degree of moral hazard in response to plan generosity. Importantly, we allow the preferences of households in the current small group and individual insurance markets to differ. If we instead assumed the two groups had the same preferences ex ante, we would limit the scope for our welfare analysis and might produce misleading predictions if their preferences in fact differed.

Estimating preferences in the individual market is straightforward: we observe each household's choice set and its realized choice of plan. However, the problem is more challenging in the small group market because of the involvement of employers and brokers in plan choice. To distinguish the preferences of households separately from those of the employer agent, we identify small group households who are forced to switch out of their plan when their employer cancels group coverage. We study the subset of these "forced switchers" who do not regain coverage through other employers or through public insurance, i.e., those whose only option is to purchase coverage on the individual market. We estimate the preferences of the forced switcher group as a proxy for the entire small group's preferences. When we estimate our demand model using alternately the individual market and forced switcher samples, we recover distinct preferences, particularly in the tendency for moral hazard.

On the supply side, we estimate a premium-setting equation. Consistent with observed premiums and medical claims in our data, we assume that insurers in the individual market set premiums equal to costs on a plan-by-plan basis, accounting for net transfers the plans receive for risk adjustment and other government subsidies. This assumption allows us to estimate the total costs of each plan, including administrative costs, using a two-stage least squares approach. Combining this measurement with a model of competition and price-setting among insurers that follows Azevedo and Gottlieb (2017), we predict the effect of pooling the small group and individual markets on enrollment, costs, and consumer surplus. In our main analysis, we model a perfectly competitive individual insurance market. In a later extension, we allow firms to compete with positive markups up to a cap set by insurance regulations.

With this model, we conduct a set of counterfactual analyses to address our key regulatory design question: how would consumers, employers, and payers fare under alternative market segmentation? At a broad level, we highlight the role of adverse selection and insurer markups in determining the outcome of pooling. At a more specific level, we quantify the welfare effect of two recent trends in small group insurance coverage. First, we feature the federal rule change implemented in 2020 that allows employers to offer individual coverage HRAs or ICHRAs. Second, we can examine the effect of a shrinking small group insurance market; from 2011 to 2015, the share of small employers offering group coverage fell almost six percentage points (Corlette et al., 2017). The economic downturn resulting from the COVID-19 pandemic may have accelerated this trend (Dafny et al., 2020).

In our first counterfactual, we simulate market outcomes in a scenario in which the small group market shuts down. Small group employees with no access to alternative group or public insurance can either purchase insurance on the individual market or choose uninsurance. In this scenario, the former small group members lose access to the tax savings and the employer premium subsidies they enjoyed, but those with sufficiently low incomes gain access to government premium and cost-sharing subsidies. With this initial analysis, our goal is to characterize the effect of pooling when employers no longer play any role in insurance provision.

We find some changes to the extent of adverse selection in the individual market. Gold plan

premiums in the individual market, for example, fall from an average of $685 to $458 per month. For households in the individual market prior to the merged outcome, we estimate that average consumer surplus increases by $18 per month per household and the share of households who choose to be uninsured falls by 3 percentage points. The magnitude of this surplus change, however, would be larger but for the preference of healthy small group members to exit coverage. We predict the lowest spending households in the small group population choose to be uninsured when given the option of individual coverage.

Given the tendency to exit coverage, we consider a second equilibrium scenario in which small group households keep the same tax benefits and the same employer contribution (in lump-sum terms) that they earned with group coverage. We pool the markets by requiring all small group employers to offer ICHRA accounts, through which small group employees purchase individual coverage using their original subsidies. We find that pooling under this policy further mitigates adverse selection in the individual market. Average gold premiums fall from $685 per month to $354, while silver and bronze premiums fall by $11 and $5, respectively. As a result of the new equilibrium premiums, consumer surplus for the average individual market household rises by $28 per month.

Small group households, now able to choose freely from all plans on the individual marketplace or opt to be uninsured, see much larger increases in surplus, rising by $263 per month. We also find employer and government spending decreases. While part of this decline follows from the reduced premiums and thus reduced premium subsidies, the main driver of the equilibrium outcomes is again an increase in uninsurance. If households do not fully recognize the long-term costs of being uninsured, the policy could generate losses outside the framework of our model.

To examine the effect of pooling without the shift to uninsurance, we conduct a third counterfactual in which we both extend ICHRA and compel small group employees to enroll in coverage. In this setting, we find that pooling the markets generates greater consumer surplus gains among individual market households. Small group households' consumer surplus gains are lower because, based on revealed preference, these consumers prefer to be uninsured. We also find that employer expenditures fall in this simulation relative to the baseline with segmented markets. This decline

raises the possibility that wage increases could also compensate those small group enrollees who lose surplus due to the mandated coverage.

From the perspective of insurers and brokers, the gains to consumers, employers, and the government from pooling come at the expense of insurer margins. We examine whether our welfare conclusions might change if, in response to pooling in the individual market, insurers mark up premiums. We quantify equilibrium consumer surplus when those markups range from 0% to 25%, the approximate upper bound under federal regulation.[3] Up to a markup of roughly 20%, we find that individual market consumers still prefer pooling—the gains from pooling with healthier small group enrollees exceed the losses from new markups. Small group consumers also prefer pooling. On net, these households benefit from decoupling their premiums from other enrollees within an employer. Their surplus also increases because individual market plans involve lower administrative costs, such as broker fees.

Overall, our simulations point to specific market integration policies that can increase consumer welfare without additional taxpayer expenditure. When employers contribute subsidies toward small group households' insurance coverage in a pooled market, they not only provide a valuable benefit to their employees but also provide an external benefit to individual market households.

Previous literature. Our analysis relates to several research areas in the economics of health insurance markets. A growing body of work studies plan choice and optimal menu design for employer-sponsored insurance, including issues of selection on moral hazard, adverse selection, and risk preferences (Einav et al., 2013; Ho and Lee, 2020; Marone and Sabety, 2020). We build a model that features these same behavioral elements but in the context of a 'managed competition' market. Instead of an employer or planner setting premiums, insurers choose premiums as they compete in a regulated private market.

Given this focus, our work also connects to a literature identifying the consequences of market design in individual insurance, including the design of subsidies (Tebaldi, 2017; Jaffe and Shepard, 2020; Polyakova and Ryan, 2019), risk adjustment (Geruso et al., 2019; Einav et al., 2019), partic-

[3]We define markups for our analysis using total costs, including administrative costs.

ipation penalties (Diamond et al., 2021), and re-classification risk (Handel et al., 2015; Atal et al., 2020). Our contribution is to analyze the importance of market segmentation.

Finally, a smaller literature considers the characteristics of small group insurance. Closest to our paper is Fleitas et al. (2021), who also consider the impact of segmenting the small group and individual markets. They have detailed data on the small group market from a single large insurer but do not observe individual market enrollees. Thus, they focus on the effect of transferring the small group into a separate individual marketplace rather than pooling the two populations. Other papers consider plan pricing, the allocative efficiency of plan designs, and the extent of re-classification risk in the market (Abraham et al., 2019; Bundorf et al., 2012; Fleitas et al., 2018).

The rest of the paper proceeds as follows: in Sections 2 and 3, we describe our institutional setting and data. We contrast the market outcomes in small group and individual insurance in Section 4. Section 5 presents the model of supply and demand in the individual market and Section 6 describes how we take this model to data. In Section 7, we discuss our results and counterfactual analyses. Section 8 concludes.

2 Institutional Detail and Setting

Both the individual and small group markets faced new regulations following the implementation of the Affordable Care Act (ACA) in 2014. In this section, we summarize the key features of these markets in Oregon that are relevant for our analyses. We provide further details in Appendix A.

2.1 Individual insurance market

Households can purchase insurance coverage in the individual market through two channels. First, they can search and select a health plan through a marketplace created under the ACA. Enrollees in Oregon use the federal healthcare.gov online platform, in an arrangement known as a "state-based exchange on the federal platform". Through this portal, Oregon residents can shop for health plans whose prices and cost-sharing levels may be subsidized to reflect their financial circumstances. Households with incomes between 100% and 400% of the federal poverty level (FPL) see plans with

reduced premiums, and those between 100% and 250% also see reduced out-of-pocket costs for some plans, reflecting a schedule of government subsidies.[4]

Households that are eligible for subsidies must purchase through the marketplace to receive them. Unsubsidized households can either use the marketplace channel or a second "off-marketplace" channel, employing an insurance broker or agent to purchase individual coverage.[5] We define the choice set of off-marketplace purchasers differently from on-marketplace enrollees by using the observed menus in our data.

Plans offered in the individual market face regulation of both premiums and the level of coverage. Since the implementation of the ACA in 2014, insurers in all states must 'guarantee issue' all new plans to all consumers—that is, the insurer cannot reject applicants based on health status or pre-existing health conditions. Premiums may vary only with family size, state-defined geographic regions, tobacco use, and age. Age-based premiums follow a standard age curve in Oregon, with a ratio of 3 to 1 from the oldest to youngest enrollee. We exploit the formulaic variation in premiums by age in later analyses.

The ACA limits plan differentiation in two ways. First, all plans must cover a set of ten essential health benefit categories, including outpatient services, emergency room visits, pregnancy and maternity visits, mental health care, and prescription drugs. And second, each plan's coverage design must fit into one of four metal tiers classified by actuarial value, defined as the percentage of health costs the plan is expected to cover. The plan tiers include bronze, silver, gold, and platinum, with actuarial values of 60%, 70%, 80% and 90%, respectively. Oregon requires all insurers entering the marketplace to offer a bronze, silver, and gold plan; carriers offering plans outside the marketplace must offer at least one bronze and one silver plan.[6] Partly as a result of these requirements, consumers in the individual market often choose plans from a large menu. For

[4]Households with incomes at or below 133% of FPL, as well as those meeting several other criteria, are eligible for Medicaid, the joint state and federal means-tested government insurance program.

[5]Some off-marketplace plans are "grandfathered", meaning they were initially purchased prior to March 23, 2010 and renewed in future years; grandfathered plans need not adhere to the benefit design requirements of the ACA. When we model demand for individual insurance, we omit those households purchasing grandfathered plans but include households purchasing ACA-compliant plans through brokers.

[6]Oregon's menu requirement is more stringent than federal rules, which mandate that insurers sell at least one gold and one silver plan in the marketplace in each geographic market they enter. Oregon also requires insurers to offer a "standardized" plan in each metal tier (Blumberg et al., 2013).

example, a buyer in the Portland area in 2015 has a choice of 31 bronze plans, 40 silver plans, and 24 gold plans, offered by 8 unique carriers.

Households that are eligible for cost-sharing subsidies must purchase a silver plan in order to receive these reduced out-of-pocket costs. These subsidies change the standard silver plan design to a more generous actuarial value of between 73% and 94% for consumers with incomes of 100% to 250% of the FPL, with the lowest incomes receiving the higher actuarial values. As noted below, this subsidy structure has meaningful implications for household enrollment decisions, and hence the costs and premiums of plans in different tiers.

2.2 Small group insurance market

Small employers, defined in Oregon as firms with up to 50 full-time employees, have the option of offering health insurance coverage for their employees. Nationally in 2015, approximately 47% of firms with 3-9 workers and 68% of firms with 10-49 workers offered coverage to employees; approximately 76% of eligible workers took up this coverage (Claxton et al., 2015). After the implementation of the ACA, the fully-insured small group market faces many of the same plan design restrictions as the individual market. Plans must cover the same essential health benefits, must be structured according to the same metal tiers, and must be 'guaranteed issue'. The small group and individual markets differ in the purchasing channel and the pricing rules. We discuss each feature in turn.

While the ACA intended states to offer a marketplace for small group employers to shop for plans— known as the Small Business Health Options program or SHOP— Oregon did not have a small group marketplace during the span of our data. Instead, small employers purchase plans through an insurance broker who typically receives a fee per enrollee from the insurance carrier.[7] In the broker-mediated market, employers choose plans to offer their employees, typically providing a much smaller choice set relative to the individual market. A typical small group offers one to

[7]Brokers typically receive a per-month per-enrollee commission plus occasional sign-on bonuses. One carrier in Oregon, for example, offered a $14.27 per-enrollee per-month payment for groups with fewer than 26 enrollees and $11.25 for plans with 26-49 enrollees. Bonuses equaled $100 for a 1-9 enrollee group, $200 for a 10-25 enrollee group, and $400 for 26-49 enrollee group (Providence Health Plan, 2011). The average small group commission in Oregon in 2016 was $19.70/enrollee per month (The Kaiser Family Foundation, 2020).

two broker-recommended plans to its employees, often from the same carrier (AHRQ, 2016). The presence of intermediaries may also shield insurance carriers from premium competition; we show later that the markups of plans offered by small employers in our sample are often meaningfully higher than those available for comparable plans in the individual market.

After choosing a menu of plans to offer, employers contract with the relevant carrier(s) and pay premiums on behalf of the group. Employees pay their share of the group premium from their pre-tax earnings—that is, all premiums for insurance obtained through an employer are exempt from federal and state taxes, regardless of whether the employer or the employee pays. This creates a tax wedge relative to the individual market, where households typically pay for insurance with post-tax dollars.[8] There is also premium variation relative to the individual market because small group market insurers are required to use 'tiered composite' community rating, described in Appendix A. The composite rating system creates a cross-subsidy within the employer pool between older and younger enrollees and between employees covering only themselves and those covering families. Finally, employers typically subsidize the cost of employee premiums, covering as much as two thirds of the premium cost (Claxton et al., 2015).

3 Data

We collect data from three sources. First, to analyze both plan enrollment and health care costs, we use claims data from the Oregon Health Authority's All Payer All Claims (APAC) dataset. Second, we use the National Association of Insurance Commissioners' SERFF database for information on plan design and premium levels. Third, we collect Medical Loss Ratio (MLR) reports from the Centers for Medicare and Medicaid Services for the insurance carriers operating in the small group and individual insurance markets in Oregon. We describe each data source and its use below.

Claims and enrollment data. From Oregon's APAC data, we collect private insurance claims and enrollment information for all small group and individual insurance plans purchased from 2010-2016. Our claims data cover out-of-pocket costs and costs to the insurer for inpatient, outpatient,

[8] Premiums in the individual market are part of itemized deductions, but subject to limitations: only medical expenses exceeding 7.5% of adjusted gross income are deductible.

and pharmaceutical claims for each covered enrollee, including spouses and dependents. The data also record the household's insurance plan choice as well as the plan's characteristics, including its carrier, metal tier, coverage period, and premiums. Focusing on the period around the introduction of the Affordable Care Act, in years 2014 to 2016, the data include 354,366 unique households in the individual market and 218,827 small group households.[9]

We use the claims data, together with the Johns Hopkins ACG (Adjusted Clinical Groups) Case-Mix System software, to construct a measure of predicted health spending for each household in each year. The ACG software predicts expected medical expenditure of each enrollee based on diagnostics and demographics data.[10] The results are then normalized to an ACG score, where a score of one corresponds to average expenditure in a reference population.

Finally, we derive a measure of household income. We do not observe income directly but instead approximate it using our data on both household demographics and the net premiums consumers pay for insurance. Specifically, for some observations in our data, we observe net marketplace insurance premiums, reflecting government subsidies that scale with household income. We use this sample to estimate a predictive model of income as a function of demographics. We apply our model estimates to predict income for the entire sample. Using this generated income measure, we derive premium subsidies and tax rates for all households in the individual and small group markets. Appendix B details our procedure.

Table 1 summarizes the demographic information for the full sample of households included in our analyses. On average, households in the small group market are younger, healthier (as measured by household ACG score), more urban, and more likely to purchase single coverage than households in the individual market.[11]

[9] We exclude households who choose plans from insurers with very low market shares or who are missing key demographic characteristics. Appendix B details our data construction procedure.

[10] We use the ACG software's "concurrent" risk measure. This measure uses diagnostics and demographics to predict expected medical expenditure in the same year. An alternative measure that uses lagged diagnostics and demographics limits the sample size in our setting due to high rates of consumer churn in the individual market.

[11] While over 70% of enrollees in both samples purchase insurance for a single household member, some of these enrollees likely have a spouse or other family member who accesses health insurance through a distinct employer or public program.

Plan characteristics. Each year, insurers operating in Oregon's individual and small group markets file details of their plan offerings via the National Association of Insurance Commissioners' SERFF database. These filings allow us to validate the characteristics of each plan chosen in Oregon's APAC data, including the levels of deductibles, copayments, and the gross premium levels set for each plan. To simplify the model of household plan choice, we bin similar plans together to create constructed plans. In the model, households choose among constructed plans, defined by a combination of insurer, rating area, metal tier, and plan type.[12] Hereafter, we refer to constructed plans simply as plans.

Table 2 summarizes plan characteristics. We observe important differences in the plan choices between the small group and individual markets. Only 14% of individual market households enroll in the more generous gold plans compared to 36% of small group households. For the least generous bronze plans, the shares are nearly reversed: small group households choose bronze plans 14% of the time, while bronze plans capture a 30% share of the individual market.

Administrative costs. Lastly, we collect Medical Loss Ratio (MLR) reports from the Centers for Medicare and Medicaid Services for the insurance carriers operating in the small group and individual insurance markets in Oregon. The reports contain state-wide enrollment and costs, including both clinical costs and administrative costs, that insurers incur in each business line in which they operate. We exploit these observed measures of administrative cost in later analyses of insurer pricing. We also use the national enrollment and financial information for each carrier that operates in Oregon to create measures of insurance participation, enrollment, revenues, and costs in states outside Oregon. We use these variables as instruments in the premium-setting regressions described in Section 6.3.

4 Descriptive Analyses

Before describing our model, we highlight the differences between the small group and individual markets in Oregon in terms of health spending, premiums, and the level of insurer markups. The

[12] We divide plan types into 'managed care' and 'not managed care'. Managed care plans include EPO and HMO plans. We define the premium of a constructed plan as the average premium of all plans grouped into the bin.

empirical state of the markets in Oregon will underlie our model estimates and help illustrate the likely consequences of pooling the two markets. We also compare the experience in Oregon to other U.S. states.

4.1 Oregon comparison

Comparing costs. The average small group household spends less on health care than an individual market household. In Table 2, we report that small group households are two percentage points more likely to have no monthly health spending (32% vs. 30%); when consuming health care, they spend an average of $524 per month versus $591 per month among individual market households. We illustrate the full distribution of health spending by year and market type in Figure 1. To control for moral hazard effects, we break out costs separately for plans of the same actuarial value in Panels A through C.

Within the silver and gold metal tiers, which compose roughly 70% of plans purchased in the individual market and 86% of plans purchased in the small group market, we again observe lower costs among small group enrollees. In particular, small group households are nearly twice as likely incur $0 of health spending in gold plans (25% versus 13%) and 42% more likely to spend $0 in silver plans (34% vs. 24%). Conditional on positive spending, the distribution of costs among small group households appears more concentrated at lower levels of spending. For bronze plans, the fraction of zero-cost consumers are similar in the two segments, while the distribution of individual market enrollees shows slightly lower costs. These findings are consistent with the reported mean and median spending levels by metal tier in Table 2.

Comparing premiums. Given the observed differences in spending, we examine how premiums differ between the small group and individual markets. To illustrate the differential, we focus on a standardized enrollee who is single, 40 years old, and too wealthy to receive government premium or cost-sharing subsidies in the individual market. Figure 2 provides an initial comparison of the premiums carriers set across the two markets. It depicts the distribution of premiums for the standardized enrollee across plans offered by each carrier in a year, metal tier, and market segment. We compare the distribution of premiums separately by metal tier in Panels A through C.

Base premiums are higher for small group plans in 2014-2015 in all tiers, despite small group enrollees' lower medical claims costs in the silver and gold plans. We interpret this pattern as consistent with more intense price competition or lower administrative costs in the individual market. Individual plan premiums increase in 2016, leading the difference in premiums between the two markets to be statistically insignificant in that year.

The base premiums, however, omit several adjustments needed to capture the true small group premiums that households face. A more precise comparison between markets would also account for: (a) pooling within an employer (b) the tax savings for employer-provided insurance; and, (c) employer contributions towards premiums. Figure 3 Panel A illustrates the effect of these adjustments. We describe our adjustment approach in Appendix G.[13]

We find that when we add adjustments for the typical group composition and the implicit tax subsidy in employer-sponsored insurance, small group premiums continue to exceed those in the individual market. Only when we add employer subsidies do we find small group plan premiums broadly fall below gross premiums in the individual market. Panel B of Figure 3 illustrates that when the employer subsidy rate is on the order of 15% in 2014 and 2015, the typical employee would be indifferent between a small group versus individual market plan.[14]

Insurer markups over health costs. Given our measurement of both costs and premiums, we can also compute a insurer markup of premiums over health costs. We call this ratio the medical markup. For this analysis in the small group market, we use the full premium paid to small group plans independent of employee tax savings or employer subsidies, since the full premium flows to the insurer as revenue.

In Figure 4 we plot the distribution of medical markups by year in the small group and individual market. In Panels A to C, we distinguish the markups by metal tier. In all tiers, we observe larger

[13] In brief, we account for the tiered composite rating in Oregon by drawing simulated groups from the distribution of small groups in Oregon, replacing one member of the group with our standardized enrollee. We take an average of the single employee premium across all simulated groups. To adjust for taxes for our standardized enrollee, we multiply the premium by $1-\tau$, where τ is the average tax rate for a single adult making the median annual income in Oregon. Finally, we assume an employer subsidy rate of 50%.

[14] This employer subsidy rate, however, assumes that consumers would not be eligible for federal premium subsidies in the individual market. For consumers eligible for government subsidies, the employer subsidy rate would have to increase to make the premiums comparable.

markups in the small group market, particularly after 2014. The differences are largest in Panels B and C for the more generous plans. For the most popular silver plans, which compose the majority of plan purchases, the median small group insurer in 2014 had a roughly 40% markup as a share of premiums. This median markup fell over time, but remained above zero in all years of our data. Unsurprisingly, in the individual market where premiums are lower and average health spending is higher, we see much lower markups as a share of premiums. The median firm had negative markups in all three years, while even the 75th percentile firm's markups were negative in 2014.

The markups we illustrate in Figure 4 do not include non-medical costs, including the costs of broker fees and plan administration. They also do not include risk adjustment payments and other subsidies carriers receive. In the supply model we develop later, we measure markups accounting for these additional costs and subsidies.

4.2 Broader comparison

While we focus our analysis in this paper on the insurance markets in Oregon, we investigate whether the same patterns—higher small group market premiums and markups—appear outside of Oregon. We repeat our descriptive analysis in premium data compiled by the Robert Wood Johnson Foundation for 33 states in 2014 and all 50 states and the District of Columbia in 2015-2016. We describe our procedure in Appendix G and report results in Table A6. In these years and states, we find, on average, a 10-15% higher base premium in the small group market than in the individual market for plans with the same actuarial value. In short, Oregon's market outcomes appear similar to outcomes in other states.[15]

In addition to comparing premiums, we also measure medical markups outside Oregon. We use the data carriers report in each state as part of MLR regulation to examine these markups.[16] We

[15] In 2017-2018, we observe a distinct pattern of premiums. Reflecting an increase in premiums in the individual market in many states, we observe individual market premiums that exceed the small group base premiums by approximately 10%. The premium gap we observe between segments in 2018 are similar to findings from Abraham et al. (2019). The marketplace premiums in these years reflect, in part, an anticipated policy change in October 2017 in which the federal government stopped reimbursing insurers for the added expense of cost-sharing reductions granted by statute to certain low-income enrollees. In addition, the temporary risk corridors program, administered under Section 1342 of the ACA, ended after the 2016 plan year.

[16] We do not compare medical loss ratios directly because these measures include adjustments for quality improvement programs and other expenses we do not observe in our main dataset.

find that medical markups are higher in the small group market relative to the individual market in most states in the years 2014 through 2016: percentage markups in the individual market typically range between a loss of 10% and a gain of 10%, while small group markups reflect a gain of 20-30%. We illustrate the average markups by sector, state, and year in Appendix Figure A1.[17]

4.3 Implications for market pooling

Our comparison of costs, premiums, and markups above suggests multiple possible outcomes from eliminating market segmentation. First, we observe that small group enrollees both spend less than individual market enrollees on average and are partially shielded from plan premiums by the implicit tax subsidy they receive and by their employer's subsidy.[18] Thus we can view market pooling as introducing a population of relatively low-cost consumers with a low exposure to premiums into the individual marketplace. If the high market share of gold plans in the small group market reflects a preference for coverage, many of these new customers may continue to choose gold plans, reducing the existing plans' average costs and premiums and potentially spurring more individual enrollees to opt for a gold plan. The costs of enrollees in silver and bronze plans, in turn, may fall as relatively sick enrollees move up to the more generous metal tiers. Under average cost pricing, these plans' premiums will fall. Interestingly, there may be little cost of this adjustment to small group enrollees, since the benefit of entering a more competitive marketplace, with lower plan markups, may more than offset the loss from being pooled with a higher-spending consumer group.

There is, however, a different possible outcome of market pooling that is less beneficial to consumers. If the popularity of gold plans in the small group market is due to employers providing generous coverage—perhaps as a recruitment device—rather than a true reflection of employee risk preferences and health needs, then small group members moving to the individual market may choose lower-coverage plans when given the option. In that case, the market share of gold plans may remain low and gold plan premiums high. Further, in the absence of employers prompting

[17] We can also compare markups in 2017, beyond the years of our main data. In that year, markups in the individual market increase, on average, relative to prior years. However, medical markups in the small group market continue to exceed those in the individual market in most states, as shown in Appendix Figure A1.

[18] While the tax treatment is a true benefit to employees, the employer contribution is not: it is likely to be passed through in reduced wages. We return to this point in our discussion of counterfactuals below.

employees to enroll, low-spending small group employees may choose to forego insurance on the individual market entirely. In such an outcome, only sick small group members would choose to enroll in the individual market, potentially raising the average costs in their chosen plans.

Distinguishing between these possible outcomes from pooling requires a model that captures the preferences and underlying characteristics of households in both the individual and the small group markets. While the differences in underlying health care needs will surely affect market outcomes under pooling, so will the extent of moral hazard (i.e., price responsiveness at the point of care) and risk aversion in the two populations; these factors can also affect plan choices and subsequent medical spending, costs, and equilibrium premiums. In the following sections we outline the model we use to estimate these objects and also describe our approach to estimation.

5 Model

We build a demand model that predicts households' plan choices and subsequent medical care utilization. We use the multiplicative moral hazard model from Einav et al. (2013); this model captures both adverse selection, where enrollees select into plans based on unobserved health status, as well as moral hazard, a spending response to insurance. We model insurer pricing in the individual market using a supply framework that features insurance plan competition, as in Azevedo and Gottlieb (2017).

5.1 Consumer demand and spending

At the beginning of each year, a household engages in a two-stage sequential choice. In Stage 1, the household—internalizing the needs and preferences of its members—chooses an insurance plan from a set of offered plans. In Stage 2, conditional on plan choice and the realization of its health care status, the household chooses the amount of medical care to consume. Following the notation and structure of Einav et al. (2013), we characterize a household by three objects, $(F_{\lambda,t}(\cdot), \omega, \psi)$, where, for clarity, we omit household-specific subscripts. $F_{\lambda,t}(\cdot)$ represents the household's expectation over its uncertain health status $\lambda \geq 0$ in period t. The household realizes the value of λ in a given period after it chooses a plan; a higher value of λ corresponds to a household with greater health

care needs. The second object, ω, represents the household's price sensitivity for medical care and can be interpreted as the household's level of moral hazard. Lastly, ψ represents the household's coefficient of absolute risk aversion.

We present this model in reverse order, beginning with Stage 2. We use a multiplicative moral hazard specification because it predicts higher moral hazard spending for sicker individuals. This assumption both seems natural in our setting and matches results from Ho and Lee (2020), who find a similar assumption fits the variation in their employer sample.

5.1.1 Stage 2: medical care utilization

At the beginning of Stage 2 in each year t, a household enrolls in an insurance plan j and realizes health λ. The household then chooses its level of medical spending $m \geq 0$ for that period to maximize its utility given by:

$$u_{j,t}(m;\lambda,\omega) = (m-\lambda) - \frac{1}{2\omega\lambda}(m-\lambda)^2 + [y_t - c_{j,t}^{OOP}(m) - p_{j,t}] + g(X_{j,t},\epsilon_{j,t}) \ . \tag{5.1}$$

In (5.1), y_t represents annual income, $c_{j,t}^{OOP}(m)$ are the out-of-pocket (OOP) costs the household pays for its medical care, and $p_{j,t}$ is the annual plan premium. We specify $c_{j,t}^{OOP}(m)$ as $(1-x_{j,t})\times m$, where $x_{j,t}$ is the percentage of spending that the insurer pays in period t under plan j—i.e., the plan's actuarial value. The final term, $g(.)$, is a function of other variables that can affect household utility.

The household's medical spending, denoted $m_{j,t}^*$, must satisfy the following first-order condition from (5.1):

$$m_{j,t}^* = \lambda + \omega\lambda x_{j,t} \ . \tag{5.2}$$

The second term in this expression implies that the amount of additional medical spending due to cost-sharing is increasing in the moral hazard parameter ω and also in household "sickness" λ. A zero value of λ will result in zero spending, even under full insurance.

5.1.2 Stage 1: insurance plan choice

In Stage 1, each household realizes its $\epsilon_{j,t}$ and chooses an insurance plan from a choice set of plans $\mathcal{J}_t$ to maximize its expected utility for the current year. The household anticipates that its health needs follow $F_{\lambda,t}$ and its health spending will be governed by optimal Stage-2 behavior. We assume the household has constant absolute risk aversion (CARA) preferences over Stage-2 utilities given optimal medical spending, denoted $u^*_{j,t}(\lambda,\omega) \equiv u_{j,t}(m^*_{j,t}(\lambda);\lambda,\omega)$. Given these assumptions, the expected utility that a household anticipates receiving from plan j at the beginning of period t is given by

$$v_{j,t}(F_{\lambda,t},\omega,\psi) = -\int exp(-\psi \times u^*_{j,t}(\lambda,\omega))dF_{\lambda,t}(\lambda)\ , \tag{5.3}$$

where ψ is the household's coefficient of absolute risk aversion and the household's optimal choice of plan is $j^* = \arg\max_{j\in\mathcal{J}_t} v_{j,t}(\cdot)$.

5.2 Insurance supply

To model insurer pricing and equilibrium supply in the individual market, we start by following the assumption of perfect competition in Azevedo and Gottlieb (2017), who adopt the Einav et al. (2013) model with additive moral hazard as an example. We assume insurance carriers in the individual market choose their plan-specific premiums to equal average costs, including both claims and administrative costs. We describe our specification of administrative costs in Section 6. In Section 7.4, we describe an extension to our main analysis in which we allow positive plan markups.

We can write down expected claims costs using our model. Given optimal medical spending in equation (5.2), a household characterized by (λ,ω) has the following expected claims costs:

$$c_{j,t} = \int (x_{j,t}\lambda + \omega x^2_{j,t}\lambda)dF_{\lambda,t}(\lambda).$$

5.3 Consumer surplus

To quantify consumer surplus, we define the certainty equivalent of plan j as e_j such that $-exp(-\psi e_j) = v_{j,t}(F_{\lambda,t},\omega,\psi)$, as in Einav et al. (2013). Integrating over the distribution of λ, we can show that

for individuals of type i:

$$e_{i,j}(F_{\lambda,t}, \omega, \psi) = -\frac{1}{\psi_i} log \left[\int exp(-\psi \times u^*_{j,t}(\lambda, \omega)) dF_{\lambda,t}(\lambda) \right] \tag{5.4}$$

so that, for households of type i, we can write ex ante consumer surplus—that is, consumer surplus before a household realizes its ϵ shock in the plan choice stage—in the following way:

$$CS_i = E_\epsilon \left(\max_j e_{i,j} \right) = log \sum_{j=1}^{J} exp \left(e_{i,j}(F_{\lambda,t}, \omega, \psi) \right) \tag{5.5}$$

Here, each household chooses from J plans in the individual market. Later, we describe an alternative approach to calculating consumer surplus in the small group market. We use this alternative measure when comparing welfare for consumers forced to switch from small group to individual coverage.

6 Empirical Model

We transform the model of insurance demand and insurer supply into a two stage empirical model for estimation. We begin by detailing the data samples used for estimation and then describe our estimation approach.

6.1 Data samples

To estimate preferences of households in the individual market, we construct a dataset of all households who purchase individual market insurance and all uninsured households in the state. We directly observe households who purchase individual market insurance in our all-payer claims data. The annual American Community Survey, run by the US Census Bureau, provides the size of the uninsured population for each rating area, year, and age. We infer additional characteristics of the uninsured population using a detailed survey of uninsured households in California, the California Health Interview Survey. Appendix B.4 details this procedure. Table A4 compares the uninsured

population to the insured individual market population. On average, insured households are older and have larger incomes.

Our counterfactual simulations also require reasonable estimates of the preferences and attributes of current small group households. We do not estimate a full model of plan and utilization choices in the small group market; in our setting, estimating such a model would require strong assumptions about the role of employers and brokers as agents. Instead, we consider the sample of small group enrollees whose employers canceled group coverage during the time period of our data. We can track these households as they choose plans in the individual market or become eligible for coverage in other markets, including large group insurance or public insurance. Among this population of "forced switchers", we focus on those households who switch to individual coverage or uninsurance. We define the uninsured share of forced switchers as those households formerly enrolled in small group insurance who do not purchase a group plan, individual plan, and do not fall into eligibility categories for public insurance. We also exclude individuals who switch to a spouse's insurance. Appendix B.2 describes how we identify the forced switcher population.

We use our forced switcher sample as a plausible proxy, conditional on observables, for the preferences of the larger small group enrollee population. In Appendix Table A3, we demonstrate balance between the demographics of the forced switchers and the demographics of the entire small group in the year before the forced switchers leave the small group market. While the sample of forced switchers appears largely similar to the small group market population overall, they differ on a few dimensions. On average, forced switchers are slightly older and sicker according to health status scores. We therefore condition on observables when we use the model to estimate switcher preferences.

6.2 Joint likelihood of plan choice and health spending

We next define the likelihood of observing (a) the plan choices of households and (b) the health care spending of subscribers to each plan j, accounting for both moral hazard and risk aversion.

6.2.1 Equations for estimation

We begin by making a distributional assumption on the household's health state, λ. We model λ as following an exponential distribution to approximate the empirical spending pattern depicted in Figure 1. If $\lambda \sim exponential(\alpha)$, plugging optimal medical spending into the utility equation (6.1) and accounting explicitly for variation in the underlying parameters and characteristics of household i yields:

$$u^*_{i,j,m,t} = \frac{1}{2}x^2_{i,j,t}\omega_i\lambda_i - (1 - x_{i,j,t})\lambda_i + y_{i,t} - p_{i,j,m,t} + g(X_{j,m,t}, \epsilon_{i,j,m,t})$$

where m indexes geographic markets and t indexes years. Suppressing (m, t) subscripts for notational simplicity, the expected utility over the distribution of λ is:

$$v_{i,j}(F_{\lambda,i}, \omega_i, \psi_i) = -\int exp(-\psi_i u^*_{i,j}) dF_{\lambda,i}(\lambda).$$

Noting that our distributional assumption implies $E(\lambda) = 1/\alpha$ and applying the order-preserving monotonic transformation $-\frac{1}{\psi_i}ln(-v_{i,j})$, we write the expected utility as [19]:

$$U_{i,j} \approx -p_{i,j} + \frac{x_{i,j}}{\alpha_i - \psi_i} + \frac{x^2_{i,j}\omega_i}{2(\alpha_i - \psi_i)} + g(X_j, \epsilon_{i,j})$$

$$U_{i,0} = g_0(\epsilon_{i,0})$$

We specify $g(X_j, \epsilon_{i,j}) = (\beta_0 X_j + \epsilon_{i,j})/(\alpha_i - \psi_i)$ so that sicker or more risk averse households put more weight on plan characteristics like carrier identity, in the same way that they put more weight on coverage. Making an analogous assumption for the outside option, we find:

$$U_{i,j} = -p_{i,j} + \frac{x_{i,j}}{\alpha_i - \psi_i} + \frac{x^2_{i,j}\omega_i}{2(\alpha_i - \psi_i)} + \frac{\beta_0 X_j + \epsilon_{i,j}}{\alpha_i - \psi_i} \tag{6.1}$$

$$U_{i,0} = \frac{\epsilon_{i,0}}{\alpha_i - \psi_i}$$

[19] This transformation also requires us to recognize that when $Ax + Bx^2$ is close to zero, we can approximate $ln(1 + Ax + Bx^2) \approx Ax + Bx^2$. We provide details in Appendix C.

In this expression, we observe four components to utility. The first three components derive from the financial terms of insurance plans: disutility from premiums; utility from covered non-discretionary spending; and utility from spending due to moral hazard. The final component allows non-financial characteristics, like carrier name, to affect utility. In addition, risk coverage also generates utility: both spending-related terms are adjusted upwards because we divide by $\alpha_i - \psi_i$ in place of α_i.

Finally, the expected cost to the insurer of enrolling a household of type i is:

$$c_{i,j} = \int (x_{i,j}\lambda_i + \omega_i x_{i,j}^2 \lambda_i) dF_{\lambda,i}(\lambda)$$

$$c_{i,j} = \frac{x_{i,j}}{\alpha_i} + \frac{\omega_i x_{i,j}^2}{\alpha_i}$$

which approaches zero when α_i is large. We assume the insurer is risk-neutral. As in the model with additive moral hazard discussed in Azevedo and Gottlieb (2017), the insurer pays the full cost of spending due to moral hazard while consumer utility reflects only half of that spending, adjusted upwards due to risk coverage.

6.2.2 Plan choice and health spending

Plan choice. Taking equation (6.1) as a starting point and multiplying through by $\alpha_i - \psi_i > 0$, we have:

$$u_{i,j} = x_{i,j} + \frac{1}{2}\omega_i x_{i,j}^2 - (\alpha_i - \psi_i)p_{i,j} + \beta_0 X_j + \epsilon_{i,j}. \tag{6.2}$$

$$u_{i,1} = \epsilon_{i,1}$$

where i denotes households and j denotes health plans. We suppress the indices (m, t) denoting geographic markets and time periods to simplify exposition. We assume $\epsilon_{i,j}$ follows a Gumbel or type I extreme value distribution. The probability that an enrollee i chooses plan j then takes the standard logit form:

$$s_{i,j} = Pr(d_{i,j} = 1) = \frac{exp(V_{i,j})}{\sum_{k=1}^{J} exp(V_{i,k})} \tag{6.3}$$

where $V_{i,j} = x_j + \frac{1}{2}\omega_i x_j^2 - (\alpha_i - \psi_i)p_j + \beta_0 X_j$, and $d_{i,j} = 1$ when household i chooses plan j. Here, plan choice $j = 1$ represents the outside good of no insurance. Thus, $V_{i,1} = 0$ since both the actuarial value and premium equal zero when the 'plan' represents a lack of insurance. Other components of the utility equation are normalized relative to the outside option.

Health spending. Based on our model above with multiplicative moral hazard and with our particular parameterization, we can write the insurer's expected responsibility for health spending using the following form:

$$c_{i,j} = (x_{i,j} + \omega_i x_{i,j}^2)\lambda_i$$

where again we have omitted subscripts (m, t).

We further assume that all enrollees have some positive health care spending and that the insurer bears some cost of enrolling even healthy consumers. However, we define a spending cutoff $\underline{c}$ such that, for $0 \leq c_i \leq \underline{c}$, the "hassle costs" of submitting claims may lead the enrollee not to submit one. The insurer similarly saves the processing cost of the claim. Thus, we interpret zero spending observations as implying small but positive health care spending by the enrollee, with an associated small cost to the insurer. We therefore employ a fixed cutoff, $\underline{c}$, and treat all observed costs before that threshold as censored.[20]

Given our assumption that λ_i follows an exponential distribution with parameter α_i, we can write:

$$f(c_{i,j}|x_{i,j}, \omega_i, \alpha_i) = \begin{cases} 1 & x_{i,j} = 0, c_{i,j} = 0 \\ 0 & x_{i,j} = 0, c_{i,j} \neq 0 \\ 1 - exp\left(-\alpha_i \left(\frac{\underline{c}}{x_{i,j}+\omega_i x_{i,j}^2}\right)\right) & x_{i,j} \neq 0, c_{i,j} \leq \underline{c} \\ \frac{\alpha_i}{x_{i,j}+\omega_i x_{i,j}^2} exp\left(-c_{i,j}\frac{\alpha_i}{x_{i,j}+\omega_i x_{i,j}^2}\right) & x_{i,j} \neq 0, c_{i,j} > \underline{c} \end{cases}$$

[20] In robustness analyses, we vary the cutoff to test the effect on our coefficients of interest. We also design a likelihood routine to recover the threshold. See Appendix E.

The joint likelihood of the household's plan choice and its health spending is:

$$\mathcal{L}(\theta) = \prod_{i=1}^{N}\prod_{j=1}^{J} f(d_{i,j}, c_{i,j}|\cdot,\theta)^{d_{i,j}} = \prod_{i=1}^{N}\prod_{j=1}^{J} P(d_{i,j}=1|\cdot,\theta)^{d_{i,j}} f(c_{i,j}|\cdot,\theta)^{d_{i,j}} \tag{6.4}$$

where $j = 1$ is the outside option, $P(d_{i,j} = 1|\cdot, \theta)$ is the probability that patient i chooses plan j, and $f(c_{i,j}|\cdot, \theta)$ is the likelihood of patient i in plan j having cost $c_{i,j}$. We derive the log-likelihood used for estimation in Appendix C.

We further parameterize $(\alpha_i, \omega_i, \psi_i)$ as functions of household observables. α_i contains observables relevant for nondiscretionary spending: the sum of the severity types in the household, as measured by ACG score; an indicator for a family member having a top-quartile ACG score; and indicators for whether the household head is older than 50 and whether the household covers dependents under its plan. We assume ω_i and ψ_i are constant across households in our sample.[21] Because enrollees eligible for cost-sharing subsidies have incentives to select silver-tier plans, we account for this possible steering effect in plan choice by including in X_j an interaction between an indicator for silver plans and an indicator that equals 1 if the household's purchase is subsidized through cost-sharing subsidies.[22]

The assumption of perfect competition, if taken literally, implies no role for an unobserved quality term in the utility equation. However, we recognize that this is an abstraction: plans in the individual market may in fact be differentiated in ways that consumers value, and premiums may (at least somewhat) respond to this unobserved heterogeneity. Following prior literature (Polyakova and Ryan, 2019; Tebaldi, 2017; Tebaldi et al., 2019), we address this issue by using the fact that insurers in the individual market are not permitted to vary premiums freely across consumers. Within each rating area, premiums for a given plan vary only by age, family size and (through subsidies) across income levels. Further, this variation is based on a pre-specified statutory formula

[21] Setting a constant ω_i, however, does not imply that all medical care is subject to the same moral hazard effect. In our specification of moral hazard as multiplicative, when a household has a higher severity score, λ, its moral hazard will be larger. Thus, households likely to have more inpatient care, for example, will have a higher moral hazard component in their expected spending.

[22] We adjust the household's premium and its coverage $x_{i,j}$ to account for subsidies. Allowing the cost-sharing subsidy eligibility to affect a household's preference for silver plans, in an additively separable way, accounts for any additional steering to silver plans from the subsidy structure that is not captured by premium and out-of-pocket cost reductions.

that does not vary by carrier or plan. That is, the institutional features of this market permit us to address premium endogeneity concerns by including carrier fixed effects to control for unobserved quality differences across carriers in each year and rating area. Remaining variation within tier stems from variation in the age composition and size of each observed household.[23]

6.3 Premium setting

We calculate the total premium revenue collected by the insurer for plan j, denoted R_j, as the share-weighted average premium charged to enrollees across markets the plan serves. We first label plan j's standardized premium, for a single 40 year old enrollee (suppressing m and t subscripts as before), as p_j; this will be useful for our counterfactual simulations below. To find the premium household i pays, we multiply this standardized premium by the rating factors $\gamma_{k,i}$ assigned to all individuals k covered in household i. Under the ACA, Oregon fixes the value of these weights according to a published age curve; we apply these weights in our calculation. We sum these weighted premiums to the plan level, multiplying the premium each household faces by the predicted probability it chooses plan j, $\hat{s}_{i,j}$. In our notation, revenue then equals:

$$R_j = \sum_{i \in \{t\}}^{N_t} \hat{s}_{i,j} * p_{i,j} = \sum_{i \in \{t\}}^{N_t} p_j * \hat{s}_{i,j} \sum_{k \in i}^{K_i} \gamma_{k,i} \tag{6.5}$$

Then, under the assumption of perfect competition, we can estimate the costs of each plan by finding the vector of parameters that equate the insurer's revenue with its total health and administrative costs. These costs exclude any claims costs addressed by risk adjustment. We assume the insurer sets premiums to break even on each plan it offers, summing expected costs across individuals and geographic markets in a given time period.[24] Thus total premium revenues

[23] Given that our utility model already controls directly for plan tier actuarial values, we did not also add plan tier fixed effects. As a robustness check, we run an additional specification in which we control for carrier, year, and rating area fixed effects. The qualitative results remain unchanged.

[24] Our assumption that carriers break even plan by plan is consistent with the approach of Azevedo and Gottlieb (2017).

for plan j equal its total costs:

$$R_j = \sum_{i \in \{t\} \forall m}^{N} \left(\hat{s}_{i,j} * \beta_{4,j} \kappa_{i,j} c^e_{i,j} \right) + \beta_5 * A_j + \eta_j \tag{6.6}$$

where $\hat{s}_{i,j}$ is the probability that household i chooses plan j, $c^e_{i,j}$ is the expected medical claims cost of household i in plan j, and $\kappa_{i,j}$ adjusts costs to account for possible government cost-sharing subsidies.[25]

The parameter vector $\beta_{4,j}$, specific to an insurance carrier and metal tier, adds flexibility to adjust observed costs for the risk adjustment rules that allow the insurance carrier to insure themselves against some fraction of the health spending risk. In addition, $\beta_{4,j}$ could also capture a markup over variable costs, as in Bundorf et al. (2012): we would interpret a value greater than one as a markup for that carrier and plan type. We hold this value fixed in our counterfactual simulations.

The second term in the premium-setting equation reflects plan-level fixed costs. A_j includes observable inputs into each carrier's plan-level administrative costs. We approximate these administrative costs in two ways. First, we use carrier by time period indicators to reflect common per-plan or per-plan-region administrative costs incurred by the carrier in each plan year. Second, as a robustness analysis, we approximate plan-level administrative costs using observed prior-year administrative cost measures that carriers report in MLR filings. From these MLR data, we compute the carrier's total per-enrollee monthly administrative cost as the sum of taxes and fees, wellness activities, and general administrative expenses. We allocate these costs to the plan level based on plan enrollment by geographic market. Finally, η_j is an element of the carrier's administrative costs that the econometrician does not observe.

We begin by estimating equation (6.6) via ordinary least squares, using 2015-16 data since

[25] We adjust for subsidies for household i by multiplying expected costs by $\kappa_{i,j} = x^o_j / x_{i,j}$, where x^o_j is the actuarial value of the plan without any cost-sharing subsidies. This term accounts for the fact that the insurer's realized costs do not include the cost-sharing subsidies that entered the model of consumer choice. Given that spending follows an exponential distribution and given our first-stage estimates, we can write $c^e_{j,t}$ (before adjusting for subsidies) as:

$$c^e_{i,j} = E[c_{i,j} | \widehat{\alpha}_i, \widehat{\omega}_i] = \frac{x_{i,j} + \widehat{\omega}_i x^2_{i,j}}{\widehat{\alpha}_i}.$$

carriers may not have reached an equilibrium in 2014. We then adopt a two-stage least squares approach to address the possibility that η_j may be correlated with the predicted market share and claims variables.[26] We use instruments, Z_j, that are assumed to be correlated with $s_{i,j}$ and $c^e_{i,j}$ but mean independent of η_j: $E[\eta_j|Z_j] = 0$. Our instruments act as demand shifters. These include the number of plans offered to households in the same market and the fraction of households in the market who are subsidized. We also use lagged values of these two instruments.[27] Lastly, we employ year fixed effects to account for time-varying administrative costs.

6.4 Consumer surplus

To compute consumer surplus, we transform our surplus expression from Section 5.3 to match our empirical specification. Integrating over our chosen distribution of λ, we define the certainty equivalent of plan j for individuals of type i as follows:

$$e_{i,j} = \frac{1}{\alpha_i - \psi_i}\Big[x_{i,j} + \frac{x^2_{i,j}\omega_i}{2} - (\alpha_i - \psi_i)p_{i,j} + \beta_0 X_j + \epsilon_{i,j}\Big] \tag{6.7}$$

The ex ante consumer surplus then becomes:

$$CS_i = E_\epsilon\left(\max_j e_{i,j}\right) = \frac{1}{\alpha_i - \psi_i} log \sum_{j=1}^{J} exp\left(x_{i,j} + \frac{x^2_{i,j}\omega_i}{2} - (\alpha_i - \psi_i)p_{i,j} + \beta_0 X_{i,j}\right) \tag{6.8}$$

Each household chooses from J plans. We compute this measure both for individual market and small group market households.[28]

When we compute small group consumer surplus as described above, we face one additional hurdle: plan choices in the small group market do not match the revealed preferences of small group

[26] For example, the instruments address the potential for unobserved quality variation within a plan—e.g. across enrollee types. This variation is not correlated with premiums because of the institutional restrictions on premium setting already described, and therefore does not bias the premium coefficient. However, the variation does affect shares and thus is correlated with η.

[27] In our robustness analysis that uses administrative costs from MLR data, we also explore a second set of instruments that affect insurer costs through economies of scale. These include the total number of subscribers of the carrier's plans in other states and the total number of subscribers of the carrier's plans in the individual market in other states.

[28] Because we do not estimate a model of premium setting in the small group market, we calculate our small group consumer surplus estimates using observed plan premiums. This approach contrasts with how we calculate consumer surplus in the individual market, both prior to combining the two markets and in counterfactuals; in both cases, we derive these estimates using simulated premiums.

households forced onto the individual market (i.e. our forced switcher sample). The distinction implies that employers influence plan choice. To allow employer preferences in the choice problem—and ultimately to derive an alternative consumer surplus value—we assume a simple model of employer plan choice. In this model, the employer chooses a default silver-tier health plan for its employees, but employees may opt for an alternative plan by paying a fee, $f_{i,j}$. We report an alternative consumer surplus measure that includes this fee in Appendix D.2. Our findings are identical to the 'no fee' case, but with a baseline small group surplus that is smaller by a fixed amount.

7 Results and Counterfactuals

7.1 Demand estimates

We report our underlying demand estimates in Table 3. The underlying estimates are difficult to interpret on their own, with the exception of the estimated preference for silver plans among consumers eligible for cost-sharing subsidies. This coefficient is large and positive in the main sample and the switcher sample. For this subsidy-eligible population, the regulatory design of the subsidies strongly influenced enrollees to choose silver plans.[29]

We use the underlying demand estimates to derive several parameters of economic interest, reported in Tables 4 and 5. Our derived estimates illustrate our model's predictions of household spending, moral hazard, and risk aversion for both the individual market and forced switcher populations.

To illustrate the empirical distribution of spending given our sample demographics, Table 4 translates the estimated parameters into expected non-discretionary spending, $E(\lambda_i)$, separately for the individual and small group markets. In these statistics, we also extrapolate the switcher

[29] Households may be eligible for both cost-sharing and premium subsidies. We account for premium subsidies directly by measuring premiums net of subsidies for each household. We also account for cost-sharing subsidies directly in the actuarial value variable. For example, the silver plan actuarial value increases from .7 to .87 for households whose income falls between 150% and 200% of the federal poverty level. We include a distinct interaction term for cost-sharing subsidy eligibility and silver plan choice because eligible consumers must purchase a silver plan to benefit from cost-sharing subsidies. The interaction term captures how advertising or enrollment navigators might encourage silver plan enrollment more than would be predicted simply through the effect on actuarial value.

preferences to apply to the entire small group market. We find that overall $E(\lambda_i)$ is lower in the small group market relative to the individual market: on average, insurers incur costs of \$419 per month for small group enrollees relative to \$547 per month in the individual market.[30] The overall difference arises because of the much higher proportion of small group enrollees under age 50 and with no dependents. If we look only at households with dependents, we find non-discretionary spending to be higher in the small group market. Because demographics and family composition explain part of the difference in medical claims costs between the two markets, we condition on these factors in our counterfactual simulations.

Table 5 provides estimates of moral hazard and risk aversion. Moral hazard is higher for the sample of switchers than it is in the individual market. Under our preferred specification, with $\underline{c} = 20$, we predict that moving switchers from zero insurance to full insurance would increase medical spending by 22%. For individual market enrollees, the increase would be 11%. Finally, the estimated CARA coefficient is 6.8×10^{-4} for the individual market and 2.7×10^{-3} for the sample of switchers. In gamble terms, our estimate implies that a household would be indifferent between receiving \$0 and a 50-50 gamble in which it earns \$100 or loses \$99.32 (for the individual market) or \$97.40 (for switchers). That is, the individual market households and switchers have quite similar levels of risk aversion. Our estimated magnitudes for both moral hazard and risk aversion are in line with estimates in the previous literature (Handel, 2013; Einav et al., 2013; Marone and Sabety, 2020; Ho and Lee, 2020).

What do these estimates imply for our counterfactual analyses? The two groups of enrollees are similar on most dimensions. We do find small group enrollees to have a slightly higher preference for silver plans if subsidized and a slightly higher moral hazard parameter, which may imply a stronger preference for high-coverage plans than in the individual market (Einav et al., 2013). These small group consumers, on average, also have lower non-discretionary spending, which could lower costs to insurers operating more generous plans. It is not clear, however, that the differences are large enough to address fully the adverse selection in the individual market. The new equilibrium will

[30] Our non-discretionary spending measure, $E(\lambda_i)$, includes both insurer and out-of-pocket costs but excludes moral hazard spending. This measure differs from the mean spending statistics in Table 2, which exclude out-of-pocket costs but include any added insurer spending from moral hazard.

depend partly on the extent to which insurers pass through claims costs to plan premiums. We discuss our pass-through estimates next.

7.2 Cost estimates

Table 6 contains the estimates of the premium setting equation, Equation 6.6. Columns 1 and 2 summarize the OLS regression of predicted plan premium revenues on predicted medical claims costs, both with and without payer-year fixed effects. Payer-year fixed effects act as our proxy for administrative costs. In Column 3 we report an instrumental variables specification in which we instrument for predicted claims costs using the number of plans in the market and the fraction of households in the market who are subsidized. Our estimated coefficient on claims costs in all three specifications lies between 0.73 and 0.75 and is highly statistically significant. In words, approximately 75% of claims costs are passed through to premiums. The coefficient on claims costs suggests that additional risk adjustment revenue, after accounting for cost-sharing subsidies, makes up the balance.

In Column 4, we run a specification that replaces the payer-year fixed effects with an alternative proxy for administrative costs. Here, we employ measures of lagged plan administrative costs and year fixed effects to capture plan-level administrative costs. We instrument for both predicted claims costs and administrative costs. The coefficient on claims falls to 0.55 and the coefficient on administrative costs is 0.62; both are statistically significant. We view this specification as a useful sanity check on our estimates that exploit fixed effects only. However, since we do not observe administrative costs for every carrier in the data, we use column (3) as the main specification for our counterfactual analyses.

7.3 Counterfactual simulations

We use the estimated model to evaluate market outcomes under three alternative market designs. First, we imagine an environment in which regulators close the small group market. Employees of these small groups must either purchase insurance on the individual market, where they may be eligible for federal premium and cost-sharing subsidies, or choose to be uninsured. For this

population, premiums are no longer tax-exempt and employers do not contribute to premium costs. We use this simulation to assess whether small group enrollees' preferences for coverage, and their relatively low health care needs compared to individual market households, can mitigate adverse selection were regulators to pool the markets.

Second, we study the effect of allowing small groups to offer their employees the choice of individual coverage with tax benefits and subsidies. Specifically, we approximate new federal regulations from 2020 allowing health reimbursement arrangements paired with individual market coverage (ICHRA). In this extension, small group employers no longer offer a group insurance option. Instead, these employers provide funds for employees to purchase coverage on their own.

Finally, third, we again allow small employers to offer extended ICHRA insurance to their employees but now require all households in this pool to purchase at least bronze-level coverage. While constraining the choice set is not welfare-maximizing for those consumers who prefer to be uninsured, we use this simulation to examine overall surplus across all consumers when employers mandate insurance participation.

Under each counterfactual environment, we predict changes in equilibrium insurance participation, plan and metal-tier level market shares, and the premiums consumers pay and insurers collect as revenue. We do so using the population of small group enrollees in Oregon in the year 2016, two years after the regulation of both the small group and individual insurance markets under the ACA. Using the new equilibrium, we measure the market outcomes relative to the baseline equilibrium in terms of employer spending, government spending in the form of both tax expenditures and premium subsidies, and consumer surplus.

To conduct these counterfactual simulations, we need a method to find new equilibrium premiums as the enrolled population changes. Under our assumption of a perfectly competitive individual insurance market, Azevedo and Gottlieb (2017) provide an algorithm to compute this equilibrium. We describe the algorithm in detail in Appendix F. In brief, we assume consumers have the same choice of carriers, metal tiers, and plan types as in the observed market. Following Azevedo and Gottlieb (2017), we augment our pool of households with a mass of 'behavioral consumers' who incur zero covered health costs and choose each available contract with equal probability; the inclu-

sion of these behavioral types ensures that all contracts are traded.[31] We then apply a fixed point algorithm in which, in each iteration, consumers choose contracts according to their preferences, taking prices as given. Prices are adjusted up for unprofitable contracts and down for profitable contracts until we reach an equilibrium.

Our key counterfactual results appear in Tables 7 and 8. Table 7 summarizes predicted consumer surplus and market shares by metal tier for each scenario while Table 8 reports equilibrium premiums. We report two measures of equilibrium premiums: standardized gross premiums and population net premiums. The standardized gross premium equals a weighted average of the premium that a single 40-year-old adult would face in each region for each plan, weighted by the empirical plan-region market shares. The net premiums represent the average net premiums that households in the sample face, accounting for the distribution of choices by age, family size, and region.

We distinguish market outcomes from the perspective of individual market enrollees as of 2016 (Panel A) and the perspective of small group enrollees (Panel B) who must shift to individual insurance under the new market designs. We describe the results of the three counterfactuals in turn.

Simulated market outcomes. Before turning to our alternative market design counterfactuals, we simulate market outcomes for the individual market under current conditions. Column (1) of Tables 7 and 8 presents these in-sample predictions. We find that adverse selection severely affects the market for gold-tier plans. On average, a 40-year old single buyer of a gold plan faces premiums of $685 per month. Only 7% of households purchase these plans.[32] We find the average consumer surplus in the individual market is $234 per month for the set of households for which the measure is defined, excluding outlier households with large health spending.[33]

[31] Our reported simulations assume 1% of the sample are behavioral types. Increasing this share to 5% has very little effect on the results.

[32] These predictions are different from the raw data summarized in Table 2, where average gold plan premiums are reported as $533. There are two reasons for this difference. First, the data in Table 2 cover the years 2014-2016 while our simulations consider only 2016. Second, in our equilibrium we assume zero markups, i.e. we do not allow insurers to experience gains or losses. In 2014, many plans in the individual market in fact reported losses.

[33] Consumer surplus is defined for the set of households where $\alpha > \psi$. We exclude the top 5% of households with the highest expected non-discretionary spending. If we include these households, average consumer surplus would equal $737 per month.

Removing small group employer coverage. When we act as a regulator and eliminate small group employer coverage—in effect pooling the small group and individual markets—we predict some reduction in adverse selection for the individual market. In addition, both the government and employers experience savings in this environment.

The effects we observe reflect lower non-discretionary spending among small group employees. When healthier and less costly enrollees enter the individual market with preferences for insurance, some may choose generous coverage, particularly if they are eligible for government cost sharing or premium subsidies that exceed the subsidies in their employer plans. However, our estimated preferences also suggest many small group households will choose to be uninsured when forced to pool. We estimate 61% of small group households will exit coverage. This uninsured population draws disproportionately from the set of lower spending households: small group households choosing uninsurance have expected non-discretionary spending of $240/month versus $671/month for insured households.

When the healthiest small group consumers choose to become uninsured, we find relatively modest benefits from pooling for individual market households. While premiums of gold plans in particular fall to $458 per month for the standardized enrollee, Tables 7 and 8 show that market shares and premiums for other tiers experience only small changes. The share of households who choose uninsurance falls by three percentage points. Those uninsured households entering coverage are healthier than the incumbent pool, with expected non-discretionary spending of $225/month versus an average of $851 across all enrollees.[34] Overall, the changes to market prices and participation generate an average surplus gain of $18 per month for an individual market household.

For households who shift from employer coverage to the individual market, the new market design offers both benefits and costs. These households lose large tax subsidies and employer contributions toward premium payments. Further, households with a high willingness to pay for insurance enter into a market which suffers from adverse selection. On the other hand, small group market households gain federal premium subsidies, access insurance premiums that are not subject to the markups present in the current small group market, and can choose to be uninsured.

[34] For this statistic, we define uninsured households as households who are predicted to choose uninsurance with greater than 50% probability under current market conditions.

For this population, we find that gross premiums for bronze and silver plans are lower in the merged market than in the former small group market. Bronze and silver plan standardized gross premiums fall by $77 and $113 per month, respectively. Premiums for gold plans, however, are $50 larger in the merged market.[35] For the average household, the gains from removing plan markups dominate losses from both pooling with a more adversely selected pool and from the loss of tax subsidies and employer contributions. We find consumer surplus for small group market households increases by $218 from the baseline estimate. This change partly reflects the new equilibrium premiums and partly reflects the expanded choice set. In the baseline market design, we assume small group households cannot choose to be uninsured. The revealed preferences of the switcher sample show that many households would prefer to exit coverage rather than participate in the market. Providing the option to opt out of the health insurance market generates consumer surplus gains to many households.[36]

Government expenditures on insurance coverage for small group employees are significantly lower in the merged counterfactual market than under current market conditions. Under employer coverage, premiums are tax-exempt. Thus, we measure government expenditures at the household level as the household's average tax rate multiplied by the observed premium of their chosen plan. In the individual market, government expenditures for a household equal expected premium subsidies. We find that government expenditures for small group employees decrease from $124 per month per household under employer coverage to $71 in the merged market. That is, although the government did not subsidize small group household premiums prior to the pooling, the value of the implicit tax subsidies they 'paid' in that environment exceeds the cost of premium subsidies for those low-income consumers who shift from the small group to the individual market. Part of this decrease reflects the prediction that a meaningful portion of small group households will choose to be uninsured in the pooled environment.

[35] The average gross premiums for a standardized 40-year-old enrollee are not equal in the small group and individual markets in column 2 of Table 8. The slight differences reflect different weighting in the two populations. In particular, the small group and individual market enrollment populations differ in (a) their carrier and plan type choices within metal tier and (b) their distribution across regions.

[36] When we assume consumers cannot choose to exit insurance, our baseline estimate of surplus is -$86. If we assume that consumers can choose to be uninsured in the baseline case, our predicted average consumer surplus would rise to $164. In addition, we assume a fixed employer contribution of 65% of post-tax premiums. If in practice employers offer more generous subsidies, our measure of consumer surplus would increase further.

Employers also see substantial savings from the removal of group coverage. Under the assumption that employers contribute 65% of (post-tax) premium costs, the average share reported in national employer surveys, we find that small employers spend an average of $278 per employee per month on health coverage. In the counterfactual scenario, we fix the contribution at $0. In total, removing small employer group coverage and forcing small group employees to choose plans on the individual market would yield approximately $419 million in yearly savings for small employers in this sample. Here, wages are likely to adjust to this change in employer spending. In part for this reason, in later counterfactuals we hold the employer's per-employee contributions fixed.

Extending ICHRA coverage. We next simulate the effect of allowing small group employees to purchase insurance on the individual market with tax-exempt and employer-subsidized premiums. In effect, we introduce a population to the individual market who are less price sensitive and have lower health spending on average than the existing individual market population. We find that extending ICHRA coverage for small employers mitigates adverse selection among gold tier plans in the individual market. In the merged market, 15% of small group employee households choose gold plans. Households enrolled in the individual market prior to the ICHRA extension see their average standardized gross premium for gold plans drop from $685 to $354 per month. The share of these households choosing gold plans increases to 11%.

In terms of consumer surplus, we find that merging the two markets under an ICHRA policy improves welfare. Individual market households, on average, see consumer surplus increase by $28 per month.[37] For small group employees, average gross premiums fall in all three metal tiers. Additionally, small group employees continue to forego insurance coverage when pooled. Small group consumer surplus increases substantially, with average gains of $263 per month, given the added choice to exit coverage.

In solving for the equilibrium premiums in the individual market, we adjust for the design of government premium subsidies, which are a function of both household income and the premium of the second-cheapest silver plan offered in the market. We also implement a simplification to

[37] Again, we compute this average across households for whom consumer surplus is measurable and we exclude outliers ($\alpha - \psi > .05$).

adjust our equilibrium premium search for the subsidy design in the small group market. We fix the small group per-household subsidy as a lump sum equal to the subsidy in the base period. Thus, we assume that small group market households receive the same employer contribution in the counterfactual and base period.[38]

Despite these subsidy policies remaining fixed, government and employer expenditures decrease to 34% and 43% of their previous levels, respectively. That many small group households choose to be uninsured drives the decline in expenditure.

If employers increased wages in response to this reduction in spending, we might find average consumer surplus increases yet further. However, we caution that much of the welfare improvement—both the gains to small group households and the government and employer savings—are due to small group households shifting to uninsurance. These households are healthier than the typical pool but nonetheless have spending needs; we predict non-discretionary spending of $358/month for those exiting coverage. If households underestimate the future cost of uninsurance, our revealed preference framework would also omit this cost. As a comparison, in our final counterfactual simulation, we quantify the welfare changes when small group enrollees face ICHRA coverage choices but must maintain coverage.

Extending ICHRA coverage with mandatory enrollment. The final column of Tables 7 and 8 considers a variant on the extended ICHRA simulations that addresses the main caveat to the previous results: the increase in the uninsured share among small group households. We now assume that small employers can induce all employees to choose to be insured, whether through influence or directive, even after shifting coverage elections to the individual marketplace. In this setting, we find that pooling the small group and individual market households further mitigates the adverse selection problem among gold-tier plans. For individual market households, the average standardized gross premium for gold plans drops from $685 to $276 per month; the share of

[38] In an alternative set of assumptions, we set small group market subsidies in a similar way to the individual market: we fix the proportion of a small group market plan's premium that is covered by premium subsidies equal to the baseline share we observe in the data. In this setting, employer contributions scale with premiums. Our results from this approach are qualitatively similar to the results presented in the text.

households choosing gold plans increases to 14%. Small group employee households, now denied the opportunity to choose uninsurance, frequently choose silver (35%) or gold plans (36%).

These new equilibrium premiums and plan choices generate a larger improvement in consumer surplus for individual market households—an increase of $44 per month— and a larger reduction in government spending for individual market enrollees. Government and employer spending on small group households continues to fall, but by less than in the previous simulation when households could choose to be uninsured. Our measure of small group consumer surplus rises as well, by an average of $57 per month. However, small group consumer surplus gains are significantly smaller than in alternative counterfactuals because some households prefer to be uninsured.

7.4 Role of insurer market power

In our main analysis, we assume perfect competition in the individual insurance market both before and after changes to the pool of potential enrollees. In this robustness analysis, we ask: how would equilibrium outcomes differ if individual market insurers charged positive markups in the pooled environment?

To answer this question, we start with a competitive individual insurance market, the same as in our baseline analysis before pooling. We then conduct a counterfactual in which we simultaneously add markups to individual market plans and add small group households to the individual insurance pool. We allow plan-by-plan markups that range from 0% to 25%. The 25% level approximates the maximum allowable medical markup under medical loss ratio (MLR) regulation.[39] For the small group pooling, we conduct the exercise in the same way as in two of the key counterfactuals from our main analysis: (a) we mimic an extension of ICHRA accounts to all small group households, and (b) we repeat the ICHRA exercise but require all small group households to purchase at least bronze-level coverage.

[39] Under MLR rules, in each market segment carriers report the sum of the total costs of their enrollees' clinical care with the costs of any quality improvements programs they conduct. Call this total cost C. The MLR constraint requires that at least 80% of premium dollars, p, collected in a market segment must contribute to paying these costs. Thus, when binding, the constraint implies $0.8 * p = C$ or equivalently $p = 1.25 * C$. In our implementation, we deviate from the exact MLR criterion in that we mark up premiums 25% over total realized costs, which include administrative costs. We define premiums net of risk adjustment payments.

To determine equilibrium premiums in this counterfactual environment, we apply a modified version of the algorithm from Azevedo and Gottlieb (2017) but now include a fixed markup by plan; the markup is fixed in that it does not vary with the equilibrium outcome.[40] In addition, our approach of fixed markups per plan does not allow cross-subsidization of plans within an insurer—e.g. an insurer cannot subsidize an unprofitable gold plan with a profitable bronze plan. As discussed in Azevedo and Gottlieb (2017), one can micro-found this restriction with a strategic model with differentiated products. If an insurer taxes one plan to subsidize another, it risks being undercut on the taxed plan and only selling the money-losing option.[41]

Individual market consumer surplus We first evaluate how our new counterfactual—one that combines both market pooling and insurer markups—affects the level of surplus of individual market consumers. The effect of markups alone is simple to predict: when premiums increase due to markups, the level of consumer surplus for individual market households falls relative to the competitive outcome. To quantify the overall surplus change from both pooling and markups, we design our measurement exercise to determine the level of markup at which the surplus gain from pooling with healthier enrollees just offsets the loss from higher premiums due to markups.

In Table 9, we report the equilibrium premiums that result from pooling under a range of markups. We start with the 0% markup case, which repeats the findings of our main analysis. Under the extended ICHRA counterfactual without markups, the standardized premiums for all plan types fall with pooling. This decrease in premiums generates an average gain of $28 in surplus for individual market households.

When we add positive markups, the equilibrium premiums increase relative to the competitive pooled environment: comparing a 25% markup to 0% markup, standardized plan premiums increase by $23, $50, and $264 for bronze, silver, and gold plans, respectively. However, for our welfare

[40] An alternative to our approach would be to apply the best-response iteration algorithm of Hastings et al. (2017), described in Appendix Section A.2 of their paper. With this alternative, we could allow markups to vary with the elasticities of the enrolled population. This approach, however, would require modification to incorporate the MLR constraint on premium setting in the individual market. Our fixed markup of 25% approximates an equilibrium in which the MLR constraint binds on all insurers.

[41] The strategic model that Azevedo and Gottlieb (2017) describe can be seen as a limiting case of the differentiated product models common in the industrial organization literature. The concept applies most readily in settings that feature many firms with relatively small scale.

analyses, the key comparison is not whether pooled premiums increase with greater markups, but whether the pooled premiums with markups exceed the premiums in the baseline environment with neither pooling nor markups. When we conduct this comparison, we find that bronze and silver plan premiums exceed the initial equilibrium level under either 10% or 25% markups. Gold plan premiums, in contrast, remain lower. Relative to an initial standardized level of $685, gold premiums equal $449 and $618 in an equilibrium environment with markups of 10% and 25%, respectively.

We translate these changes in equilibrium premiums into changes in average consumer surplus. Relative to the surplus gain of $28 for individual market households in the competitive setting, adding 10% markups lowers the gain from pooling to $13 in surplus. With 25% markups, individual market households lose an average of $8 from pooling. That is, the costs of markups outweigh the benefit of adding healthier enrollees to the pool. The cost and benefits roughly equal each other at a markup of 20%.

Small group market consumer surplus Even with the introduction of insurer markups, consumers in the individual market may benefit from pooling because of the entry of healthier small group households. We now consider how markups affect the surplus of small group households. In the counterfactuals involving the use of ICHRAs, the small group enrollees enter with employer subsidies and tax-advantaged premiums. These features reduce small group enrollees' effective price sensitivity. We allow small group households to choose plans in the individual market based on their preferences for premiums and other plan attributes, accounting for fixed plan markups.

In our original ICHRA counterfactual scenario with a competitive individual market, we find small group households gain $263 from pooling even though they enter a more adversely selected market. Two main features of the pooled market generate this gain. First, small group households benefit from the option to exit coverage. And second, premiums in the small group market exceed those in the individual market for the same coverage level, reflecting higher markups over medical costs.[42] Those medical markups are higher either because of weaker competitive effects or because

[42] A third feature, that small group households no longer face 'tiered composite' community rated premiums in the pooled market, explains $8 of the difference in surplus.

of higher administrative costs, notably in the form of brokers' fees. The Kaiser Family Foundation (2020) reports a $9.82 per-member per-month differential between broker fees in the small group and individual insurance markets in Oregon in 2016.[43] Thus, when we allow markups in the individual market, the total welfare effect for small group households will depend on how the composition of the choice set changes and how the marked-up premiums in the individual market compare with the baseline small group premiums.

Under the extended ICHRA counterfactual, we find small group households gain surplus from pooling even when there are strictly positive markups in the individual market. On average, households gain surplus of $263 with 0% markups, $253 with 10% markups, and $238 with 25%, relative to the baseline small group market. A key driver of this surplus is the household's new option to exit coverage in the extended ICHRA counterfactual.

To separate this added-choice effect from the effect of changing markups, we repeat the ICHRA experiment but remove the small group household's choice to become uninsured. In our 'forced insurance with ICHRA' counterfactual with $0 individual market markups, we find small group consumers gain $57 in surplus relative to the baseline when pooled in the individual market. If we layer on markups of 10% and 25% in the individual market, the gain in surplus to small group consumers equals $45 and $28, respectively.

We observe a positive surplus change even with market power in the individual market and even when small group enrollees cannot exit coverage. The gain stems from the net value to small group consumers from decoupling their premiums from other enrollees within an employer and from the fact that higher administrative costs in the small group market can lead to higher premiums in that market, even if both markets share the same percentage markup. Given the average household size of 1.54 members in the small group in 2016, the differential in average broker fees alone can explain $15 of the surplus gains per month when markups equal 25%.

[43]The reported levels of per-member per-month (PMPM) broker fees vary by insurance carrier and plan year. In data reported by the Kaiser Family Foundation for Oregon in 2016, broker fees in the individual market equaled $9.82 PMPM and $19.70 PMPM in the small group market.

8 Conclusion

We assess the impact of segmentation on market outcomes in US health insurance. We focus on an ideal laboratory to study the effect of segmentation: the division between individual market coverage through insurance marketplaces and employer coverage in the small group market. Our analysis highlights two economic mechanisms at work. First, following a standard adverse selection analysis, we can compare the welfare gains and losses from pooling as a function of the relative costliness and preferences of the households in the two segments. Second, we can consider the welfare costs of agency in the employer market. The employer's choice of coverage on behalf of its employees, and the role of broker intermediaries in that choice, can generate higher markups in the segment.

Using detailed data on plan choices and health spending that allows us to track subscribers across insurance plans and markets, we estimate preferences for both individual market and small group market households. We then use these estimates, along with a model of insurance plan premium setting, to simulate market outcomes in counterfactual scenarios where we integrate the markets.

We identify particular market integration policies that can improve welfare for households in both market segments. In an initial simulation we show that removing small group coverage would prompt many small group employees to choose to be uninsured, with relatively small changes to market conditions in the individual market. Employer and government savings, however, would be large. Our more policy-relevant simulations explore the effect of extending the ICHRA rule, where employers contribute funds toward plans purchased on the individual market. We show that this change would mitigate severe adverse selection in the individual market. Market pooling would also benefit small group households and would reduce employer and government spending.

Why do small group households benefit from pooling with sicker individual market households? In our analysis, small group members avoid higher markups in their segment by shifting to individual insurance. We show that one driver of these markups is higher administrative costs in the form of brokers' fees. In addition, we illustrate that in Oregon in 2016, even if insurers responded to pooling in the individual market by marking up premiums, individual market and small group households

would still prefer pooling for a range of markups.

Finally, while our measurement focuses specifically on the small group and individual market segments in Oregon, the tools we develop can apply more broadly. For example, the differential in premiums between the small group and individual markets in other states, as we report in Section 4.2, suggests that a similar motivation for pooling may exist in those locations. One could also use our framework to study related policies, such as expanding eligibility for Medicare to populations younger than 65 years old (Rae et al., 2021). We emphasize that the extent of the welfare change from such pooling depends not only on the differential in health spending, but also on household preferences and markups in each insurance segment.

References

Abraham, Jean M, Anne B Royalty, and Coleman Drake, "Plan choice and Affordability in the Individual and Small-Group Markets: Policy and Performance—Past and Present," *Health Affairs*, 2019, *38* (4), 675–683.

AHRQ, "2016 Medical Expenditure Panel Survey-Insurance Component, Tables of Private-Sector Data by Firm Size, Industry Group, Ownership, Age of Firm, and Other Characteristics, Table Series I.A.2," 2016.

Atal, Juan Pablo, Hanming Fang, Martin Karlsson, and Nicolas R Ziebarth, "Long-term Health Insurance: Theory Meets Evidence," 2020. National Bureau of Economic Research Working Paper.

Azevedo, Eduardo M. and Daniel Gottlieb, "Perfect Competition in Markets With Adverse Selection," *Econometrica*, 2017, *85* (1), 67–105.

Blumberg, Linda J, Shanna Rifkin, Sabrina Corlette, and Sarah J Dash, "Stabilizing Premiums Under the Affordable Care Act: State Efforts to Reduce Adverse Selection," Technical Report 2013.

Bundorf, M Kate, Jonathan Levin, and Neale Mahoney, "Pricing and Welfare in Health Plan Choice," *American Economic Review*, 2012, *102* (7), 3214–48.

Claxton, Gary, Matthew Rae, Michelle Long, Nirmita Panchal, Anthony Damico, Kevin Kenward, and Heidi Whitmore, "Employer Health Benefits: 2015 Annual Survey," *Kaiser Family Foundation and Health Research and Educational Trust*, 2015.

Corlette, Sabrina, Jack Hoadley, Kevin Lucia, and Dania Palanker, "Small Business Health Insurance and the ACA: Views from the Market 2017," Technical Report 2017.

Cutler, David M and Sarah J Reber, "Paying for Health Insurance: the Trade-off Between Competition and Adverse Selection," *Quarterly Journal of Economics*, 1998, *113* (2), 433–466.

Dafny, Leemore S, Yin Wei Soon, Zoë B Cullen, and Christopher T Stanton, "How Has Covid-19 Affected Health Insurance Offered by Small Businesses in the US? Early Evidence from a Survey," *NEJM Catalyst Innovations in Care Delivery*, 2020, *1* (4).

Diamond, Rebecca, Michael J. Dickstein, Timothy McQuade, and Petra Persson, "In-

surance without Commitment: Evidence from the ACA Marketplaces," 2021. NBER Working Paper.

Einav, Liran, Amy Finkelstein, and Pietro Tebaldi, "Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies," 2019. NBER Working Paper.

_ , _ , Stephen P Ryan, Paul Schrimpf, and Mark R Cullen, "Selection on Moral Hazard in Health Insurance," *American Economic Review*, 2013, *103* (1), 178–219.

Fleitas, Sebastian, Caitlyn Fleming, Gautam Gowrisankaran, and Tony LoSasso, "Selection, Subsidies, and Welfare in Health Insurance: Employer-Sponsored Health Insurance Versus the ACA Marketplaces," 2021. Working Paper.

_ , Gautam Gowrisankaran, and Anthony Lo Sasso, "Reclassification Risk in the Small Group Health Insurance Market," 2018. National Bureau of Economic Research Working Paper.

Geruso, Michael, Timothy Layton, and Daniel Prinz, "Screening in Contract Design: Evidence from the ACA Health Insurance Exchanges," *American Economic Journal: Economic Policy*, 2019, *11* (2), 64–107.

Hall, Mark A and Michael J McCue, "The Health of the Small-Group Insurance Market," *Issue brief (Commonwealth Fund)*, 2018, pp. 1–9.

Handel, Ben, Igal Hendel, and Michael D Whinston, "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk," *Econometrica*, 2015, *83* (4), 1261–1313.

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

Hastings, Justine, Ali Hortaçsu, and Chad Syverson, "Sales Force and Competition in Financial Product Markets: The Case of Mexico's Social Security Privatization," *Econometrica*, 2017, *85* (6), 1723–1761.

Ho, Kate and Robin S Lee, "Health Insurance Menu Design for Large Employers," 2020. National Bureau of Economic Research Working Paper.

Jaffe, Sonia and Mark Shepard, "Price-linked Subsidies and Imperfect Competition in Health Insurance," *American Economic Journal: Economic Policy*, 2020, *12* (3), 279–311.

Marone, Victoria and Adrienne Sabety, "Should There Be Vertical Choice in Health Insurance

Markets?," 2020. Working Paper.

Oregon Health Insurance Marketplace, "Small Business Health Care Tax Credit," 2020.

Polyakova, Maria and Stephen P Ryan, "Subsidy Targeting with Market Power," 2019. National Bureau of Economic Research Working Paper.

Providence Health Plan, "Producer Compensation Plan: Fully Insured Commercial Business with Attachment B," 2011.

Rae, Matthew, Cynthia Cox, Krutika Amin, and Tricia Neuman, "How lowering the Medicare eligibility age might affect employer-sponsored insurance costs," *Kaiser Family Foundation*, 2021.

Tebaldi, Pietro, "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA," 2017. Becker Friedman Institute for Research in Economics Working paper.

_ **, Alexander Torgovitsky, and Hanbin Yang**, "Nonparametric Estimates of Demand in the California Health Insurance Exchange," 2019. National Bureau of Economic Research Working Paper.

The Kaiser Family Foundation, "Health Insurance Market Reforms: Rate Restrictions," 2012.

_ , "State Health Facts: Health Coverage and Uninsured," 2017.

_ , "State Health Facts: Broker Compensation by Health Insurance Market," 2020.

Tikkanen, Roosa, Robin Osborn, Elias Mossialos, Ana Djordjevic, and George Wharton, "International Profiles of Health Care Systems," Technical Report, The Commonwealth Fund 2020.

U.S. Census Bureau, "Table 1: Coverage Numbers and Rates by Type of Health Insurance: 2013, 2016, and 2017," 2018.

9 Tables and Figures

Table 1: Summary statistics on demographics variables

	Individual Market			Small-Group Market		
Variable	**Mean**	**S.D.**	**Median**	**Mean**	**S.D.**	**Median**
Single-membered	0.70			0.75		
Married, no dependent	0.14			0.08		
Not married, with dependent(s)	0.07			0.07		
Married, with dependent(s)	0.09			0.10		
Number of dependents	1.93	1.09	2.00	2.21	1.25	2.00
HH health status score	1.39	2.44	0.62	1.25	2.20	0.60
Income (as ratios of the FPL)	2.46	0.29	2.47	2.38	0.29	2.35
Age	46.96	11.75	48.00	42.62	11.28	42.00
Over-50	0.42			0.28		
Living in rating areas 1, 2, or 3	0.69			0.78		
Number of						
unique HHs	354,366			218,827		
HH-year observations	512,515			383,036		

Note: This table presents demographic summary statistics on the population of households in Oregon choosing insurance plans in both the individual and small group markets in years 2014-2016. We compute the sum of risk score for members of a household, where we predict each member's risk using the Johns Hopkins' ACG software. We calculate the number of dependents for the subset of households who have dependents. Rating areas 1-3 include the urban areas of Portland, Eugene, and Salem, respectively. Rating areas 4-7 include largely rural areas of the state.

Table 2: Summary statistics on insurance variables

	Individual Market			Small-Group Market		
Variable	**Mean**	**S.D.**	**Median**	**Mean**	**S.D.**	**Median**
Spending = 0						
Overall	0.30			0.32		
Bronze plans	0.46			0.43		
Silver plans	0.24			0.34		
Gold plans	0.13			0.25		
HH spending (if nonzero)						
Overall	591	1,431	127	524	1,293	127
Bronze plans	388	1,196	67	461	1,283	93
Silver plans	571	1,394	132	488	1,260	115
Gold plans	932	1,752	273	586	1,333	158
Monthly HH premiums						
Overall	364	248	280	432	275	317
Bronze plans	294	215	216	333	206	251
Silver plans	360	225	279	408	249	297
Gold plans	533	314	428	504	313	360
Market shares						
Bronze plans	0.30			0.14		
Silver plans	0.55			0.50		
Gold plans	0.14			0.36		
Number of insurers active						
in rating areas 1-3 (mean)	7.0			8.0		
in rating areas 4-7 (mean)	6.5			7.3		
Number of						
unique HHs	315,150			73,583		
HH-year observations	444,255			105,984		

Note: This table presents insurance summary statistics on the population of households in Oregon choosing insurance plans in both the individual and small group markets in years 2014-2016. In both markets, we omit households who choose grandfathered plans, catastrophic plans, platinum plans, or plans that are not observed in the SERFF data. We exclude 16,641 household-year observations from the small group sample for having a platinum plan. Our monthly spending variable includes all medical costs covered under the insurance plan but omits patient out-of-pocket expenses. Household premiums in the individual market reflect gross premiums by plan. Small group premiums reflect the gross premiums paid by the employer per household; the household's tax subsidy or employer subsidy are not included in the statistics shown.

Table 3: Main parameter estimates

	Main		Switchers	
	(1)	**(2)**	**(1)**	**(2)**
α				
Low HH health status score [0/1]	-0.156	-0.190	0.208	0.185
	(0.002)	(0.002)	(0.016)	(0.016)
HH health status score	-0.474	-0.451	-0.440	-0.426
	(0.000)	(0.000)	(0.004)	(0.004)
Dependents [0/1]	-0.367	-0.356	-0.441	-0.426
	(0.001)	(0.001)	(0.012)	(0.012)
Over-50 [0/1]	-0.796	-0.785	-0.675	-0.661
	(0.001)	(0.001)	(0.010)	(0.011)
ω				
Constant	-2.187	-1.504	-1.505	-1.154
	(0.010)	(0.006)	(0.069)	(0.053)
ψ				
Constant	-2.686	-2.558	-1.321	-1.288
	(0.012)	(0.010)	(0.025)	(0.024)
β_0				
Payer fixed-effects	✓	✓	✓	✓
Subsidized silver plan [0/1]	0.456	0.448	0.795	0.781
	(0.004)	(0.004)	(0.037)	(0.037)
$\underline{c}$	\$20	\$50	\$20	\$50
Number of				
HH-year observations	1,044,742		14,426	
insured HH-year observations	444,255		5,184	

Note: This table reports the maximum likelihood estimates from the demand specification in Equation 6.4. Columns under "Main" contain estimates from the population of individual market subscribers in the years 2014-2016. Columns under "Switchers" contain estimates from the sample of tracked households that we observe switching away from the small group market after years 2014 and 2015. As described in Section 6.1, we define the switcher population as households belonging to small groups that exited the insurance market in the prior year. For each sample, we run two specifications, defined by the cost-censoring threshold $\underline{c}$: in (1) $\underline{c} = \$20$ and in (2) $\underline{c} = 50$. Healthy [0/1] is an indicator equal to one if a household's health status score is below the 30th percentile of the distribution of scores. We define the household health status score as the sum of scores for insured members in the household, where we predict health status using Johns Hopkins' ACG software. Dependents [0/1] is an indicator equal to one if there is a dependent in the household. Over-50 [0/1] is an indicator equal to one if the primary subscriber is older than age 50. Subsidized silver plan [0/1] is an indicator equal to one if the relevant insurance plan is a silver plan and the household's purchase is subsidized through cost-sharing subsidies.

Table 4: Derived estimates - non-discretionary spending

	Individual			Small-Group		
Household type	**Mean**	**S.D.**	**Share**	**Mean**	**S.D.**	**Share**
E[λ_i]						
Overall	5.47	13.88		4.19	11.72	
No dependent, under-50	2.52	7.99	0.46	1.75	5.66	0.59
With dependent(s), under-50	6.18	13.90	0.12	7.11	15.18	0.13
No dependent, over-50	7.98	17.11	0.38	6.54	14.87	0.24
With dependent(s), over-50	13.59	22.27	0.04	16.19	23.98	0.04
N	512,515			383,036		

Note: This table describes the distribution of $E[\lambda_i]$ as implied by the maximum likelihood estimates in Table 3 separately for the individual and small group markets. The 'overall' row reports the expected underlying health costs across the full sample, using household level covariates in our specification of the parameters of the exponential distribution for λ_i. In the subsequent rows, we break down the sample by household type and compute $E[\lambda_i]$ within type. We also report the share of the sample that each household type represents.

Table 5: Derived estimates - moral hazard and risk aversion

	Main		**Switchers**	
	(1)	**(2)**	**(1)**	**(2)**
ω_i				
Constant	0.112	0.222	0.222	0.315
	(0.001)	(0.001)	(0.015)	(0.017)
ψ_i				
Constant	0.068	0.077	0.267	0.276
	(0.001)	(0.001)	(0.007)	(0.007)
Loss interpretation of ψ_i				
Constant	99.323	99.232	97.401	97.316
	(0.008)	(0.008)	(0.063)	(0.063)
$\underline{c}$	\$20	\$50	\$20	\$50
N	512,515		383,036	

Note: This table reports the derived moral hazard and risk aversion parameters implied by the maximum likelihood estimates in Table 3. Columns under "Main" denote estimates from the population of individual market subscribers in the years 2014-2016. Columns under "Switchers" denotes estimates from the sample of tracked households that we observe switching away from the small group market after years 2014 and 2015. As described in Section 6.1, we define the switcher population as households belonging to small groups that exited the insurance market in the prior year. For each sample, we run two specifications, defined by the cost-censoring threshold $\underline{c}$: in (1) $\underline{c} = \$20$ and in (2) $\underline{c} = 50$. ω_i is the multiplicative moral hazard parameter. ψ_i is the CARA risk aversion parameter. We also report risk aversion using the loss interpretation: we compare the utility of (a) a 50/50 gamble between losing X dollars and gaining \$100 with (b) the certainty equivalent utility of \$0. We report X in the table in dollars.

Table 6: Premium-setting equation

	(1)	**(2)**	**(3)**	**(4)**
Medical costs	0.745	0.737	0.733	0.547
	(0.026)	(0.027)	(0.024)	(0.054)
Administrative costs (t-1)				0.617
				(0.166)
Year FEs				✓
Payer-Year FEs		✓	✓	
N	240	238	238	186
ψ				
1^{st}-stage F-stat			18.329	4.948
R^2	0.976	0.983	0.983	0.975

Note: This table contains the estimates of our premium setting model across payer-metal tier insurance offerings. A market is defined as a calendar year and rating area combination. The model's predicted total monthly premiums and costs are in 100s of dollars. Column 1 presents an ordinary least squares (OLS) regression of premium on cost. Column 2 presents an OLS regression of premium on cost and includes payer-year fixed effects. Column 3 presents a two-stage least squares (2SLS) regression where the instrumented variable is a plan's predicted cost and the instruments are the number of plans in the same market and the fraction of households in the market who are subsidized. Column 3 also includes payer-year fixed effects. Column 4 presents a 2SLS where the instrumented variables are a plan's predicted cost as well as its predicted administrative costs from the previous year and the instruments are: the number of plans in the same market; the fraction of households in the market who are subsidized; the prior two variables for the same plan in the previous year; the total number of subscribers of the carrier's plans in other states; and, the total number of subscribers of the carrier's plans in the individual market in other states. Column 4 also includes year fixed effects. For details on the construction of a plan's predicted administrative costs, see Appendix B.

Table 7: Counterfactual Results: Outcomes

Counterfactual:	Base	SG closes	Extended ICHRA	Forced Insurance
Panel A: Individual Market (N = 178,157)				
Welfare				
Consumer surplus, change from base	0	18	28	44
Government expenditure	67	60	57	52
Market Shares				
Uninsurance	0.57	0.54	0.53	0.51
Bronze	0.17	0.16	0.16	0.16
Silver	0.20	0.20	0.20	0.20
Gold	0.07	0.09	0.11	0.14
Panel B: Small Group Market (Base: N = 55,374; Merged: N = 125,527)				
Welfare				
Consumer surplus, change from base	0	218	263	57
Government expenditure	124	71	42	75
Employer expenditure	278	0	119	238
Market Shares				
Uninsurance	0.00	0.61	0.54	0.00
Bronze	0.14	0.10	0.14	0.28
Silver	0.49	0.18	0.17	0.35
Gold/Platinum	0.25	0.11	0.15	0.36

Note: This table shows the effects of merging the individual and the small group markets in 2016 under different counterfactual scenarios. All reported numbers are averages over households and are reported at the monthly level. Small group "base" numbers are calculated using observed choices and premiums. All other numbers are predicted using the counterfactual algorithm. To account for outliers, consumer surplus is reported for the set of households where $\alpha - \psi > 0.05$. In the small group market "base" column, we omit households who choose grandfathered or other plans for which we do not observe premiums or metal tier in the data. In the alternative counterfactuals, we can simulate premiums and metal tiers for all households, including those omitted in the base category. We label this sample count the "merged N".

Table 8: Counterfactual Results: Premiums

Counterfactual:	Base	SG closes	Extended ICHRA	Forced Insurance
Panel A: Individual Market (N = 178,157)				
Standardized Gross Premiums				
Bronze	177	175	172	163
Silver	219	205	208	195
Gold	685	458	354	276
Population Net Premiums				
Bronze	226	227	224	215
Silver	294	277	287	275
Gold	1,379	838	631	485
Panel B: Small Group Market (Base: N = 55,374; Merged: N = 125,527)				
Standardized Gross Premiums				
Bronze	250	173	170	161
Silver	312	199	200	186
Gold	392	442	320	245
Population Net Premiums				
Bronze	111	162	1	0
Silver	136	197	8	3
Gold	172	784	162	49

Note: This table shows the effects on premiums of merging the individual and the small group markets in 2016 under different counterfactual scenarios. All reported numbers are averages over households and are reported at the monthly level. Small group "base" numbers are calculated using observed choices and premiums. All other numbers are predicted using the counterfactual algorithm. In the small group market "base" column, we omit households who choose grandfathered or other plans for which we do not observe premiums or metal tier in the data. In the alternative counterfactuals, we can simulate premiums and metal tiers for all households, including those omitted in the base category. We label this sample count the "merged N". Gross premiums reflect standardized premiums for a 40-year old without subsidies for each plan in the household's choice set. We average over these standardized premiums weighting by plan choice probabilities. Population net premiums reflect the average premium faced by observed households accounting for age adjustments and subsidies and weighting by plan choice probabilities.

Table 9: Counterfactual Results: Varying Markups

Panel A: Individual Market (N = 178,157)							
Counterfactual:	Base	Extended ICHRA			Forced Insurance		
Markup:	0	0	10	25	0	10	25
Welfare							
Consumer surplus, change from base	0	28	13	-8	44	32	18
Government expenditure	67	57	61	67	52	55	59
Standardized Gross Premiums							
Bronze	177	172	181	195	163	171	181
Silver	219	208	227	258	195	211	229
Gold	685	354	449	618	276	329	412
Panel B: Small Group Market (Base: N = 55,374; Merged: N = 125,527)							
Welfare							
Consumer surplus, change from base	0	263	253	238	57	45	28
Government expenditure	124	42	51	64	75	84	98
Employer expenditure	278	119	122	126	238	248	255
Standardized Gross Premiums							
Bronze	250	170	178	191	161	168	176
Silver	312	200	215	238	186	197	210
Gold	392	320	410	568	245	287	354

Note: This table shows the effects of varying a fixed markup parameter on counterfactual outcomes. All reported numbers are averages over households and are reported at the monthly level. Small group "base" numbers are calculated using observed choices and premiums. All other numbers are predicted using the counterfactual algorithm. To account for outliers, consumer surplus is reported for the set of households where $\alpha - \psi > 0.05$. In the small group market "base" column, we omit households who choose grandfathered or other plans for which we do not observe premiums or metal tier in the data. In the alternative counterfactuals, we can simulate premiums and metal tiers for all households, including those omitted in the base category. We label this sample count the "merged N". Gross premiums reflect standardized premiums for a 40-year old without subsidies for each plan in the household's choice set. We average over these standardized premiums weighting by plan choice probabilities.

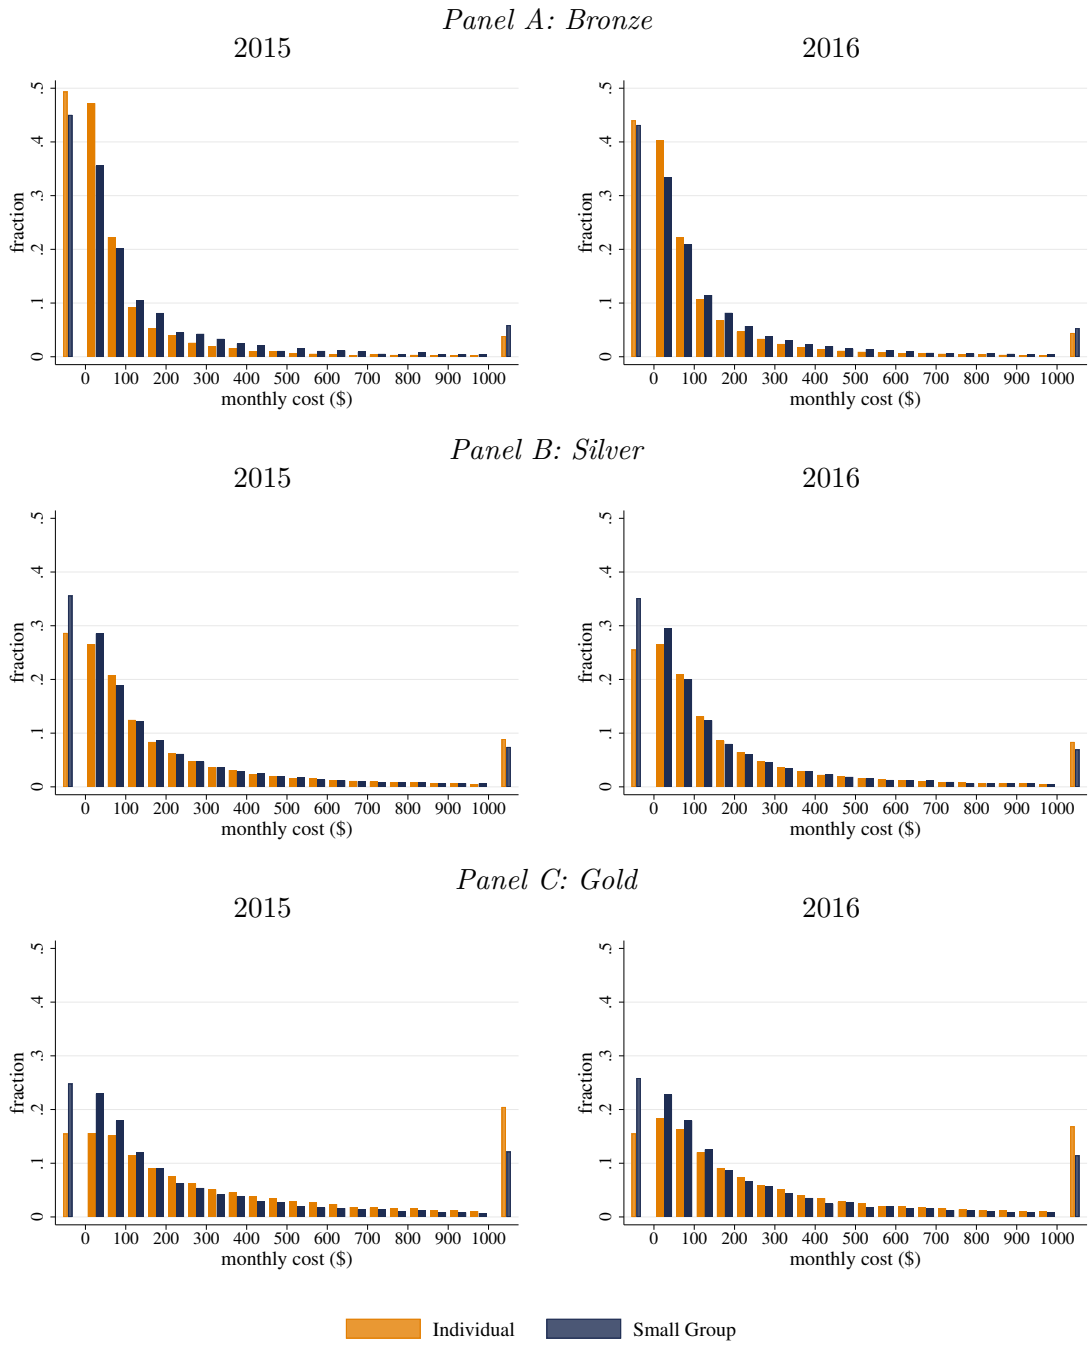


Figure 1: Distribution of monthly medical costs

Note: This figure depicts the distribution of monthly medical cost across households. Panels A, B, and C are of subscribers who purchased a bronze, silver, and gold plan respectively . The graphs on the left and right show the cost distributions for the year 2015 and 2016 respectively. For each histogram, a bar depicts the fraction of households who fall into that range of costs across the sample. The bars on the far left depict the fraction of households with zero monthly medical cost. The bars on the far right depict the fraction of households with an average of more than $1,000 of monthly medical costs. Interior bins have equal width of $50 and start from $1. The lighter bars reflect households who choose plans in the individual market while the darker bars bars reflect households in the small group market.

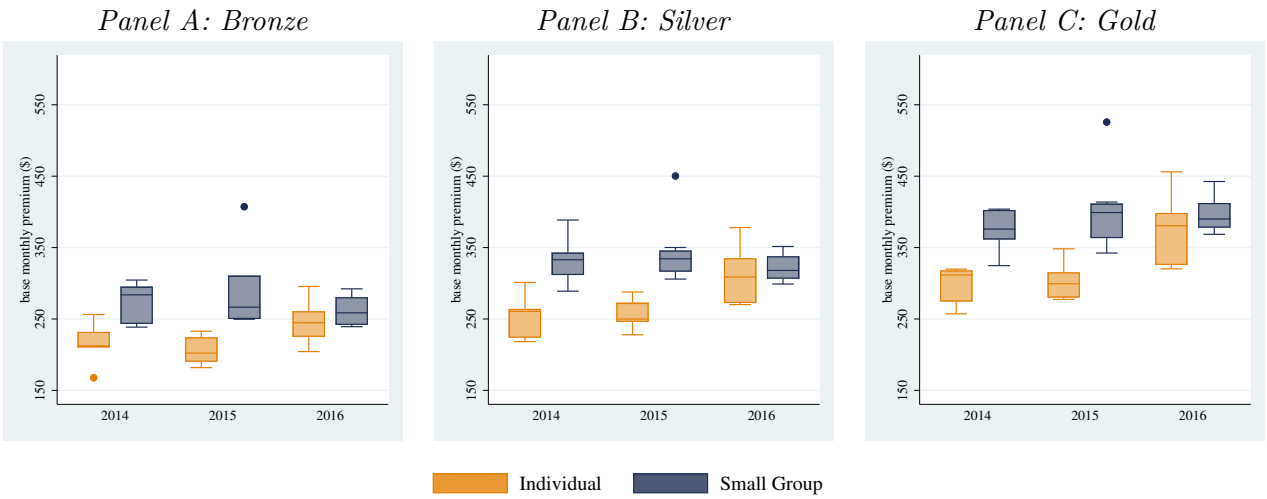
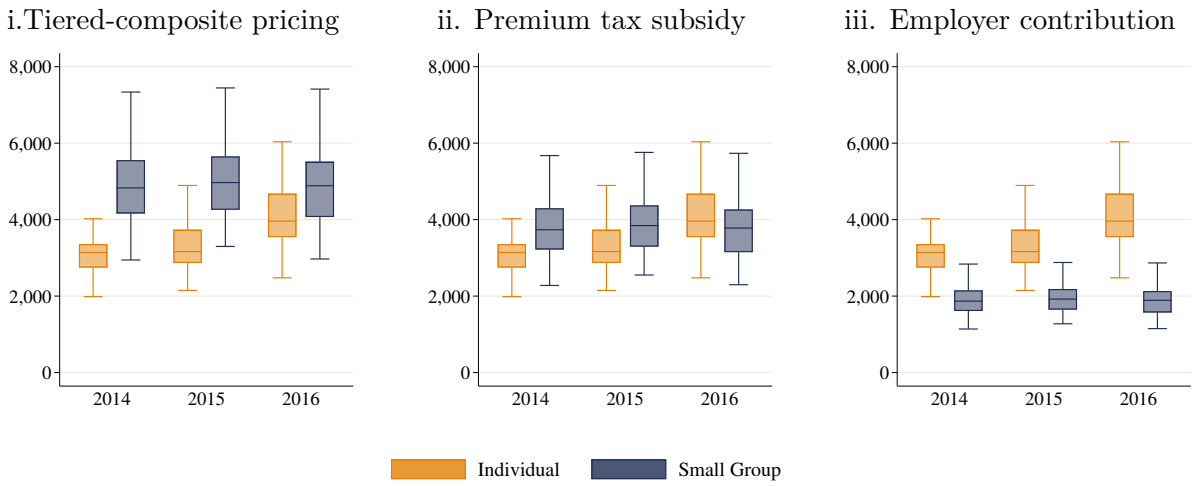


Figure 2: Base monthly premium

Note: This figure depicts the distribution of base premiums across payers. The figures from left to right show the distributions for bronze, silver, and gold plans respectively. For each panel, from left to right, each sub-panel is for the year 2014, 2015, and 2016 respectively. For each sub-panel, the box on the left is of plans in the individual market and the box on the right is of plans in the small group market. We calculate the base premium for a payer as the average premium for a non-smoking single 40-year-old adult, where we take the average across all plans that households purchase from the payer.

Panel A: Comparison of subsidy schemes



Panel B: Employer subsidy required to equate premiums across markets

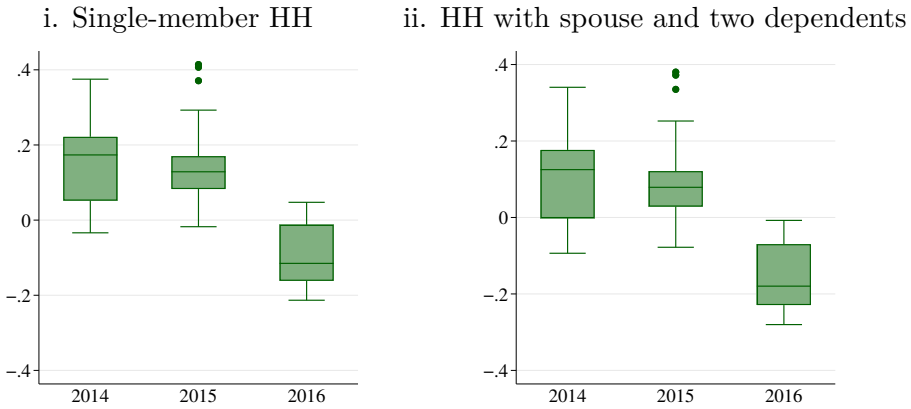


Figure 3: Distribution of subsidies

Note: This figure compares premiums in the individual and small group market after accounting for different subsidies households may receive in the small group market. For each constructed plan available in the small group market, we identify an identical constructed plan in the individual market. We then simulate for each plan the average premium that a 40-year-old single subscriber would pay in the two markets under different subsidy schemes. The plot on the left assumes that the only difference in pricing is in the tiered-composite pricing in the small group market. The plot in the middle combines tiered-composite pricing with the implicit premium tax subsidy in the small-group market. The plot on the right combines tiered-composite pricing, the premium tax subsidy, as well as the employer's premium contribution. We set the contribution at 50% for illustration. Panel B plots the employer subsidy (as a share of premiums) that is necessary to equate the individual market premium with the small group premium; the small group premium reflects tier-composite rating and tax subsidies. The figure on the left depicts the required subsidy share for a 40-year-old single subscriber. The figure on the right depicts the share for a household that includes a spouse and two dependents.

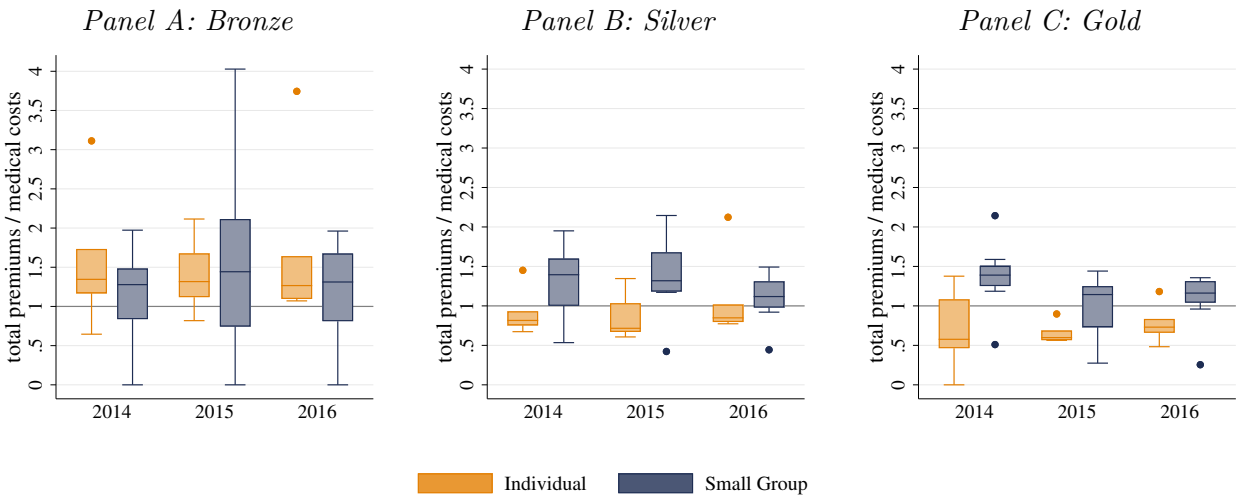


Figure 4: Medical markup (total premiums over medical costs)

Note: This figure depicts the distribution of base medical markups across payers. The figures from left to right show the distributions for bronze, silver, and gold plans respectively. For each panel, from left to right, each sub-panel is for the year 2014, 2015, and 2016 respectively. For each sub-panel, the box on the left is of plans in the individual market and the box on the right is of plans in the small group market. We calculate the base premium for a payer as the average premium for a non-smoking single 40-year-old adult, where we take the average across all plans that households purchase from the payer. We calculate the medical markup as the ratio of a payer's total premium revenue divided by the total medical cost the payer incurs. This cost does not account for risk adjustment payments or other transfers to the insurer.

Online Appendix for "Market Segmentation and Competition in Health Insurance"

Michael J. Dickstein, Kate Ho, Nathaniel Mark

October 18, 2021

Dickstein: New York University and NBER, michael.dickstein@nyu.edu; Ho: Princeton University, NBER, and CEPR, kate.ho@princeton.edu; Mark: U.S. Department of Justice, nathaniel.mark@usdoj.gov. The views expressed herein do not represent the views of the U.S. Department of Justice. Throughout this paper, we use the term "market" in ways that do not necessarily reflect the product and geographic boundaries of antitrust markets.

A Institutional Details

A.1 Individual market

Pricing constraints. In 2012, prior to the implementation of ACA regulations, only six states required insurers to issue individual insurance to any applicant and only seven states used a form of community rating to prohibit premium setting based on health status (The Kaiser Family Foundation, 2012). Oregon did not have a state-level 'guaranteed issue' rule prior to the ACA; only those households who met the requirements of a federal statute could be guaranteed coverage.[2] Carriers in Oregon could also price based on health status, though not as freely as in some states. Along with 10 other states and the District of Columbia, Oregon applied rate bands in the individual market prior to the ACA, prohibiting insurance companies from charging premiums beyond a specified share of the average premium in the market (The Kaiser Family Foundation, 2012).

The implementation of the ACA in 2014 harmonized formerly divergent state-based regulation of the individual market. Under current law, insurers in all states must guarantee issue all plans to all consumers.[3] In setting rates, carriers may only vary premiums with family size, state-defined geographic regions, tobacco use, and age. Further, premiums may only vary by age following a standard age curve, with a ratio of 3 to 1 from the oldest to youngest enrollee. We exploit the formulaic variation in premiums by age in demand analyses.

Plan design. Plans offered in the individual market after the ACA are also subject to regulation on both covered benefits and patient out-of-pocket costs. Under the ACA, all plans must cover a set of ten essential health benefit categories, including outpatient services, emergency room visits, pregnancy and maternity visits, mental health care, and prescription drugs. All plans offered fall into metal tiers classified by their actuarial value, the percentage of health costs the plan is expected to cover. The plan tiers are Bronze, Silver, Gold and Platinum, with actuarial values of 60%, 70%, 80% and 90%, respectively.

Oregon added additional regulation in the marketplace beyond the federal requirements on two dimensions: required tier offerings and plan standardization (Blumberg et al., 2013). First, while the federal regulations require insurers to sell at least one gold and one silver plan in the marketplace in each geographic market they enter, Oregon requires that all insurers entering the marketplace must offer a bronze, silver, and gold plan. If an insurer chooses to offer individual plans outside the marketplace, it must offer at least one bronze and one silver plan. Second, Oregon requires insurers to offer a standardized plan in each of the required metal tiers. The standardized plan features a specific cost-sharing and benefit design; under the broader federal regulations, insurers have flexibility to design plan copayments, deductibles, and other benefits within a metal tier so long as it achieves the specified actuarial value for the tier.

Households eligible for cost-sharing subsidies must purchase a silver plan in order to receive them. These subsidies shift the standard silver plan design to a more generous actuarial value of between 73% and 94% for consumers with incomes on 100% to 250% of the FPL, with the lowest incomes receiving the higher actuarial values.

[2] For example, under federal law, individuals leaving group coverage of at least 18 months duration could not be turned down for individual coverage provided they enrolled within 63 days of losing group coverage.

[3] 'Grandfathered' plans—that is, plans that existed before the ACA and currently in effect—were not subject to these new regulations.

A.2 Small group insurance market

Small employers, defined as firms with up to 50 full-time employees, have the option of offering health insurance coverage for their employees. In 2015, approximately 1 in 4 full-time employees worked for a small employer. The share of these employees covered by employer-provided insurance equaled 35% for workers at firms with 3-24 workers and 49% for firms between 25 and 50 workers. These rates combine both the likelihood of firms' offering coverage to classes of employees and employees taking up that coverage. 47% of firms with 3-9 workers and 68% of firms with 10-49 workers offered coverage to at least some of their workers. Within the set of small firms that offer insurance, approximately 76% of eligible workers take-up coverage (Claxton et al., 2015).

The small group market is subject to many of the same plan design restrictions as the individual market. Plans must cover the same essential health benefits, must be structured according to the same metal tiers, and must be "guaranteed issue". The small group and individual markets differ in the purchasing channel and the pricing rules.

Purchasing channel. As in the individual market, the ACA intended states to offer a marketplace for small group employers to shop for plans, known as the Small Business Health Options program or SHOP. During the span of our data, from 2014 through 2016 Oregon did not operate a SHOP marketplace; all small groups purchase coverage through agents.

Employees in a small group face two levels of agency: the employer and the insurance broker. A typical small group offers one to two broker-recommended plans to its employees, often from the same carrier (AHRQ, 2016). By contrast, consumers shopping in the individual market face a much larger choice set.

After choosing which plans to offer, employers contract with the relevant carrier(s) and pay premiums on behalf of the group. Employees pay their share of the group premium from their pre-tax earnings– that is, all premiums for insurance obtained through an employer are federal and state tax-exempt, regardless of whether the employer or the employee pays. This creates a tax wedge relative to the individual market, where households typically pay for insurance with post-tax dollars.[4] Employers also typically subsidize the cost of employee premiums, covering as much as two thirds of the premium cost (Claxton et al., 2015).[5]

Premium setting. Oregon requires insurers in the small group market to use a form of community rating known as 'tiered composite' rating. Employees at a firm do not face premiums proportional to their family size or the ages of members of their family. Instead, these ages and family sizes contribute to a group level premium cost which is then divided into four premium prices, based on the household size: (a) employee only, (b) employee and spouse, (c) employee and children, and (d) employee, spouse, and children.

To determine these four premium levels, the employer creates a list of all households who would be covered in the plan, including the family size and the ages of each household member.[6] With this

[4]Premiums in the individual market are part of itemized deductions, but subject to limitations: only medical expenses exceeding 7.5% of adjusted gross income are deductible.

[5]Small businesses with fewer than 25 employees are also eligible for tax credits of up to 50% of premium costs if they satisfy a number of qualifications, including: they must buy a plan certified for SHOP, average employee pay must be less than $50,000, and the employer must cover at least 50% of the premium (Oregon Health Insurance Marketplace, 2020).

[6]The prices also vary geographically based on the location of the employer. If an employer has multiple locations that span more than one pricing region in the state, employees may face different premiums by location.

list, and for a chosen plan design (e.g. a silver managed care plan), the insurer applies its baseline plan-specific premium to each employee and dependent according to the individual's age. Summing all of these premiums generates the total amount the employer pays the insurer for coverage under the chosen plan.

To determine the four premiums that households will face (before any employer subsidy), the insurer needs two measures. The first is the total age-based premiums in the group. The second is a sum of 'rating factors'; state regulators fix the factors to a specific level by family type. The rating factor equals 1 for employee-only, 2 for employee and spouse, 1.85 for employee and children, and 2.85 for employee, spouse, and children. The rating factors do not vary with the number of children. With these measures, the insurer calculates the average employee premium by dividing the total age-based premiums by the total rating factors. The four actual employer premium levels equal this average employee premium multiplied by the respective rating factor for each household size category.

This tiered composite system creates a cross-subsidy within the employer pool between older and younger enrollees and between employees covering only themselves and those covering families. The extent of the subsidy depends on how well the fixed rating factors match the differential in premium costs the families contribute to the small group pool.

B Dataset Construction

B.1 Household dataset creation

Our primary data come from Oregon Health Authority's All Payer All Claims (APAC) dataset for the years 2014-2016. These data include all medical claims, drug claims, and insurance plan enrollment records for commercially insured individuals in the state of Oregon. The medical and drug claim data include financial variables, diagnoses, and dates of claims. The insurance plan enrollment records are at the individual – month level and include details of both medical and drug insurance plans. Before 2016, reporting of all claims and enrollment records were mandatory by law for all commercial insurers. In 2016, the US Supreme Court, in *Gobeille v. Liberty Mutual Insurance Company*, created an exemption for reporting for self-insured plans. This exemption does not affect our research, as we focus on the fully-insured individual and small group insurance markets.

The sample we use in the analysis is a subset of the APAC dataset. The original insurance plan enrollment data include approximately 50 million observations per year (49,952,770, 48,957,463, and 38,364,318 in 2014, 2015, and 2016 respectively). We keep only unique medical plan observations (leaving 27,882,303, 27,037,082, and 20,034,911 observations). Next, we collapse the data to the individual – year observation level from the individual – month observation level (2,160,301, 2,035,972, 1,643,776).

We further restrict the individual - year observation level dataset to reflect the goals of our analysis. We restrict the data to large carriers to enable us to solve for equilibrium insurer behavior in our empirical framework. We define a large carrier as a carrier with more than 5% of individuals in any of the years in the individual market or small group market. Eight carriers qualify as large. This removes 3.3% and 2.1% of individuals in the individual and small group markets. We further restrict our sample to include only individuals in the individual and small group markets, leaving 506,356, 466,320, and 494,539 individuals in 2014, 2015, and 2016. Lastly we collapse the dataset

to the household – year observation level from the individual – year observation level. After these changes, the sample includes 312,122 households in 2014, 294,745 households in 2015, and 320,273 households in 2016.

B.2 Switcher dataset creation

To estimate the preferences of small group market households, we use a sample of small group households whose plans were closed in 2014 or 2015. We then track these households as they choose plans in the individual market, under the assumption that we can estimate small group consumers' preferences, conditional on observables, in this setting. The APAC data allow us to conduct this exercise, as they include an identifier that allows us to track individuals across plans, markets, and years. We exploit this feature of the data to determine which plans close, which households move to the individual market in the following year, and which plans attract market share in the individual market.

We identify households whose plans were closed by tracking the members of employer-health plans in the following year. We define an employer-plan as a closed plan if 80% or more of its members are either uninsured, in a plan in the individual market, or obtain coverage through a spouse or as a dependent in the following year. We identify 6,779 closed employer-plans. These plans had a total of 16,631 enrollees.

Table A1 describes the destinations of these enrollees in the year after their plan closed. 5,321 choose plans in the individual market. 1,207 obtain coverage through a spouse or as a dependent, 920 obtain coverage in other markets, including group coverage or Medicare coverage.[7] 9,183 enrollees do not obtain coverage in any of the above markets. These enrollees may be uninsured, may be covered under Medicaid, may have moved out of Oregon, or may be untracked for some other reason. We assume that these enrollees choose the outside option of uninsurance.

B.3 Variable construction

We collect household characteristics, costs, and plan choices from the claims data and the insurance plan enrollment data. We collect insurer administrative costs from public medical loss ratio data.

Rating area. Rating areas are sets of counties. We assign rating areas to individuals using zipcodes from claims data and city names from the plan enrollment data.

First, we assume that all members of a household live in the same location. We assign the mode of zipcodes observed in a household to every member of that household. We also assign the mode of city names in a household to every member of that household. Second, we assign rating areas in the following order:

An individual's rating area is assumed to be the rating area where their city is located. We allow the individual's city name to approximately match with the city name in an external dataset according to the minimal optimal string alignment distance with a maximum distance of 1. The external database of city names and their rating areas was partially manually created and partially obtained from publicly available census datasets. Some individuals are not assigned a rating area based on their city name due to missing city names or incorrectly spelled city names. We assign these individuals to the rating area that is associated with their zipcode. If the zipcode falls within

[7] We assign Medicare coverage for an enrollee who turns 65 in the plan year.

two rating areas, we choose the rating area with the largest share of zipcode population. We use the HUD-USPS zipcode-county crosswalk from 2010 census data for this matching.

Employer-health plan identifier. For the small group market, we construct an employer-health plan identifier. For most insurance carriers, this identifier corresponds to a variable in the data named contract number. For two carriers, there are multiple health plans associated with each contract number. For these carriers, we set the employer-health plan identifier equal to a contract number-health plan index.

Income. While we do not observe household income, we develop a predictive model to approximate income for enrollees in the individual and small group markets. Our goal is to use this measure of income to identify a household's eligibility for government premium and cost-sharing subsidies in the individual market.

For a subset of households in our APAC data, we observe the household's net premium– that is, premium net of subsidies.[8] Because government premium subsidies scale with income, we can compute the household's income as a function of demographics (age and family composition) and the observed net premiums. We recover a predicted measure of income for 73,440 households with observed net premiums.

Using the recovered incomes, we estimate a predictive Tobit model of income. In the model, incomes y_{imt} follow a normal distribution, conditional on household characteristics x_{imt}:

$$\begin{aligned} y_{imt} &= \beta' x_{imt} + u_{imt} \\ u_{imt} &\sim N(0, \sigma^2) \end{aligned} \tag{B.1}$$

However, the econometrician observes a truncated version of income:

$$y^*_{imt} = min\{y_{imt}, y^{max}_{imt}\}$$

The upper bound on the predicted measure of income, y^{max}_{imt}, is the lowest income that would make premium subsidies equal to zero. In almost all cases, $y^{max}_{imt} = 4$. We estimate the model parameters β, σ^2 via maximum likelihood. We use household characteristics including family size, city of residence, year, age, household health status score (both mean and max within family), whether the household is insured for all 12 months of the year, and the 10th and 50th income percentile of the household's zipcode.[9] We also include non-linear transformations of these characteristics. If households have missing age or gender, we estimate income using a Tobit model with these variables omitted.

We use the estimated model to predict the income of households in both the individual and small group markets. Using the estimated income of each household, we determine the household's federal and state average tax rates using the NBER's Internet TAXSIM version 27. We account for both marital status and the number of dependents when estimating average tax rates. We assume that married households file taxes jointly and that households have zero tax deductions.

[8] We confirm that the premiums we observe for these households are inclusive of subsidies by comparing the reported premiums to public data from SERFF that contain gross premiums by plan and household composition.

[9] When zipcode is unknown, we match percentiles of income using the household's city of residence.

Metal tier. Metal tiers and health plans are not consistently reported by all insurance carriers in the enrollment data. When they are missing, we estimate the metal tier of a household's chosen plan using their contract number, premiums, zipcodes, and external data on the number of subscribers in each metal tier and year by insurer and zipcode from Oregon's Department of Consumer and Business Services, Division of Financial Regulation.

Our determination of metal tier follows three steps. First, we set the household's metal tier to equal the metal tier of their health plan when available. Second, we use the metal tier reported by the insurer when it is available. Third, in the individual market, we set all remaining household's metal tiers to equal the mode of assigned metal tiers in their contract, if the household is in a contract that is likely to have consistent metal tiers for all of its members. We define these contracts as those in carrier-markets where more than 95% of households have the metal tier that is most common in their contract. For data from most insurance carriers, these three steps generate numbers that match with external counts of metal tiers from Oregon's Division of Financial Regulation. For these carriers, we assume all other households to be in grandfathered plans without ACA-regulated actuarial values. For other carriers with less well populated metal tier identifiers, we use claims data, observed premiums, and data from Oregon's Division of Financial Regulation (DFR) to deduce metal tiers by plan.

Medical spending. To handle outliers, we winsorize medical spending above the 99th percentile of individual market medical spending ($8,937 per month). Here, we define medical spending as all medical costs covered under the insurance plan, omitting patient out-of-pocket expenses. Additionally, when we report statistics relating to medical costs, including expected non-discretionary spending ($E[\lambda]$), we winsorize using the same upper bound.

Administrative cost. We define administrative cost measures using publicly available data collected by the Centers for Medicare and Medicaid Services (CMS). CMS requires insurers to report administrative costs annually at the state-market level as an input into regulation of insurers' medical loss ratios. We collect this insurer-state level data for the individual market for the years 2014-2016.

From the CMS data, we define total per-enrollee monthly administrative costs as the sum of the following constituent elements of administrative costs: (1) taxes and fees, from Section 3 of the CMS form, which include federal and state taxes deductible from premiums, community benefit expenditures, contributions to the Federal Transition Reinsurance Program, and various other regulatory fees; (2) wellness activities, from Section 4 of the CMS form, which include activities to improve and promote health outcomes and prevent hospital readmission as well as expenses related to health information technology; and, (3) general administrative expenses, from Section 5 of the CMS form, which include direct sales salaries and benefits, agents and broker fees, and all other non-medical general and administrative expenses. We employ the total administrative cost measure in our empirical analyses reported in Table 6.

We report summary statistics of these cost variables in Table A2. In addition, in the table we report other measures available in the CMS data. We compute statistics on the count of annual subscribers in Oregon that each insurer covers in each possible business line. Finally, we present statistics on the number of states, including Oregon, where payers that operate in Oregon are active. We classify an insurer as active in a state when it files a report with CMS in that state and year.

B.4 Uninsured population

Uninsured households are not observed in our data. We use the American Community Survey (ACS) to infer the number of uninsured households by county and age and we use the California Health Interview Survey (CHIS) to determine the likely characteristics of uninsured households. The ACS includes estimates of the uninsured populations by age, sex, and geography groups. The CHIS provides measures of the joint distribution of characteristics of the California uninsured population.

Data. To infer the size of the uninsured population, we use the 2014-2016 American Community Survey 1-year Estimates Health Insurance Coverage Status by Sex by Age datasets from the U.S. Census Bureau. These data include estimates of the number of civilian noninstitutionalized persons in each age and sex group that report "No health insurance" and "With health insurance." The estimates are available at the county level if county populations are above 65,000. Estimates are also available at the Public Use Microdata Area (PUMA) level. Counties with populations below 65,000 are proper subsets of PUMAs.

We chose to use the mixture of 1-year county level and 1-year PUMA level data over using only 5-year county data. Using 5-year county data offers improved precision of estimates, especially for small counties. However, the 5-year ACS relies on a moving average of 5 years of survey results. Therefore, its estimates are biased upwards by higher uninsurance rates in pre-2014 years.

To infer the characteristics of the uninsured population, we use the 2014-2016 California Health Interview Survey. The survey provides data on the joint distribution of family size and type, income, urban share, and health conditions for uninsured households. Unlike the ACS, the CHIS observations are at the household level. In using CHIS, we assume that the empirical distribution of characteristics observed in the CHIS is similar to the distribution of characteristics in Oregon. In 2016, California and Oregon had similar uninsurance rates (7% and 6%, respectively). This similarity remains in cross sections of the income distribution. Within the subset of incomes below 200% of the Federal Poverty Line, California had a non-elderly uninsurance rate of 13%. Oregon had a non-elderly uninsurance rate of 11% (The Kaiser Family Foundation, 2017).

Procedure. Using the ACS, we estimate the uninsured populations in each age group and rating area. Since rating areas are a collection of counties, we conduct these estimates at the county level and aggregate to the rating area. If a county has a small population, the ACS does not include county-level data. For these counties, we infer county age group uninsured populations using PUMA-level ACS data. Specifically, a small county's uninsured population for an age group is calculated as the percent of the PUMA population that is in that county multiplied by the uninsured population in the PUMA for the age group.[10] Out of 36 counties, 21 had populations below 65,000 and had to be estimated.

We use the CHIS to map the number of individuals who are uninsured to the number of households who are uninsured. For each year, age group, and rating area, we estimate the expected number of adults in a household using the CHIS data. We estimate the number of uninsured households by dividing the number of uninsured individuals by the expected number of adults in a household.

[10]For greater accuracy, we use population data from the American Community Survey 5-year estimates.

Next, we distribute characteristics to the uninsured households. First, uninsured households are distributed to bins according to the empirical joint distribution of characteristics from the CHIS, conditional on metropolitan area status of the rating area, year, and age group. Bins are defined by rating area, year, age group (< 17, 17-23, 24-33, 34-43, 44-53, 55-64, ≥ 65), income (< 2.5, ≥ 2.5), marital status, and whether the household survey respondent has health conditions.[11]

Lastly, we draw specific values of characteristics for each uninsured household randomly from the insured population of households in the same bin. Here, we assume that the distributions of characteristics are similar between the insured and uninsured populations once we condition on metropolitan area status, age, income, marital status, and health conditions.

C Estimation Details

C.1 Further details on deriving equations for estimation

As noted in Section 6.2.1, the assumption that $\lambda \sim exp(\alpha)$ implies

$$u^*_{j,t} = \frac{1}{2}x^2\omega\lambda - (1-x)\lambda + y_t - p_{t,j} + g(X_{j,t}, \epsilon)$$

Now the expected utility over the distribution of λ is:

$$v_{j,t}(F_{\lambda,t}, \omega, \psi) = -\int exp(-\psi u^*_{j,t}) dF_{\lambda,t}(\lambda).$$

If $\lambda \sim exp(\alpha)$ so that $E(\lambda) = 1/\alpha$, we can apply the properties of the exponential distribution along with the monotonic transformation $-\frac{1}{\psi}ln(-v_{j,t})$ to find the order preserving utility function[12]:

$$U_{j,t} = y_t - p_{t,j} + \frac{1}{\psi}ln\big[\frac{\alpha - \psi(1-x) + \psi\frac{1}{2}x^2\omega}{\alpha}\big] + g(X_{j,t}, \epsilon)$$

.

Comparing this utility to the utility from the outside option

$$U_{0,t} = y_t + \frac{1}{\psi}ln\big[\frac{\alpha - \psi}{\alpha}\big] + g_0(\epsilon_{0,t}),$$

we obtain the utility of the inside goods relative to the (non-stochastic component of the) outside option as

$$U_{j,t} = -p_{t,j} + \frac{1}{\psi}ln\big[1 + \frac{\psi x}{\alpha - \psi} + \frac{\psi}{2}\big(\frac{x^2\omega}{\alpha - \psi}\big)\big] + g(X_{j,t}, \epsilon)$$

[11] Health conditions used are Asthma, Diabetes, High Blood Pressure, and Heart Disease. A rating area is defined as metropolitan if most of its counties are designated metropolitan counties by the U.S. Census Bureau. These are rating areas 1, 2, 3, and 7. The ACS reports 0 uninsurance in some metropolitan area-year-age-income-marital status-family status-health status groups. We found this to be unreasonable and replaced those uninsurance estimates with: first, the average of uninsurance rates across other years in the same bin, holding everything else stable (139); second, if the first method is unavailable, the average of other uninsurance rates across other metropolitan status, holding everything else stable (114).

[12] Marone (2020) describes this step as "estimating demand in certainty equivalent units". In our setting, when $X \sim exp(\alpha)$, $kX \sim exp(\alpha/k)$. Further, $exp(kX) \sim Pareto(1, \alpha/k)$ so that $\int exp(kX) df_X = E(exp(kX)) = \frac{\alpha}{\alpha - k}$ provided $\alpha > k$.

and

$$U_{0,t} = g_0(\epsilon_{0,t}).$$

Recognizing that when $Ax + Bx^2$ is close to zero, we can approximate $ln(1 + Ax + Bx^2) \approx Ax + Bx^2$., we write our utility expression as:

$$U_{j,t} \approx -p_{t,j} + \frac{x}{\alpha - \psi} + \frac{x^2\omega}{2(\alpha - \psi)} + g(X_{j,t}, \epsilon).$$

We specify $g(X_{j,t}, \epsilon) = (\beta_0 X_{j,t} + \epsilon_{j,t})/(\alpha - \psi)$ so that sicker or more risk aversion consumers put more weight on plan characteristics like carrier identity, in the same way that they put more weight on coverage. Making an analogous assumption for the outside option, we find:

$$U_{j,t} \approx -p_{t,j} + \frac{x}{\alpha - \psi} + \frac{x^2\omega}{2(\alpha - \psi)} + \frac{\beta_0 X_{j,t} + \epsilon_{j,t}}{\alpha - \psi} \tag{C.1}$$

$$U_{0,t} = \frac{\epsilon_{0,t}}{\alpha - \psi}$$

C.2 Likelihood derivation

Section 6.2.2 provides a joint likelihood for the household's plan choice and its health spending that can be rewritten as

$$\mathcal{L}(\theta) = \prod_{i=1}^{N} \left\{ \left(\frac{1}{\sum_{k=1}^{J} exp(V_{i,k})} \right)^{d_{i,1}} \prod_{j=2}^{J} \left[\left(\frac{exp(V_{i,j})}{\sum_{k=1}^{J} exp(V_{i,k})} \right) \left(1 - exp\left(\frac{-\alpha_i \underline{c}}{x_j + \omega_i x_j^2} \right) \right)^{\mathbf{1}\{c_{ij} \leq \underline{c}\}} \right.\right.$$
$$\left.\left. * \left[\left(\frac{\alpha_i}{x_j + \omega_i x_j^2} \right) exp\left(-c_{i,j} * \frac{\alpha_i}{x_j + \omega_i x_j^2} \right) \right]^{\mathbf{1}\{c_{ij} > \underline{c}\}} \right]^{d_{i,j}} \right\} \tag{C.2}$$

With this notation, we write the log-likelihood:

$$L(\theta) = \sum_{i=1}^{N} \sum_{j=1}^{J} d_{i,j} ln \left(\frac{exp(V_{i,j})}{\sum_{k=1}^{J} exp(V_{i,k})} \right) + d_{i,j} \mathbf{1}\{j \neq 1\} \mathbf{1}\{c_{i,j} \leq \underline{c}\} ln \left(1 - exp\left(\frac{-\alpha_i \underline{c}}{x_j + \omega_i x_j^2} \right) \right)$$
$$+ d_{i,j} \mathbf{1}\{j \neq 1\} \mathbf{1}\{c_{i,j} > \underline{c}\} \left[ln \left(\frac{\alpha_i}{x_j + \omega_i x_j^2} \right) - c_{i,j} * \left(\frac{\alpha_i}{x_j + \omega_i x_j^2} \right) \right] \tag{C.3}$$

Given the exponential distribution, we need to constrain $\frac{\alpha_i}{x_j+\omega_i x_j^2} > 0$ as well. The score function of our log likelihood is:

$$\frac{\partial L(\theta)}{\partial \theta} = \sum_{i=1}^{N}\sum_{j=1}^{J} d_{i,j}\left[\frac{\partial V_{i,j}}{\partial \theta} - \frac{1}{\sum_{k=1}^{J} exp(V_{i,k})}\sum_{k=1}^{J} exp(V_{i,k})\frac{\partial V_{i,k}}{\partial \theta}\right] +$$
$$d_{i,j}\mathbf{1}\{j\neq 1\}\mathbf{1}\{c_{i,j}\leq \underline{c}\}\left[\left(\frac{-exp\left(\frac{-\alpha_i \underline{c}}{x_j+\omega_i x_j^2}\right)}{1-exp\left(\frac{-\alpha_i \underline{c}}{x_j+\omega_i x_j^2}\right)}\right)\left(\left(\frac{-\underline{c}}{x_j+\omega_i x_j^2}\right)\frac{\partial \alpha_i}{\partial \theta} + \left(\frac{\alpha_i \underline{c} x_j^2}{(x_j+\omega_i x_j^2)^2}\right)\frac{\partial \omega_i}{\partial \theta}\right)\right] +$$
$$d_{i,j}\mathbf{1}\{j\neq 1\}\mathbf{1}\{c_{i,j}> \underline{c}\}\left[\frac{\partial \alpha_i}{\partial \theta}\frac{1}{\alpha_i} - \frac{\partial \omega_i}{\partial \theta}\frac{x_j^2}{x_j+\omega_i x_j^2} - \frac{c_{i,j}}{x_j+\omega_i x_j^2}\left(\frac{\partial \alpha_i}{\partial \theta} - \frac{\alpha_i x_j^2}{x_j+\omega_i x_j^2}\frac{\partial \omega_i}{\partial \theta}\right)\right] \tag{C.4}$$

C.3 Specification

We further parameterize $\alpha_i, \omega_i, \psi_i$ as a function of household and/or plan level observables:[13]

$$\begin{aligned} ln(\alpha_i) &= W_{1,i}\beta_1 \\ ln(\omega_i) &= W_{2,i}\beta_2 \\ ln(\psi_i) &= W_{3,i}\beta_3 \end{aligned} \tag{C.5}$$

We simplify notation by defining the following vector:

$$\theta = \begin{pmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \\ \beta_3 \end{pmatrix}$$

Under this specification, the score function has the form

$$\frac{\partial L(\theta)}{\partial \theta} = \sum_{i=1}^{N}\sum_{j=1}^{J} d_{i,j}\left[\frac{\partial V_{i,j}}{\partial \theta} - \frac{1}{\sum_{k=1}^{J} exp(V_{i,k})}\sum_{k=1}^{J} exp(V_{i,k})\frac{\partial V_{i,k}}{\partial \theta}\right] +$$
$$d_{i,j}\mathbf{1}\{j\neq 1\}\mathbf{1}\{c_{i,j}\leq \underline{c}\}\left[\left(\frac{-exp\left(\frac{-\alpha_i \underline{c}}{x_j+\omega_i x_j^2}\right)}{1-exp\left(\frac{-\alpha_i \underline{c}}{x_j+\omega_i x_j^2}\right)}\right)\left(\left(\frac{-\alpha_i \underline{c}}{x_j+\omega_i x_j^2}\right)\begin{pmatrix}0\\W_{1,i}\\0\\0\end{pmatrix} + \left(\frac{\omega_i \alpha_i \underline{c} x_j^2}{(x_j+\omega_i x_j^2)^2}\right)\begin{pmatrix}0\\0\\W_{2,i}\\0\end{pmatrix}\right)\right] +$$
$$d_{i,j}\mathbf{1}\{j\neq 1\}\mathbf{1}\{c_{i,j}> \underline{c}\}\left[\begin{pmatrix}0\\W_{1,i}\\0\\0\end{pmatrix} - \frac{x_j^2}{x_j+\omega_i x_j^2}\begin{pmatrix}0\\0\\W_{2,i}\\0\end{pmatrix} - \frac{\alpha_i c_{i,j}}{x_j+\omega_i x_j^2}\left(\begin{pmatrix}0\\W_{1,i}\\0\\0\end{pmatrix} - \frac{\omega_i x_j^2}{x_j+\omega_i x_j^2}\begin{pmatrix}0\\0\\W_{2,i}\\0\end{pmatrix}\right)\right] \tag{C.6}$$

where

[13] Here, β_1 is K_1 x 1, β_2 is K_2 x 1, β_3 is K_3 x 1. x_j is a scalar.

$$\frac{\partial V_{i,k}}{\partial \theta} = \begin{pmatrix} X_k \\ -exp(W_{1,i}\beta_1)p_{i,k} \cdot W_{1,i} \\ \frac{1}{2}exp(W_{2,i}\beta_2)x_k^2 \cdot W_{2,i} \\ exp(W_{3,i}\beta_3)p_{i,k} \cdot W_{3,i} \end{pmatrix} \tag{C.7}$$

D Consumer Surplus

D.1 Derivation

In Section 7.3 in the main text, we define the certainty equivalent utility, e_{ijt}, for household i, plan choice j, and year t. Here we provide more details of its derivation. To simplify notation, we suppress market m subscripts in our description.

As in (Einav et al., 2013), we define e_{ijt} such that $-exp(-\psi_i e_{ijt}) = v_{ijt}$. Thus, $e_{ijt} = -\frac{1}{\psi_i}log(-v_{ijt})$. Here, v_{ijt} is the expected utility of choice j. In Section 5.1, we derived a form for v_{ijt} by combining our chosen functional form for utility with the assumption that the household's underlying health care need, λ_i, follows an exponential distribution:

$$v_{ijt}(F_{\lambda,t}, \omega_i, \psi_i) = -\int exp(-\psi_i \times u^*_{ijt}) dF_{\lambda,t}(\lambda) \tag{D.1}$$

$$= -exp\left[-\psi_i(y_t - p_{jt} + g(X_{jt}, \epsilon))\right]\left(\frac{\alpha_i}{\alpha_i + \psi_i(\frac{1}{2}x_{jt}^2\omega_i - (1 - x_{jt}))}\right) \tag{D.2}$$

where ψ_i is the household's coefficient of absolute risk aversion. Substituting v_{ijt} into our definition of e_{ijt}, we find:

$$e_{ijt} = -\frac{1}{\psi_i}log\left[exp\left(-\psi_i\left[y_t - p_{jt} + g(X_{jt}, \epsilon)\right]\right)\left(\frac{\alpha_i}{\alpha_i + \frac{\psi_i}{2}x_{jt}^2\omega_i - \psi_i(1 - x_{jt})}\right)\right] \tag{D.3}$$

$$= y_t - p_{jt} + g(X_{jt}, \epsilon) + \frac{1}{\psi_i}log\left(\frac{\alpha_i + \frac{\psi_i}{2}x_{jt}^2\omega_i - \psi_i(1 - x_{jt})}{\alpha_i}\right) \tag{D.4}$$

Similarly, the certainty equivalent utility for the outside option under the same approach equals:

$$e_{i0t} = y_t + g(X_{0t}, \epsilon) + \frac{1}{\psi_i}log\left(\frac{\alpha_i - \psi_i}{\alpha_i}\right) \tag{D.5}$$

We then normalize e_{ijt} by the deterministic component of the certainty equivalent utility for the outside option. Again, recognizing that when $Ax + Bx^2$ is close to zero, we can approximate $ln(1 + Ax + Bx^2) \approx Ax + Bx^2$, we find:

$$e_{ijt} \approx -p_{jt} + \frac{\omega_i}{2(\alpha_i - \psi_i)}x_{jt}^2 + \frac{1}{\alpha_i - \psi_i}x_{jt} + \frac{\beta_0 X_{jt}}{\alpha_i - \psi_i} + \frac{1}{\alpha_i - \psi_i}\epsilon_{jt} \tag{D.6}$$

$$e_{i0t} = \frac{\epsilon_{0t}}{\alpha_i - \psi_i} \tag{D.7}$$

We want to compute consumer surplus for each household i. In notation:

$$CS_{it} = E_\epsilon \left[\max_j e_{ijt} \right] \tag{D.8}$$

We substitute our expression for the certainty equivalent utility into CS_{it}. Multiplying and dividing by $\alpha_i - \psi_i$, we find:

$$\begin{aligned} CS_{it} &= E_\epsilon \left[\max_j e_{ijt} \right] && \text{(D.9)} \\ &= E_\epsilon \left[\frac{1}{\alpha_i - \psi_i} \max_j (\alpha_i - \psi_i) * e_{ijt} \right] && \text{(D.10)} \\ &= \frac{1}{\alpha_i - \psi_i} E_\epsilon \left[\max_j (\alpha_i - \psi_i) * e_{ijt} \right] && \text{(D.11)} \\ &= \frac{1}{\alpha_i - \psi_i} log \left(\sum_{j=0}^{J} \left[exp \left(-(\alpha_i - \psi_i) p_{jt} + \frac{\omega_i}{2} x_{jt}^2 + x_{jt} + \beta_0 x_{jt} \right) \right] \right) && \text{(D.12)} \end{aligned}$$

D.2 Consumer surplus in the small group market

To address the main caveat of our small group market consumer surplus measure – observed choices in the small group market do not fit the plan choices of small group employees forced to switch to the individual market – we assume a simple model of employer plan choice. In the model, the employer chooses a default silver-tier plan for its employees, but employees may opt for an alternative plan by paying a fee, $f_{i,j}$.[14]

If we define $V_{i,j}$ as the mean latent utility that household i obtains when $f_{i,j} = 0$, then the mean latent utility function for household i choosing plan j including the fee is:

$$V'_{i,j} = V_{i,j} - (\alpha_i - \psi_i) f_{i,j}$$

We recover α_i, ψ_i, and the remaining parameters of $V_{i,j}$ using our estimation of small-group preferences based on the forced switcher population. We then compare the predicted choices using these parameters to the observed choices in the small group market overall using a logit model. Specifically, we specify the fee as:

$$f_{i,j} = W_{i,j}^4 \beta_6$$

We recover the parameters of $f_{i,j}$ by fixing all other choice parameters at the switcher-sample estimates and maximizing the following likelihood:

$$L(\beta_6) = \sum_{i=1}^{N} \sum_{j=1}^{J} d_{i,j} ln \left(\frac{exp(V'_{i,j})}{1 + \sum_{k \in \mathcal{J}_i} exp(V'_{i,k})} \right) \tag{D.13}$$

The variables $W_{i,j}^4$ include interactions between plan metal tier and household characteristics such as age, health status score, household size, and employer-group size.

We report estimates of β_6 in Table A5. We estimate that fees are positive and large for non-silver plans, suggesting that employers induce their employees into silver plans at greater frequency

[14] Normalizing silver tier fees to be zero has consequences for our consumer surplus estimates. This assumption is informed by the empirical observation that 50% of small group employees are enrolled in silver plans.

than the estimated preferences would suggest. For example, using results from specification (1) in Table A5, households would have to pay $145 per month to receive bronze plan coverage. Our estimates also suggest that healthier households pay higher fees to choose an alternative plan. Older employees are estimated to pay higher fees for bronze plans, but not for higher coverage plans, suggesting that employers may encourage older employees into high coverage.

When accounting for the estimated fees using specification (5), our measure of small group consumer surplus drops significantly from our main estimate of -$86 to -$179.[15] In both of these measures, we assume that households do not have the option to opt out of insurance coverage. This generates negative estimates, as the revealed preferences of the small group switchers show that many households would prefer to be uninsured. Indeed, if we assume that uninsurance is an option for households, our consumer surplus measures increase markedly. Under a zero-fee assumption, consumer surplus increases from -$86 to $164; with fees, the surplus increases from -$179 to $134.

E Cost Censor

In the main specification, the spending cutoff is fixed. As a robustness exercise, we allow the data to recover the cutoff $\underline{c}$. Insurers may not submit a claim when a cost draw falls below the cost cutoff $\underline{c}$. The probability that an insurer does not submit a claim when cost is below the cutoff is $G(\underline{c}|x_j,\omega_i) = P(c_{i,j} = 0|x_j,\omega_i,c_{i,j} \leq \underline{c})$. Thus, the density of a cost observation is:

$$f(c_{i,j}|x_j,\omega_i,\alpha_i) = \begin{cases} 1 & x_j = 0, c_{i,j} = 0 \\ 0 & x_j = 0, c_{i,j} \neq 0 \\ \left[1 - exp(-\frac{\alpha_i \underline{c}}{x_j + \omega_i x_j^2})\right] G(\underline{c}|x_j,\omega_i) & x_j \neq 0, c_{i,j} = 0 \\ \frac{\alpha_i}{x_j + \omega_i x_j^2} exp\left(-c_{i,j}\frac{\alpha_i}{x_j + \omega_i x_j^2}\right) & x_j \neq 0, c_{i,j} \neq 0 \end{cases}$$

The probability that costs are not submitted, given a cost below the cutoff, $G(\underline{c}|x_j,\omega_i)$, is an unknown object. We estimate this object.

We assume that the probability of not submitting a claim is independent of moral hazard. The object is estimated non-parametrically, conditioning on only on actuarial value: $\hat{G}(\underline{c}|x)$.

For each potential actuarial value, we estimate $\hat{G}(\underline{c}|x)$ non-parametrically using Gaussian kernel methods. First, we define the number of cost observations associated with the actuarial value and the number of cost observations associated with the actuarial value for which insurers did not submit a claim.

$$N_x = \sum_{i=1}^{N}\sum_{j=1}^{J} d_{ij}\mathbf{1}\{x_j = x\}$$

$$N_x^0 = \sum_{i=1}^{N}\sum_{j=1}^{J} d_{ij}\mathbf{1}\{c_i = 0, x_j = x\}$$

[15] As we do when computing the main estimate, here we omit both outlier households (in terms of expected health spending) and households for whom consumer surplus is undefined when computing average surplus. Additionally, small group consumer surplus estimates are calculated using the premiums households face in the tiered pricing environment. Removing tiered pricing increases our estimate of consumer surplus from -$86 to -$78.

Second, the empirical probability and cumulative distributions of non-zero cost observations below the cutoff are estimated.

$$\hat{F}(z|x_j = x) = \frac{1}{N_x - N_x^0} \sum_{m=1}^{M} K(\frac{c_m - z}{h})$$

$$\hat{f}(z|x_j = x) = \frac{1}{(N_x - N_x^0)h} \sum_{m=1}^{M} k(\frac{c_m - z}{h})$$

The probability that claims are not submitted can be recovered from these objects.

$$\hat{G}(\underline{c}|x_j) = \frac{N_x^0}{N_x^0 + (N_x - N_x^0)\hat{F}(\underline{c}|x_j)}$$

$$\hat{g}(\underline{c}|x_j) = \frac{\partial \hat{G}(\underline{c}|x_j)}{\partial \underline{c}} = \frac{-N_x^0(N_x - N_x^0)\hat{f}(\underline{c}|x_j)}{[N_x^0 + (N_x - N_x^0)\hat{F}(\underline{c}|x_j)]^2}$$

The log likelihood has the form:

$$L(\theta, \underline{c}) = \sum_{i=1}^{N}\sum_{j=1}^{J} d_{i,j} ln\left(\frac{exp(V_{i,j})}{\sum_{k=1}^{J} exp(V_{i,k})}\right)$$
$$+d_{i,j}\mathbf{1}\{j \neq 1\}\mathbf{1}\{c_{i,j} = 0\}\left[ln(1 - exp(-\frac{\alpha_i \underline{c}}{x_j + \omega_i x_j^2})) + ln\left(\hat{G}(\underline{c}|x_j)\right)\right]$$
$$+d_{i,j}\mathbf{1}\{j \neq 1\}\mathbf{1}\{c_{i,j} \neq 0\}\left[ln\left(\frac{\alpha_i}{x_j + \omega_i x_j^2}\right) - c_{i,j} * \left(\frac{\alpha_i}{x_j + \omega_i x_j^2}\right)\right]$$

The gradient of this likelihood with respect to θ remains unchanged from the specification where $\underline{c}$ is assumed. The derivative of the likelihood with respect to $\underline{c}$ is:

$$\frac{\partial L(\theta, \underline{c})}{\partial \underline{c}} = \sum_{i=1}^{N}\sum_{j=1}^{J} d_{i,j}\mathbf{1}\{j \neq 1\}\mathbf{1}\{c_{i,j} = 0\}\left[\frac{\alpha_i}{x_j + \omega_i x_j^2}\frac{exp(-\frac{\alpha_i \underline{c}}{x_j + \omega_i x_j^2})}{1 - exp(-\frac{\alpha_i \underline{c}}{x_j + \omega_i x_j^2})} + \frac{\hat{g}(\underline{c}|x_j)}{\hat{G}(\underline{c}|x_j)}\right]$$

To estimate θ and $\underline{c}$, we use a penalized maximum likelihood approach. A standard maximum likelihood approach would overfit the data by setting $\underline{c}$ very high. That is, a high enough $\underline{c}$ ensures that the likelihood of observed *any* cost $c_{i,j}$ is 1. To avoid overfitting, we penalize the magnitude of $\underline{c}$, maximizing the penalized likelihood function:

$$PL(\theta, \underline{c}) = L(\theta, \underline{c}) - \Psi\left[log(\underline{c}\sigma\sqrt{2\pi}) + \frac{(log(\underline{c}) - \mu)^2}{2\sigma^2}\right]$$

Where the penalization parameter Ψ, and hyperparameters μ, σ^2 are specified before estimation. This method can be derived from the assumption that $\underline{c}$ has a log normal prior distribution.

$$\underline{c} \sim lognormal(\mu, \sigma^2)$$

Then, the likelihood of observing the data and $\underline{c}$ under the assumed prior is equal to the penalized likelihood function:

$$PL(\theta, \underline{c}) = log(P(\text{observe data}|\theta, \underline{c})) + \Psi log(P(\text{observe } \underline{c}|\mu, \sigma^2)$$

$$PL(\theta, \underline{c}) = L(\theta, \underline{c}) - \Psi\left[log(\underline{c}) + \frac{(log(\underline{c}) - \mu)^2}{2\sigma^2}\right]$$

F Counterfactual Algorithm

In this appendix section, we provide more detail on the algorithm we use to compute our counterfactual equilibrium. The approach mirrors that of Azevedo and Gottlieb (2017) for a competitive insurance market. We adjust the algorithm for the specific regulatory environment in Oregon and for our specification of expected utility.

We begin by collecting the parameters from both our supply and demand side estimation. From our maximum likelihood routine, which relies on both household plan choices and observed health care spending to identify the parameters of demand, we collect θ:

$$\theta = \begin{pmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \\ \beta_3 \end{pmatrix}$$

where the derived parameters of utility depend on θ under our specification:

$$\begin{aligned} ln(\alpha_i) &= W_{1,i}\beta_1 \\ ln(\omega_i) &= W_{2,i}\beta_2 \\ ln(\psi_i) &= W_{3,i}\beta_3 \end{aligned} \tag{F.1}$$

Here, β_0 is the parameter vector multiplying the plan-specific indicator variables in the utility specification.

We also collect (β_4, β_5), the parameters of our supply-side price-setting equation, Equation 6.6, where under perfect competition insurance carriers set total premium revenue equal to the sum of the portion of health care costs the insurer bears and the insurer's allocated administrative costs for operating the plan:

$$\widehat{R}_{jmt} = \sum_{i \in \{m,t\}}^{N_{mt}} \left(\widehat{s}_{ijmt} * \widehat{\beta}_4 \kappa_{ij} \bar{c}_{ijmt}\right) + \widehat{\beta}_5 A_{jmt} \tag{F.2}$$

In the equilibrium search, our goal is to find p_{jmt}, the plan-market-year baseline premium the insurer sets to equate revenue and costs in equilibrium for all plans and markets. We can translate the normalized baseline premium to (a) gross premiums for household i, p_{ijmt}, and (b) net (subsidized) premiums for household i, p^s_{ijmt}, using regulated age-rating factors in effect in Oregon in our sample

period. We label these factors $\gamma_{k,i}$ for household member k in household i:

$$\begin{aligned} p_{ijmt} &= p_{jmt} \sum_{k \in i}^{K_i} \gamma_{ki} \\ p^s_{ijmt} &= max\{p_{ijmt} * (1 - subs^p_{ijmt}) - subs^l_{ijmt}, 0\} \end{aligned} \tag{F.3}$$

We simplify the premium subsidy formula for the purpose of searching for counterfactual premiums. In the true formula, premium subsidies for individual market households depend both on household income relative to regulated thresholds and the premium for the second cheapest silver plan offered in a market and year. To simplify our equilibrium search, while still capturing the key features of premium subsidies, we fix subsidy levels before the equilibrium search algorithm. We allow for two types of subsidies. The first is a proportional subsidy $subs^p_{ijmt}$ to be applied to the gross premium. The second is a lump sum subsidy $subs^l_{ijmt}$.

In the main specification, we set $subs^l_{ijmt} = 0$ for individual market households (or small group market households in counterfactuals where the small group market closes) and set the proportional subsidy $subs^p_{ijmt}$ to equal the subsidized share of premiums a household would face in the observed market. Under this method, we capture the differential elasticity of subsidized consumers but still allow a degree of price-sharing to constrain insurers from raising price. We omit the specific demand shift for an insurer that sets its premium to rank as the second cheapest silver plan. For small group market households in the extended ICHRA counterfactuals, we set $subs^p_{ijmt} = 0$ and set $subs^l_{ijmt}$ equal to the household's tax benefit and employer contribution for their enrolled plan in the observed data. This allows for the employer contribution to remain constant as plan purchasing shifts to the individual market. The tax benefit is equal to the household's average tax rate times the gross premium. The employer contribution is equal to 65% of the remaining (post-tax) premium.

F.1 Steps in the equilibrium search

With these components, our algorithm proceeds as follows:

1. Compute each household's expected utility for plan j given the lth guess at the household's effective premium, $p^{l,s}_{ijmt}$.[16]

$$\hat{V}_{ijmt} = x_{jt} + \frac{1}{2}\hat{\omega}_i x^2_{jt} - (\hat{\alpha}_i - \hat{\psi}_i) p^{l,s}_{ijmt} + \hat{\beta}_0 Z_{jmt}$$

2. Compute predicted shares, $\hat{s}_{ijmt}$:

$$\hat{s}_{ijmt} = \frac{exp(\hat{V}_{ijmt})}{\sum_{k=0}^{J_{mt}} exp(\hat{V}_{ikmt})}$$

   where $k = 0$ represents the outside good, with $\hat{V}_{i0mt} = 0$. Here J_{mt} equals the number of insurance options in market m in year t.

[16] In the notation below, we leave off an i subscript from x_{jt}, the actuarial value of plan j at t. In our implementation, we allow x_{jt} to vary by household for those low income households eligible for cost sharing subsidies.

3. Compute expected costs for household i under plan j:

$$\bar{c}_{ijmt} = E[c_{ijmt}|\hat{\alpha}_i, \hat{\omega}_i] = \frac{x_{jt} + \hat{\omega}_i x_{jt}^2}{\hat{\alpha}_i}$$

4. Compute total costs to the insurer, the right-hand side of the price-setting equation:

$$\widehat{R}_{jmt} = \sum_{i \in \{m,t\}}^{N_{mt}} \left(\widehat{s}_{ijmt} * \widehat{\beta}_4 \kappa_{ij} \bar{c}_{ijmt} \right) + \widehat{\beta_5} A_{jmt}$$

We can also compute the left-hand side of the pricing equation using prices:

$$\widehat{R}_{jmt} = p_{jmt}^l \sum_{i \in (m,t)}^{N_{mt}} \left(\hat{s}_{ijmt} \left(\sum_{k \in i}^{K_i} \gamma_{ki} \right) \right)$$

Thus, to find the next guess at the normalized premium vector, p_{jmt}^{l+1}, we can combine the two sides of the equation as:

$$p_{jmt}^{l+1} = \frac{\sum_{i \in \{m,t\}}^{N_{mt}} \left(\widehat{s}_{ijmt} * \widehat{\beta}_4 \kappa_{ij} \bar{c}_{ijmt} \right) + \widehat{\beta}_5 A_{jmt}}{\sum_{i \in (m,t)}^{N_{mt}} \left(\hat{s}_{ijmt} \left(\sum_{k \in i}^{K_i} \gamma_{ki} \right) \right)}$$

5. We test for convergence by comparing the plan-specific price vector p_{jmt}^{l+1} from the $(l+1)$st iteration against the lth iteration across all plans. We compute the mean across plans $j = 1, ...J_{mt}$ in all markets m and t, which total J possible plans:

$$\frac{1}{J} \sum_{\forall (j,m,t)}^{J} \left(\frac{p_{jmt}^{l+1} - p_{jmt}^l}{p_{jmt}^l} \right)$$

If the mean percentage difference in prices (excluding the outside good) is less than .001, we stop the search algorithm and define the equilibrium price vector, p^{eql}, of length J. An entry in p^{eql} is equal to $p_{jmt}^{eql} = p_{jmt}^{l+1}$. If our condition on the mean percentage difference in premiums is not satisfied, we return to Step 1, using p_{jmt}^{l+1} in place of p_{jmt}^l.

F.2 Counterfactual sample

For each counterfactual scenario we examine, we define the set of households able to choose a plan in the individual insurance market in 2016. We include households in the data who purchase an individual market or small group plan plus the set of uninsured households in that year. In addition to this sample, we also add a group of households we label "behavioral types," as in Azevedo and Gottlieb (2017). In each market-year pair (m, t), we add a total number of additional behavior types equal to 1% of households in (m, t). These households incur no health care costs and choose each of the J_{mt} plans in the market-year with equal probability: $s_{ijmt}^{behav} = 1/J_{mt}$. We also assume there is no additional administrative cost to serve these households; that is, A_{jmt} remains unchanged. Following Azevedo and Gottlieb (2017), we must include behavioral types in the counterfactual sample to guarantee existence of the equilibrium.

G Premium Comparison

In this section, we describe how we compute and compare household premiums in the small group and individual insurance markets.

G.1 Premiums in the individual market

For the individual market, we observe each each household's composition in terms of the number of household members and each household member's age. For each plan in the household's choice set in its rating area and in a given year, we collect the standardized premium for those plans. The premium calculation involves the following steps:

1. For each member of a subscriber-HH, collect the age-rating factor according to the standardized age curve for Oregon.[17]

2. Multiply the person-specific age-rating factor against the standardized base premium for the plan.

3. Sum the age-based premium for household members to find the household premium for that plan. Following ACA regulations, for households with more than three dependents we include only the age-specific premiums for the three eldest dependents when computing the household's total premium.

G.2 Premiums in the small group market

Computing household-level premiums in the small group market requires knowing not only the household's size and the ages of household members, but also information on all households in the observed small group. We compute a household's premium for each possible plan as follows:

1. For each member of a subscriber-HH, collect the age-rating factor according to the standardized age curve for Oregon.

2. Multiply the person-specific age-rating factor against the standardized base premium for the plan.

3. Sum the age-based premium for household members to find the household premium for that plan. All dependents' age-specific premiums contribute to the household's total premium.

4. Save the total premium computed above by household for *all* households in the small group. Sum these premiums to find the total "plan side premium"– that is, the total amount of premium dollars the small group pays the insurer to cover the households in the group. Call this sum $TotPrem_{plan}$.

5. To find the premium that each household faces, we divide up the total plan-side premium by weights:

[17] Oregon follows a state-specific age-rating curve. The exact rating factor or multiplier for each age can be found here: https://www.cms.gov/CCIIO/Programs-and-Initiatives/Health-Insurance-Market-Reforms/Downloads/StateSpecAgeCrv053117.pdf

(a) Assign to each household a family rating factor: single = 1, single with children = 1.85, married without children = 2, and married with children = 2.85.

(b) Sum these family rating factors across all households in a small group. For example, if a group consisted of two single households and one household with two parents and children, the groups rating factor would equal 1+1+2.85 = 4.85. Call this sum $TotFRF$.

(c) Define one of four possible household premiums as: single = $1 * \frac{TotPrem_{plan}}{TotFRF}$, single with children = $1.85 * \frac{TotPrem_{plan}}{TotFRF}$, married without children = $2 * \frac{TotPrem_{plan}}{TotFRF}$, and married with children = $2.85 * \frac{TotPrem_{plan}}{TotFRF}$. Households thus don't pay their specific premium contribution, but only one of these four premium levels that divide the entire group premium pool by these fixed factors.

G.3 Comparing premiums in the individual and small group markets

In Figure 3, we illustrate the differences in annual premiums comparing plans in the small group and individual insurance market by year. To do so, for each constructed plan available in the small group market, we identify an identical constructed plan in the individual market that shares the same carrier, plan type, metal tier, and geographic region. Our goal is to find the difference in premiums between small group and individual coverage for an example household that contains a 40-year-old single subscriber who is not eligible for federal premium subsidies in the individual market.

We carry out this analysis in three steps. In our first step, we compare the premiums without any tax adjustments or employer subsidies. To find our household's premium in the small group market, we apply the tiered-composite pricing system as if our household were in a small group with other households. We simulate the other members of the small group using the empirical distribution of household size and age composition that we observe in Oregon:

1. We loop through all observed small groups in Oregon
2. For each group we observe, we replace one random household in the group with our example 40-year old single adult.
3. We use the steps we describe in Appendix Section G.2 to compute our example household's required premium
4. We average these required premiums over the empirical distribution of small groups.

In our plot, we show the distribution of both the small group and individual market premiums that the household would face for all possible plans in a given year.

In our second step, we assume our single household has the median income for single adults in Oregon. We compute the average tax rate this household would face and adjust the small group premium to reflect the household's tax savings from using money for non-taxed insurance premiums rather than taxable income. There is no similar tax advantage in the individual market.

Finally, third, we add, in addition to the tax subsidy, a subsidy from the employer. We assume a 50% employer subsidy, which is the minimum employer's premium contribution that allows smaller businesses to earn the Small Business Health Care Tax Credit.

G.4 Comparing premiums in the individual and small group markets outside Oregon

To compare premiums in the small group and individual insurance markets outside Oregon, we exploit the 'HIX Compare' data (hereafter 'HIX') compiled by the Robert Wood Johnson Foundation. The dataset provides plan-level detail on cost-sharing terms, premiums, and deductibles for nearly all plans offered in the health insurance marketplaces beginning in 2014. We divide our analyses on plans available in the individual insurance and small group insurance markets into two periods, 2014-2016 and 2017-2018.

In 2014, HIX data contains plans from 33 states with both small group and individual insurance plan data. In 2015-2018, all states and the District of Columbia appear with at least one plan in the silver and gold metal tiers in both the small group and individual markets.

In the HIX data, we save raw premiums for a 27-year-old single insurance enrollee for each plan offered in the individual and small group markets. We define a plan as a unique combination of year, state, rating area, carrier, metal tier, plan type, and individual deductible level. We compare the small group and individual markets by averaging the premiums we observe in each business segment over a specific subset of observations; we then match equivalent subsets between business lines.

We run our comparisons between the small group and individual markets using three subset definitions. First, in our most granular comparison, we compare the two markets at the plan level, matching plans by year, state, rating area, carrier, metal tier, plan type, and individual deductible amount. We match many plans across the two business lines. In 2015, for example, we find 2,411 bronze plans, 2,114 silver plans, and 1,378 gold plans that are offered in both the small group and individual market. At the year-state-metal level, we observe an average of 27 matched plans and a median of 16. In effect, under this definition, we must limit our analysis to carriers that offer nearly identical plan designs in both business lines.

Our second definition compares plans across the business lines in a state, year, and rating region, where the plans share only the same carrier and metal tier. That is, we find the mean at the carrier-metal level first in a geographic market and year and then compare these means across the business lines. Because we aggregate to the carrier-metal tier level, we can now use data in our calculations from carriers that offer slightly different types of plans within a metal tier across the business lines.[18] In a typical year-state-metal combination, we observe 29 matched plans on average, with a median of 18.

Finally, our third definition aggregates to a higher level before comparing across business lines. Specifically, we aggregate at the year, state, rating region, and metal tier level. For example, we aggregate all 52 individual market silver plans that carriers offered in rating area 1 in Alaska in 2014. We similarly aggregate the 8 small group silver plans carriers offered Alaska's region 1 in 2014. We then compare these means across business lines. In 2014, we observe 390 year-state-rating region combinations per metal tier. In 2015-2018, we observe roughly 500 such combinations.

Our results for years 2014-2016 appear in Table A6. Using our most granular definition (definition 1), reported in Panel A of Table A6, we see that in 2014-2015, across metal tiers, the small group premiums tend to be $20-30 higher per month than the equivalent individual market plan. As a share of the small group premiums, these differences amount to an increase of roughly 7-12%.

[18] By requiring the plans to have the same metal designation, all the plans we aggregate together share the same actuarial value. However, the plans we aggregate may differ by PPO vs. HMO plan type, and may have a different mix of deductible and cost-sharing terms to achieve the same actuarial value.

In 2016, the small group premiums again exceed the individual market premiums, but by a smaller amount: premiums are on the order of 1-3% higher.

In Panels B and C, we show premium differences using definitions 2 and 3 above, respectively. The results appear similar to Panel A, but slightly higher. Recall, in the later panels, we use more of the observed plan premiums because we do not require the carrier to be operating the same plan across both markets. Instead, we compare mean premiums at either the year-state-region-metal-carrier level or the year-state-region-metal level across the two business lines. In these comparisons, we find premiums are higher by approximately $20-40 per month, amounting to roughly 10-15% higher premiums in the small group vs. individual insurance market for the same actuarial value plan in 2014-2016.[19]

We repeat this analysis in the HIX data for Oregon alone. HIX contains premium information for Oregon only in years 2015-2016. As in our main data, we again find that small group premiums exceed comparable premiums in the individual marketplaces. In Oregon, under the plan-level and carrier-metal definitions in the years we study, the gap between small group and individual insurance premiums is slightly larger than the national average.

Finally, as a robustness, we repeat our premium comparison for years 2017-2018. We omit an extensive discussion of these years in our main analysis both because we do not have claims data for these years and because there were two important regulatory changes in 2017. First, prior to 2017, the federal government paid compensation to insurers who enrolled consumers eligible for cost-sharing subsidies. Eligible consumers who purchase silver plans face reduced cost sharing (higher actuarial value plans) by statute, leading to added costs borne by the insurer; beginning in 2017, the federal government no longer reimbursed insurers for these added costs. Many insurers responded to this anticipated change by "silver loading" in 2017, meaning they raised the silver plan premiums to compensate for the added costs. Second, the temporary risk corridors program, administered under Section 1342 of the ACA, ended after the 2016 plan year. The program limited both the gains and losses to insurers offering plans in the early years of the individual marketplaces, with the goal of encouraging insurer participation.

In part reflecting these regulatory changes, we find the unsubsidized individual market premiums exceed the comparable small group premiums in 2017-2018, across all plan tiers. Using the same definitions we apply to 2014-2016, we find higher baseline premiums on the order of 10-15% in the individual market. These differences appear both in Oregon and nationally.

[19] As a robustness, we also compute the minimum across plans that fit in a business line under definitions 2 and 3. We then compare the minimums across the small group and individual markets. That is, rather than average the premiums at the year and market level within a metal tier, we find the minimum across all plans or across all plans within a carrier. The qualitative findings are similar when using the minimum order statistic in place of the mean.

H Appendix Tables and Figures

Table A1: Full sample of forced switchers

Destination	**Grandfathered**	**Non-Grandfathered**	**Total**
Individual Market	4378	943	5321
Coverage Through Spouse	1000	207	1207
Large Group Market	565	122	687
Other	184	49	233
Uninsured	7433	1750	9183
Total	13560	3071	16631

Note: This table describes the destination of small group enrollees whose plans where closed in 2014 or 2015. "Other" includes coverage as a dependent or coverage through Medicare, another small group plan, an association plan, or other miscellaneous types of plan. "Uninsured" comprises enrollees who do not fall into any other category and are therefore untracked.

Table A2: Administrative costs

Variable	Mean	S.D.	Median
Per-enrollee monthly admin costs in the individual market ($)			
Total	71.53	28.24	62.86
Taxes and fees	21.23	25.64	11.96
Wellness activities	4.14	2.99	3.64
General administrative expenses	46.17	16.22	42.85
Number of annual subscribers in Oregon			
Individual market	25,931	31,702	16,778
Small-group market	21,883	13,197	19,812
Large-group market	84,770	92,063	42,990
Other markets	40,772	68,557	6,543
Number of states where the payer is active	7.5	14.8	2

Note: This table displays measures of insurers' administrative costs. We define all cost measures from publicly available data collected by the Centers for Medicare and Medicaid Services (CMS). We collect this insurer-state level data for the individual market for the years 2014-2016. In the table, we report total per-enrollee monthly administrative costs and the constituent elements of administrative costs: (1) taxes and fees, which include federal and state taxes deductible from premiums, community benefit expenditures, contribution to the Federal Transition Reinsurance Program, and various other regulatory fees; (2) wellness activities, which include activities to improve and promote health outcomes and prevent hospital readmission as well as expenses related to health information technology; and, (3) general administrative expenses, which include direct sales salaries and benefits, agents and broker fees, and all other non-medical general and administrative expenses. In addition, in the table we report other measures available in the CMS data. We compute statistics on the count of annual subscribers in Oregon that each insurer covers in each business line. Finally, we present statistics on the number of states, including Oregon, where payers that operate in Oregon are active. We classify an insurer as active in a state when it files a report with CMS in that state and year.

Table A3: Demographics of the switchers

	Small-Group		Switchers	
Demographic variable	**Mean**	**S.D.**	**Mean**	**S.D.**
Single-membered	0.75		0.73	
Married, no dependent	0.08		0.08	
Not married, with dependent(s)	0.07		0.09	
Married, with dependent(s)	0.10		0.10	
Number of dependents	2.21	1.25	2.30	1.32
HH sum of health status scores	1.25	2.20	1.33	2.56
Income (as ratios of the FPL)	2.38	0.29	2.39	0.28
Age	42.62	11.28	45.56	11.10
Over-50	0.28		0.38	
Living in rating areas 1, 2, or 3	0.78		0.67	
Number of subscriber-year observations	383,036		16,631	

Note: This table compares the demographic characteristics of the population of "forced switchers" who lose small group coverage to the larger population of small group enrollees in years 2014 - 2016. The first column reports summary statistics for the full small group sample. The second column includes statistics for all households in the switcher sample in the year prior to their exit from small group coverage. We compute the sum of health status scores for all members of a household, where we predict each member's score using the Johns Hopkins' ACG software. Number of dependents is calculated with the subset of households who have dependents. Rating areas 1-3 include the urban areas of Portland, Eugene, and Salem, respectively. Rating areas 4-7 include largely rural areas of the state.

Table A4: Demographics of the uninsured

	Uninsured Households		Insured Households	
Demographic variable	**Mean**	**S.D.**	**Mean**	**S.D.**
Single-membered	0.51		0.70	
Married, no dependent	0.16		0.15	
Not married, with dependent(s)	0.13		0.06	
Married, with dependent(s)	0.20		0.09	
Number of dependents	1.97	1.34	1.91	1.10
HH sum of health status scores	1.90	3.34	1.44	2.49
Income (as ratios of the FPL)	2.37	0.29	2.46	0.29
Age	37.65	12.41	47.00	11.80
Over-50	0.22		0.45	
Living in rating areas 1, 2, or 3	0.68		0.69	
Number of subscriber-year observations	600,487		444,255	

Note: This table compares the demographic characteristics of uninsured households and households who chose plans in the individual market in years 2014 - 2016. We compute the sum of household health status scores for all members of a household, where we predict each member's score using the Johns Hopkins' ACG software. Number of dependents is calculated with the subset of households who have dependents. Rating areas 1-3 include the urban areas of Portland, Eugene, and Salem, respectively. Rating areas 4-7 include largely rural areas of the state.

Table A5: Estimated model of small group fees

	Specifications				
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Constant					
Bronze	1.449	2.042	2.069	1.930	1.906
	(0.022)	(0.111)	(0.111)	(0.112)	(0.121)
Gold	1.276	0.252	0.262	0.380	0.334
	(0.022)	(0.104)	(0.104)	(0.105)	(0.114)
Platinum	2.003	1.515	1.303	1.061	1.260
	(0.029)	(0.130)	(0.133)	(0.134)	(0.143)
Age					
Bronze		-0.015	-0.015	-0.015	-0.016
		(0.003)	(0.003)	(0.003)	(0.003)
Gold		0.027	0.029	0.030	0.030
		(0.003)	(0.003)	(0.003)	(0.003)
Platinum		0.013	0.025	0.024	0.025
		(0.003)	(0.004)	(0.004)	(0.004)
HH sum of health scores					
Bronze			-0.335	-0.338	-0.343
			(0.026)	(0.026)	(0.027)
Gold			-0.380	-0.385	-0.395
			(0.023)	(0.023)	(0.025)
Platinum			-0.690	-0.705	-0.656
			(0.030)	(0.030)	(0.031)
Group size					
Bronze				0.008	0.008
				(0.001)	(0.001)
Gold				-0.005	-0.005
				(0.000)	(0.000)
Platinum				0.014	0.015
				(0.001)	(0.001)
Family size					
Bronze					0.025
					(0.049)
Gold					0.053
					(0.050)
Platinum					-0.228
					(0.053)
Number of observations	46,310	46,310	46,310	46,310	46,310

Note: This table reports the maximum likelihood estimates from the employer preferences specification in Equation D.13. In this analysis, we omit households who choose grandfathered or other plans for which we do not observe premiums or metal tier. All variables are interactions between metal tier indicators (silver-tier excluded) and household characteristics. The household characteristics that we use include: the primary subscriber's age; the sum of health scores of all individuals in the household; the number of households in the household's small group (contract number); and the number of individuals in the household.

Table A6: National comparison of small group and individual market premiums, 2014-2016

Panel A: Plan-level definition								
			SG-IM in $/month			% Difference		
Year	Metal	Count	Mean	Median	Std	Mean	Median	Std
2014	Bronze	451	21.97	17.32	21.48	11.70	9.52	9.96
2014	Gold	319	21.32	20.88	33.40	7.23	6.29	10.91
2014	Silver	390	23.17	18.98	26.00	9.35	7.60	10.04
2015	Bronze	2411	22.24	21.37	30.71	10.80	8.71	14.61
2015	Gold	1378	25.84	16.99	44.24	8.54	5.41	14.29
2015	Silver	2387	30.89	29.55	38.50	12.03	10.40	15.15
2016	Bronze	1573	2.95	0.00	36.24	2.12	0.00	14.72
2016	Gold	792	1.57	2.47	51.48	1.38	0.57	13.36
2016	Silver	1694	5.42	1.48	42.78	3.05	0.59	14.60
Panel B: Carrier-metal definition								
			SG-IM in $/month			% Difference		
Year	Metal	Count	Mean	Median	Std	Mean	Median	Std
2014	Bronze	559	13.76	21.31	88.57	11.28	11.66	16.83
2014	Gold	741	21.24	28.77	90.37	9.76	9.31	15.20
2014	Silver	741	21.32	30.75	87.87	11.67	10.94	16.30
2015	Bronze	2432	32.21	33.70	34.53	15.06	14.35	15.76
2015	Gold	2244	47.27	44.70	57.51	16.24	14.36	18.70
2015	Silver	2417	45.59	41.07	48.90	17.65	14.49	17.94
2016	Bronze	1932	10.99	7.58	37.82	5.46	3.12	15.41
2016	Gold	1831	16.14	16.26	56.19	5.81	4.89	15.48
2016	Silver	1885	20.18	15.76	44.99	7.71	5.40	15.31
Panel C: Rating area-level definition								
			SG-IM in $/month			% Difference		
Year	Metal	Count	Mean	Median	Std	Mean	Median	Std
2014	Bronze	322	12.11	20.40	75.03	9.99	10.00	19.49
2014	Gold	390	25.44	33.06	71.23	10.56	11.59	15.76
2014	Silver	390	25.83	28.85	58.76	12.11	11.60	16.40
2015	Bronze	498	30.85	31.43	24.14	13.06	13.00	10.52
2015	Gold	499	39.52	40.75	36.40	12.34	12.46	11.48
2015	Silver	499	57.42	59.72	32.18	21.37	21.27	12.19
2016	Bronze	498	6.26	5.74	33.67	3.25	2.29	12.76
2016	Gold	499	27.29	25.09	50.10	8.68	7.27	14.35
2016	Silver	499	38.25	39.30	42.85	14.17	13.06	15.15

Note: This table reports the difference in premiums between the small group and individual insurance markets from HIX Compare data. We collect baseline premiums for a 27-year-old single enrollee in a plan for 33 US states in 2014 and all US States and DC in 2015-2016; premiums for the plan are specific to the year, state, rating region, carrier, metal tier, plan type, and deductible level. Panel A reports statistics for a plan-level comparison. Panel B reports statistics after aggregating plans to the carrier-metal-market. Panel C reports statistics after aggregating plans to the metal-market in each business segment.

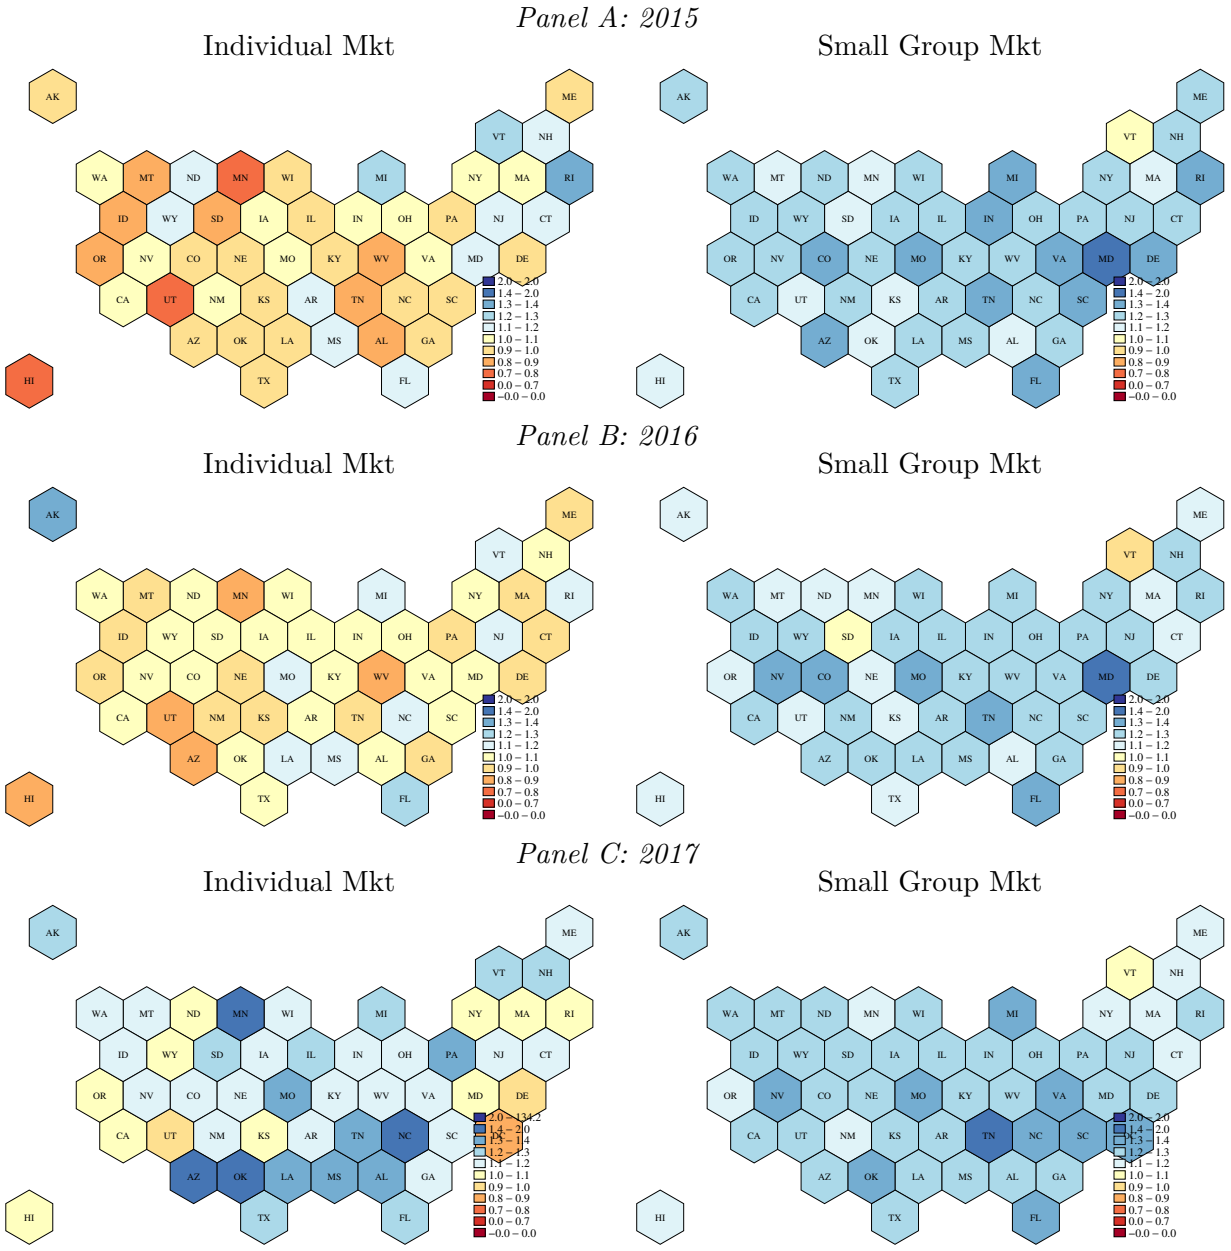


Figure A1: Average medical markups by state, year, and insurance segment

Note: This figure depicts the medical markup by state, year, and insurance market segment computed using Medical Loss Ratio data. The medical markup is calculated as the ratio of premium revenue divided by the total medical costs insurers incur. The cost in this ratio does not account for risk adjustment payments or other transfers to the insurer.