

Market Segmentation and Competition in Health Insurance

Michael J. Dickstein

New York University and NBER

Kate Ho

Princeton University, NBER, and CEPR

Nathaniel Mark

U.S. Department of Justice

April 4, 2023

Abstract

In the United States, households obtain health insurance through distinct market segments. To explore the economics of this segmentation, we consider the effects of pooling coverage provided through small employers and through individual marketplaces. We model households' demand for insurance and health care, along with insurers' price-setting, to predict equilibrium choices and premiums. Applying our model to data from Oregon, we find that pooling can mitigate adverse selection in the individual market and benefit small group households without raising taxpayer costs. Our estimates provide insight into the effects of new regulations that allow employers to shift coverage to individual marketplaces.

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] Varying degrees of insurance segmentation also exist in other developed countries, including Germany, where wealthy households and workers in specific occupations obtain private coverage instead of statutory social insurance, and Australia and Canada, where 46% and 67% of households, respectively, elect additional private coverage to supplement public schemes (Tikkanen et al., 2020).

We explore how a segmented insurance system affects the degree of adverse selection, the level of insurance premiums, and ultimately consumer surplus. Focusing on the U.S., we consider 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.^[2] 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 (ICHRA), employers fund tax-exempt accounts through which employees purchase individual coverage. If adopted widely, ICHRA could transform the U.S. insurance system, eliminating the employer's role in plan choice.

We thank conference and seminar participants at Boston College, Columbia University, Stanford University, University of Texas at Austin, University of Pennsylvania, NYU Stern IO Day, and the NBER IO Winter meeting 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, Dexin Li, and Nicole Ng for excellent research assistance. 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. Edited by Chad Syverson.

¹These 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).

²In 2022, Vermont separated its markets to take advantage of subsidies for the individual market available under the American Rescue Plan Act (Hall and McCue, 2021).

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 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 employees. In a typical market, employers choose one or two plans to offer to eligible households (Agency for Healthcare Research and Quality, 2016); the chosen level of coverage could reflect an employer’s incentive to provide greater tax-advantaged benefits to employees or a broker’s incentive to recommend plans that pay higher sales commissions. These agency frictions—and the associated administrative costs, like broker fees—can increase the gap between plan premiums and health costs.

We quantify the welfare consequences of market segmentation using data from Oregon’s insurance markets in years 2014 through 2016. These dates span the first three years of operation of the individual marketplaces created under the Affordable Care Act. Our data include information on premiums, household enrollment status by year and market segment, and health spending, for all private insurance plans in Oregon. We also observe household transitions between different types of private insurance, including between individual coverage and group coverage.

We begin our empirical analysis 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 age-specific gross premiums that are 7% higher than they would be in the individual market.³ Combining these facts, and controlling for government subsidies, we find that insurer markups over health care costs are higher in the small group market in our time period. We externally validate this finding by showing similar patterns exist in data from states outside Oregon in years 2014 through 2016. In 2017, after the expiration of two federal premium support programs, markups remain higher in the small group market in 36 states, but the differences between markets shrink.

³We 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.

When comparing plan choices, 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 \$811 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 and covered health spending per household averages only \$439 per month.

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, which themselves depend on households’ characteristics and preferences for 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.

On the demand side, we predict a household’s plan choice and subsequent medical spending using a model that features risk aversion, moral hazard, and heterogeneity in a household’s underlying health needs (Einav et al., 2013). Estimating preferences in the individual market is straightforward: we observe each household’s choice of plan and realized level of health spending. We can match these data to our model predictions in a likelihood framework. 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 exploit the set of small group households who lose employer coverage in our time period and must purchase coverage on the individual market. We show these “forced switcher” households appear similar along observable demographics to the small group sample as a whole. 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 for the individual market. In our model, insurers choose premiums as a function of predicted patient costs and insurer-specific administrative costs. We do not, however, require plan-specific premiums to equal average costs; we allow flexibility to account for either net transfers the plans receive from risk-related programs or plan-level markups over medical costs. Under our specification, a coefficient on predicted medical costs above one indicates a positive medical markup. We estimate a coefficient of 0.73 using our data, reflecting net subsidies from government risk-related programs rather than markups.

We combine our supply and demand side estimates with a model of competition to predict the effect of pooling the small group and individual markets on enrollment, costs, and consumer surplus. Given the evidence in our setting against net markups, we proceed in our analysis with a model of competition and price-setting among insurers in the pooled market that follows [Azevedo and Gottlieb \(2017\)](#). However, acknowledging that other time periods or geographic markets might differ in their degree of competition, we also simulate the new equilibrium were insurers to respond to pooling by charging positive (fixed) markups, up to the maximum level permitted under the ACA. While we allow insurers to re-price, we fix the set of plans they offer in the new equilibrium to be the same as the plans offered in the individual market before pooling.

In our counterfactual analyses, we highlight the role of adverse selection and insurer markups in determining the overall welfare change from pooling. We show how household preferences—including a household’s degree of moral hazard, risk aversion, and underlying health spending needs—can interact with insurer competition to determine which consumers win or lose from pooling. We quantify these effects under a specific policy that requires all small group employers to offer ICHRA accounts. In this setting, we also impose a mandate under which small group households must use these accounts to enroll in individual coverage; they cannot choose to be uninsured. When small group households purchase insurance on the individual market, they keep the same tax benefits and the same employer contribution that they earned with group coverage.

Under this policy, we find that pooling the markets mitigates adverse selection in the individual market and generates consumer surplus gains among individual market households equal to \$44

per month. Small group households gain \$58 per month in consumer surplus.⁴

Our model predictions depend on the observed plan choices and costs for households in Oregon and the pricing behavior of insurers operating in the state. We generalize our findings on the value of pooling to settings where (a) health spending needs and preferences in the small group and individual market populations differ from our Oregon setting and (b) insurers add markups to premiums. Not surprisingly, the surplus gains to individual market households fall as small group households have greater health spending needs or insurers charge positive markups in a pooled environment. However, we also find that even if small group consumers have identical spending needs to individual market households, pooling nonetheless reduces adverse selection because small group members, at all health levels, are predicted to have stronger preferences for more comprehensive plans.

Overall, we show that for an empirically relevant range of market characteristics, the adoption of ICHRA plans benefits the average household and improves adverse selection in the individual market. Thus, when employers subsidize 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, 2022). 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.

⁴An important caveat to our approach, which also applies to other models in the literature, is that we assume consumers make rational decisions. We do not model issues like consumer search frictions or mistakes (Brot-Goldberg et al., 2017; Handel and Kolstad, 2015). However, our counterfactual analysis is a step towards addressing this issue. When we impose a mandate that small group employees must enroll in a plan, we address one possible consumer mistake—the choice to remain uninsured. We ask whether small group consumer surplus still increases with market pooling when tied to a mandate. We leave further analysis of policies that target consumer mistakes to future research.

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, 2022; Jaffe and Shephard, 2020; Polyakova and Ryan, 2019), risk adjustment (Geruso, Layton, and Prinz, 2019; Einav, Finkelstein, and Tebaldi, 2019), participation penalties (Diamond et al., 2021), and re-classification risk (Handel, Hendel, and Whinston, 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, Royalty, and Drake, 2019; Bundorf, Levin, and Mahoney, 2012; Fleitas, Gowrisankaran, and Lo Sasso, 2020).

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 Sections 7, 8, and 9, we discuss our results, simulate counterfactual pooling outcomes, and generalize our findings to other market settings. Section 10 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 online platform, healthcare.gov, 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.⁵

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.⁶ 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

⁵Households with incomes at or below 138% of FPL, as well as those meeting several other criteria, are eligible for Medicaid, the joint state and federal means-tested government insurance program.

⁶Some 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.

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.⁷ Partly as a result of these requirements, consumers in the individual market often choose plans from a large menu. For 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

⁷Oregon’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).

a small group marketplace during the span of our data. Instead, small employers purchase plans through brokers who typically receive a per-month, per-household fee from the insurance carrier for each household they enroll.⁸ 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 two broker-recommended plans, often from the same carrier (Agency for Healthcare Research and Quality, 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. Their exemption from federal and state taxes creates a tax wedge relative to the individual market, where households typically pay for insurance with post-tax dollars.⁹ In addition, unlike standard community rating in the individual market, small group market insurers must 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).

2.3 Risk-related programs

Three risk-related programs operate in either the small group or individual insurance markets during our sample period: risk adjustment, risk corridors, and reinsurance. The programs work in similar ways to limit the downside risk to insurers from participating in the market by transferring funds from insurers who enroll households with lower health spending to insurers that enroll households with higher health spending in a given year. We provide a detailed description of each program in Appendix H.

⁸The average small group commission in Oregon in 2016 was \$19.70/enrollee per month (The Kaiser Family Foundation, 2020).

⁹Premiums 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.

Both risk adjustment and the risk corridors program operate in the small group and individual markets. The programs determine transfers an insurer pays or receives separately in each market segment. For risk adjustment, a plan’s risk score—an ex ante measure of enrollees’ expected health spending—determines whether the plan pays into the risk adjustment program or receives payments from the program. The risk corridors program, in effect from 2014 to 2016, similarly transferred funds between plans, but did so as a function of realized spending relative to a target level of 80% of a plan’s premium revenues.^[10] Although the language of the ACA allowed the risk corridors program to pay out more than the amount of penalties it collected, in practice payouts were set to be revenue neutral between 2014 and 2016.^[11] The likelihood of reductions in risk corridor payouts created uncertainty for insurers. Finally, the reinsurance program, which operated between 2014 and 2016, collected funds from insurers offering fully insured plans in any segment, including the individual, small group, and large group markets. Only individual market plans, however, could receive transfers from the reinsurance program.

We design our model of premium setting to estimate the proportion of costs that are covered by risk-related payments, rather than assuming it, and we allow for uncertain expectations on the part of carriers.

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 use Medical Loss Ratio (MLR) reports from the Centers for Medicare and Medicaid Services to define measures of administrative costs. We describe each data source below.

¹⁰If a plan’s health care costs plus the costs of qualifying quality improvement programs fell below 97% or above 103% of the target threshold, the plan either paid a penalty to the program or received a payout.

¹¹In 2014 in Oregon, the program paid out only \$12.90 million of the \$102.23 million owed to individual market insurers and \$660,326 of the \$5.23 million owed to small group insurers.

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 to 2016. Our claims data cover out-of-pocket costs and costs to the insurer for inpatient, outpatient, 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 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.¹²

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.¹³ The results are then normalized to an ACG score, where a score of one corresponds to the 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 I 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

¹²We 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.

¹³We 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.

households in the individual market¹⁴

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 type¹⁵. In Appendix Table A7, we show that our constructed plans are nonetheless sufficiently detailed that there is little variation in plan premiums, cost sharing terms, or networks within a constructed plan. Hereafter, we refer to constructed plans simply as plans.

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 specifying our model, we conduct several descriptive analyses to highlight the differences between the small group and individual markets in Oregon in terms of plan choices, health spending,

¹⁴While over 70% of enrollees in both samples purchase insurance for a single household member, some of these enrollees may have a spouse or other family member who accesses health insurance through a distinct employer or public program.

¹⁵We 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.

premiums, and the level of insurer markups. These differences will motivate our modeling choices and illustrate the key factors that underpin our estimates of the consumer surplus change from market pooling: household preferences, the extent of adverse selection, and the level of markups in the existing market segments.

4.1 Comparing plan choices and costs

We begin by examining plan design choices and costs between the small group and individual insurance markets. As we report in Table 2, 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.

Despite enrolling in more comprehensive coverage, however, small group households spend less on health care each month. 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—that is, the tendency to spend more on health care when enrolled in a more generous plan—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 69% 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 to 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 is 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.

The observed heterogeneity of households’ health spending between the two markets, coupled with differences in plan preferences, has immediate welfare implications for market segmentation. Any non-random division of the population along these dimensions will leave opportunity for welfare-enhancing swaps, including the possibility of Pareto improvements. For example, consider the case in which we introduce a relatively low-cost set of small group households with a lower premium elasticity and higher degree of risk aversion into the individual marketplace. If these newly shifted households continue to choose gold plans, their enrollment will reduce the existing plans’ average costs and premiums, potentially spurring more individual enrollees to opt for a gold plan. The average costs of enrollees in silver and bronze plans, in turn, may fall as relatively sick enrollees leave these plans for the now cheaper gold metal tier. Under average cost pricing or fixed markups, silver and bronze plan premiums will fall. The cost to small group enrollees from this pooling depends, in part, on the equilibrium premiums they faced in the segmented market. We explore this issue next.

4.2 Comparing premiums and markups

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 a 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 silver and gold plans. 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 out-of-pocket premiums that households face. The household’s premium depends on (a) pooling within an employer (b) the tax savings for employer-provided insurance; and, (c) employer contri-

butions towards premiums. Appendix [G](#) describes these adjustments.^{[16](#)} 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. Ultimately, the relative differences in premiums between the two segments will depend on the size of the employer’s subsidy relative to possible government subsidies in the individual market.

Given our measurement of both costs and premiums, we can also compute an 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 premium paid for small group plans independent of employee tax savings or employer subsidies, since the total premium flows to the insurer as revenue. In Figure [3](#) 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 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.

From Figure [3](#), it appears that markups are large in the small group market and that many individual market plans failed to collect sufficient premium revenue to cover health care expenses. However, these markups offer an incomplete picture of true profitability because (a) they do not include non-medical costs, including the costs of broker fees and plan administration, and (b) they do not include risk adjustment payments and payments from the risk corridor and reinsurance programs, both of which operated in the individual marketplaces from 2014–2016. In Figure [4](#) Panel B, we add risk-related payments to individual market insurer revenue and plot the adjusted markups. Again, small group markups exceed those seen in the individual market on average,

¹⁶In 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.

though the distributions are closer.

Given the importance of risk-related payments, the possibility of insurer markups, and the potential for administrative costs to affect premium setting, we allow for these elements when we model the insurer’s premium choices and solve for equilibrium premiums in the individual market.

4.3 Comparison to other markets

Finally, while we focus our analysis in this paper on the insurance markets in Oregon, we investigate how premiums and markups compare across the small group and individual markets in other states. 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 Appendix Table [A10](#). 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 2017-2018, we observe a different pattern: individual market premiums exceed small group base premiums by approximately 10%. Our findings for 2018 are similar to results from [Abraham, Royalty, and Drake \(2019\)](#). Individual market premiums increased in these years in many states reflecting, 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 ended after the 2016 plan year.

We can translate these premiums differences to markups using the data carriers report in each state as part of MLR regulation.^{[17](#)} We 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. In 2017, average markups in the individual market increase relative to prior years, reflecting higher premiums. However, even when premiums in the individual market exceed those in the small group market, markups remain larger in the small group market in 36 of 50 states. These higher markups

¹⁷We 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. In this analysis, we also do not adjust for risk corridors and reinsurance payments.

reflect lower premiums but also lower medical costs among small group enrollees. We illustrate the average markups by sector, state, and year in Appendix Figure [A1](#).

Given the heterogeneity that exists across years and geographies in the competitiveness of the market segments, we later use our model to explore how different markups may affect our conclusions about the benefit of pooling. In the following sections we outline our model and describe our approach to estimation.

5 Model

We build a demand model that predicts households' plan choices and subsequent medical care utilization in the individual market. 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.^{[18](#)} In Stage 1, the household chooses an insurance plan from a set of available options. 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.

¹⁸A household in our setting includes a head of household and any dependents covered by a chosen plan. We do not model the potential choice of a family member to obtain coverage distinct from the rest of the household.

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:

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

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, $g(\cdot)$, 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 [\(5.1\)](#):

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

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 . If $\lambda = 0$, medical spending will equal zero 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 set of options, \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 follow 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

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

where ψ is the household's coefficient of absolute risk aversion and the household's optimal plan choice 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 with 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. Under perfect competition, insurance carriers in the individual market choose their plan-specific premiums to equal average costs, including both expected claims and administrative costs. In our specification, which we describe in Section [6](#), we also allow flexibility for premiums to differ from costs with a fixed markup that can either be below one, reflecting net subsidies to the insurer in the form of risk-related payments, or above one, reflecting positive plan markups over health costs.

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:

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

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 :

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

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:

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

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.

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 a count of the uninsured population in each rating area and year, differentiated by 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 estimates of the preferences and attributes of current small group households. To estimate these preferences, we proceed by collecting a 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 or 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.

In addition to the switcher sample, we use the observed choices of the larger sample of small group households to explore how much the employer’s plan choice differs from the preferences we recover from the switcher population. Our employer model, described in detail in Appendix [D.2](#), generates estimates of “fees” employees would have to pay to deviate from the employer’s plan choice to their preferred option. These fees will be non-zero whenever plan choices in the small group market do not match the revealed preferences of the forced switcher sample.¹⁹ We allow these fees to vary by household characteristics, including age, health status score, household size, and employer-group size. We use the estimates to generate an alternative measure of consumer surplus in the small group market.

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 in the individual market, accounting for both moral

¹⁹The differences between the employer’s preferences and the household’s preferences could reflect incentives of the employer for more generous coverage—say for recruiting purposes—or could reflect a household’s behavioral biases due to procrastination, the choice architecture, or the incentive to churn out of coverage ([Diamond et al., 2021](#)). We do not take a stand on the drivers of these differences.

hazard and risk aversion.

6.2.1 Equations for estimation

We begin by specifying the household's health state, λ , using an exponential distribution to approximate the empirical spending pattern depicted in Figure [1](#). If λ follows an exponential distribution with parameter α , then 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:

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

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 [20](#):

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

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

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:

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

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

²⁰This 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 to affect utility. In addition, risk coverage also generates utility: the positive-valued terms adjust upward when we divide by $\alpha_i - \psi_i$ instead 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:

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

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

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 type I extreme value distribution. The probability that an enrollee i chooses plan j then takes the standard logit form:

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

where $V_{i,j} = x_{i,j} + \frac{1}{2}\omega_i x_{i,j}^2 - (\alpha_i - \psi_i)p_{i,j} + \beta_0 X_j$, 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 chosen specification of the household’s health state, 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) . Implicitly, we match observed costs to predicted costs for each plan generosity level, $x_{i,j}$, and for observable household health characteristics. In this set-up, we treat plans offered by different insurers symmetrically.

In addition, we 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 censored²¹

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

²¹In 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](#)

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

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

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

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](#).

Finally, we exploit our detailed claims data to specify the household parameters, $(\alpha_i, \omega_i, \psi_i)$. In particular, our data contain demographic information on each household member, details of their past and present diagnoses, and a record of the frequency and type of health spending they undertake, including office visits, laboratory tests, prescription drugs, and hospitalizations. As we describe in Section [3](#), we can summarize each household member's health state by computing a severity score, the ACG score, as a function of the member's claims. We then specify α_i using this measure along with demographics relevant to predict non-discretionary health spending: the sum of the ACG scores in the household; an indicator for whether a family member has 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 each market segment.^{[22](#)}

Other plan characteristics. In addition to the utility from coverage and the disutility from premium payments, we allow enrollees to have additional preferences for distinct plans. Because

²²Setting 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.

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.²³

Under perfect competition there is no role for an unobserved quality term in the utility equation. In our model, where we already allow flexibility for premiums to differ from average costs, we also allow plans in the individual market to be differentiated in ways that consumers value. Following prior literature (Polyakova and Ryan, 2019; Tebaldi, 2022; Tebaldi, Torgovitsky, and Yang, 2023), we address the resulting endogeneity issue in demand estimation 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 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.²⁴

6.3 Premium setting

We develop a simple model of premium-setting that captures the relationship between observed premiums by plan, expected health costs for those enrollees likely to select each plan, and carrier-specific administrative costs. We use the estimates from this model, along with assumptions about plan markups, to re-compute equilibrium premiums in a counterfactual environment with different populations of potential enrollees.²⁵

²³We 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.

²⁴Given 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.

²⁵We chose not to assume Bertrand competition between plans in this selection market for three reasons. First, our pricing assumption from Azevedo and Gottlieb (2017) guarantees that our fixed point algorithm identifies a unique equilibrium. Second, the coefficient on medical costs that we estimate in Equation (6.6) suggests that markups are

To begin, 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:

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

Next, we can estimate the costs of each plan by finding the vector of parameters that equate the insurer's revenue with its total expected health and administrative costs. If we ignored payments to the plan from risk-related government subsidies and assumed perfect competition—that is, no insurer markups—we could equate premiums to costs such that the insurer broke even on each plan it offered, summing expected costs across individuals and geographic markets in a given time period²⁶. However, given both the presence of risk-related payments to insurers in the individual market in 2014 through 2016 and the possibility of markups, we add an additional element of flexibility in our empirical model. Below, we equate total premium revenues for plan j to its total costs but we add a coefficient β_4 on expected health spending:

$$R_j = \sum_{i \in \{t\} \forall m}^N \left(\hat{s}_{i,j} * \beta_4 K_{i,j} c_{i,j}^e \right) + \beta_5 * A_j + \eta_j \quad (6.6)$$

where $\hat{s}_{i,j}$ is the probability that household i chooses plan j , $c_{i,j}^e$ is the expected medical claims

not a first order issue in the individual market in our setting. Finally, similar to the approach of [Tebaldi \(2022\)](#), we depict the empirical relationship between premiums and health care costs across plans in Figure [3](#). In our setting, after adjusting for risk-related payments, we find premiums roughly equate to insurer spending on enrollees' health care.

²⁶The assumption that carriers break even plan by plan is consistent with the approach of [Azevedo and Gottlieb \(2017\)](#).

cost of household i in plan j , and $\kappa_{i,j}$ adjusts costs to account for possible government cost-sharing subsidies²⁷

The parameter vector β_4 allows the expected claims costs to be higher or lower than the premium revenue, after accounting for administrative costs by carrier, year, and market. When $\beta_4 < 1$, our predicted health spending does not directly equate to premium levels because of risk-related payments to the insurer and/or true ex ante expectations that differ from $c_{j,t}^e$. If $\beta_4 > 1$, the insurer's premiums could be thought of as incorporating a medical markup over expected variable health care costs, similar to the interpretation in Bundorf, Levin, and Mahoney (2012).

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. Given that insurers report administrative costs to justify their premium rates, we allow true administrative costs to differ in unobserved ways from the quantities we observe in regulatory filings.

We begin by estimating equation (6.6) via ordinary least squares, omitting data from 2014, the first year of the marketplaces, because insurers' premium-setting may not have reached an equilibrium in that year. We then adopt a two-stage least squares approach to address the possibility that

²⁷We adjust for subsidies for household i by multiplying expected costs by $\kappa_{i,j} = x_j^o/x_{i,j}$, where x_j^o is the lower actuarial value of the plan without any cost-sharing subsidies. Implicitly, we assume the insurer's realized costs do not include the share of out-of-pocket costs the federal government pledged to subsidize. Given that spending follows an exponential distribution and given our first-stage estimates, we can write $c_{j,t}^e$ (before adjusting for subsidies) as:

$$c_{i,j}^e = E[c_{i,j}|\hat{\alpha}_i, \hat{\omega}_i] = \frac{x_{i,j} + \hat{\omega}_i x_{i,j}^2}{\hat{\alpha}_i}.$$

η_j may be correlated with the predicted market share and claims variables.^[28] We use instruments, Z_j , that we assume to be correlated with $s_{i,j}$ and $c_{i,j}^e$ 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.^[29] 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:

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

The ex ante consumer surplus then becomes:

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

Each household chooses from J plans. We compute this measure both for individual market and small group market households.^[30]

²⁸For 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 .

²⁹In 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.

³⁰Because 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.

7 Estimation results

7.1 Demand estimates

We begin by reporting our underlying demand estimates in Table 3. The 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.³¹

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 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 \$560 per month in the individual market.³² 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.

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 no insurance to full insurance would increase medical spending by 22%. For individual market enrollees, the increase would be 11%. We can

³¹We 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.

³²Our 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.

identify these distinct levels of moral hazard in the two samples by combining information on households' observed plan choices and total health spending. For example, suppose we have two households, A and B, that choose plans with the same generosity. In addition, the members of both households have the same predicted non-discretionary health needs, λ_i , due to their diagnoses and demographics. However, household A has a larger observed spending level than household B. Our model would predict larger moral hazard for household A to rationalize the spending difference. In Appendix Figure [A2](#), we show such a pattern exists at the segment level: when we compare spending in the individual market versus the small group switcher populations, for the same plan generosity and underlying λ_i , we observe a higher median level of spending among small group switchers.

For risk aversion, we find 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). The slightly higher risk aversion in the switcher sample reflects a pattern in the data in which, conditional on underlying non-discretionary health spending needs, λ_i , we observe a higher share of switchers choosing more generous plans, such as gold plans, than do individual market households. This share is even higher than would be implied by selection on their higher moral hazard.³³ We report this difference in plan choices in Appendix Figure [A3](#). Overall, 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, 2022](#); [Ho and Lee, 2020](#)).

Finally, we report estimates from an additional version of demand in the small group market. Under the assumption that we can recover consumers' true preferences from the forced switcher sample, we can compare (a) the choices that we predict employees would make with the freedom to choose any plan with (b) the employer's observed choice. Using a discrete choice framework, described in Appendix [D.2](#), we estimate the fees an employee would have to pay to deviate from

³³More precisely, when we condition on λ_i , both selection on moral hazard and risk aversion can explain the greater preference for more generous plans. However, moral hazard also enters the portion of the likelihood that matches observed to predicted spending, as described above.

the plan choice of their employer. We report these fees in Appendix Table [A6](#). We estimate the fees to be positive and large—between \$97 and \$110 per household per year on average—for non-silver plans, suggesting that employers choose silver plans at a greater frequency than households would prefer if choosing independently. In our counterfactual analyses, we compute pre-integration small group consumer surplus both with and without these fees.

7.2 Cost estimates

Table [6](#) contains the estimates of the premium setting equation, Equation [6.6](#), for the sample of individual insurance plans. 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.

In section [6](#), we described how our estimated coefficient on claims costs in the premium setting equation allows us to capture risk-related transfers that insurers either pay to or receive from government programs, particularly the risk corridors and reinsurance programs. If the coefficient exceeds one, our estimates also capture a plan-level markup the carrier adds to its premiums above its total health spending. 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 in the individual market. This pass-through rate suggests that, after accounting for cost-sharing subsidies, additional risk-related transfers make up the balance. If there is a plan markup, it is more than offset by these factors.

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 Model fit

Finally, we test the fit of our model using a hold-out sample. We re-estimate our model on an 80% random sample of individual market households. Using these estimates, we predict the plan choices and spending levels of the remaining 20% of households and compare them to the observed levels. We report this comparison in Appendix Table [A9](#). Our estimated shares line up well with the observed shares by metal tier, albeit with a slightly larger predicted share for the outside option of 56% vs. 49% in the hold-out sample. Correcting for outliers, our model prediction also matches the observed health spending pattern by plan generosity. We predict slightly higher average monthly spending per plan in part because, as noted above, we over-predict the share of healthy consumers who choose to be uninsured.

8 Policy discussion

We conduct three successive counterfactual simulations that use our estimated model to evaluate the costs and benefits of pooling households across market segments. Our goal is to approximate new federal regulations that allow employers to offer tax-advantaged reimbursement accounts that employees can use to pay for medical expenses and individual insurance premiums. We provide more details of the regulation of these accounts, known as Individual Coverage Health Reimbursement Arrangements (ICHRA), in Appendix [A](#). In short, small group employers no longer offer a traditional group insurance option. Instead, they provide funds for employees to purchase coverage individually through a health insurance marketplace.

8.1 Simulation approach

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 enrollees in Oregon in the year 2016, two years after

the regulation of both the small group and individual insurance markets under the ACA. We use the full sample of small group enrollees for our simulation; in our estimation, we relied on the preferences of small group households forced to switch out of group coverage³⁴

To conduct the counterfactual simulations, we need a method to determine equilibrium premiums when the potential enrollee pool changes. Starting with the assumption of a 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 inclusion of these behavioral types ensures that all contracts are traded³⁵. 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. In our policy simulations, we apply this methodology accounting for net government payments to insurers. Later, we also consider simulations with fixed markups above zero.

8.2 Counterfactual Results

Our baseline 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 premiums³⁶. We report two measures of equilibrium 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 premiums that households face after subtracting government subsidies and after accounting for the distribution of choices by age, family size, and region. We distinguish market outcomes from

³⁴In our data, we cannot identify households who are offered a small group plan but choose not to enroll. Thus, in our simulation analysis we can only include employees who opted for small group coverage.

³⁵Our reported simulations assume 1% of the sample are behavioral types. Increasing this share to 5% has very little effect on the results.

³⁶Appendix F.3 describes the bootstrap procedure we use to compute standard errors for our counterfactual simulations.

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.

Simulated market outcomes. Before turning to our 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 plans³⁷ 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.³⁸

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 a reduction in adverse selection for the individual market. In addition, both the government and employers experience savings in this environment.

These effects reflect the low non-discretionary spending among small group employees. Some of these households 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 healthier small group households will choose to be uninsured when forced to pool. We estimate that 61% of small group households will exit coverage; those exiting have expected non-discretionary spending of \$240/month versus \$671/month for insured households.

Because 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 individual market enrollee, Tables 7 and 8 show that market shares and premiums for other tiers experience only small changes. Overall, the

³⁷The predicted mean gold premium is different from the raw data summarized in Table 2 for two reasons. First, Table 2 covers the years 2014-2016 while our simulations consider only 2016. Second, in our equilibrium we do not allow insurance plans to experience losses net of risk-related transfers, leading to higher predicted premium levels.

³⁸Consumer 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.

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. At least in Oregon in the time period of our sample, the size of the choice set is not a key driver of surplus changes. The number of carriers offering plans in the small group and individual markets at baseline are similar—an average of 8 vs. 7 carriers, respectively, in the urban regions. 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.³⁹ For the average household, our baseline model suggests the gains from the new market design outweigh its costs. Consumer surplus for small group market households increases by \$218 from the baseline estimate. In our second specification, which we outline in Section 6.4, we allow for fees in the small group market before pooling. This version necessarily generates a larger consumer surplus change for small group enrollees, here equal to \$311.

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 exempt from taxes. Thus, we measure government expenditures for a household as the household’s average tax rate multiplied by the observed premium of its 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. Part of this decrease

³⁹The 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.

reflects the prediction that a meaningful portion of small group households choose to be uninsured in the pooled environment.

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. Under that assumption, the counterfactual would yield approximately \$419 million in yearly savings for small employers in this sample. However, in practice, wages would likely 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 using tax-exempt dollars that are partly subsidized by employers. We find that this extended 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. For small group employees, average gross premiums fall in all three metal tiers, and employees frequently forego insurance coverage when pooled, as they did in the previous counterfactual. Using our baseline model we find that small group consumer surplus increases substantially, with average gains of \$264 per month, given the added choice to exit coverage.

Finally, we explore how the ICHRA policy change may affect government expenditures. Integrating small group enrollees leads to lower predicted premiums in all metal tiers in the individual market, and thus lower government subsidy levels.

Under current ICHRA regulations, small group households are not eligible for government subsidies in the individual market but benefit from a tax exclusion for (a) their employer's contributions

to the ICHRA account and (b) their spending on health insurance premiums over the level of the employer’s contributions.⁴⁰ Thus, both before and after integration, the total premium paid is tax deductible. We assume the employer holds fixed their expenditures per insured employee after integration. That is, if the employer contributed \$500/month before integration, and post-integration the premium falls below \$500, we assume the employer will continue to fund the employee’s ICHRA account at \$500, covering some out-of-pocket medical expenses as well as premiums. Our value of tax expenditures thus accounts for the total contribution to ICHRA accounts as well as any additional tax-exempt wage dollars an employee spends on premiums above the ICHRA level. In Table 7 we find government expenditure falls \$35 or 28% for small group households with an ICHRA account. Here, the savings come from pre-merger households for whom tax-exempt premium payments exceeded the employer contribution. Those households paid the excess directly using pre-tax cafeteria plans. When premiums fall, the tax expenditure on these amounts falls.

We caution that some of the welfare gains we find for small group households arise when a fraction of these households choose uninsurance. Households that choose to be uninsured 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 underestimated 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 employers offer ICHRA accounts and employees face a mandate for coverage.

Extending ICHRA coverage with mandatory enrollment. The final column of Tables 7 and 8 adds an employer coverage mandate to our simulated environment with ICHRA accounts. We now assume that small employers can induce all previously insured 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

⁴⁰As we describe in Appendix A employees with ICHRA accounts can use pre-tax dollars for insurance premiums or medical spending above the employer’s contributions to the account. In such a case, they would transfer a portion of their salary to an employer’s cafeteria plan and pay for their health spending with these funds.

per month; the share of 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%) under the lower equilibrium prices.

These new equilibrium premiums and plan choices generate a larger improvement in consumer surplus for individual market households—an increase of \$44 per month. Our baseline measure of small group consumer surplus rises as well, by an average of \$58 per month. However, small group consumer surplus gains are significantly smaller than in alternative counterfactuals because some households prefer to be uninsured. In terms of government expenditure, spending on subsidies for individual market enrollees again falls, now by 22% or \$15/month. Among small group households, government expenditure falls by \$37/month.

Who wins and who loses? While our third counterfactual generates overall surplus gains for both individual market and small group enrollees, it may not make all consumers better off. We now explore the distributional consequences of pooling, quantifying how pooling affects households that differ in their underlying health needs and income.

We report our baseline surplus estimates by household type in the first two columns of Appendix Table A8. In the individual market, we find the biggest surplus gains accrue to the sickest households: households in the first vs. fourth quartile of health severity see a surplus gain from pooling of \$32/month vs. a gain of \$84/month. For small group enrollees, households gain at all health levels because of the lower premiums, but the sickest households see smaller gains. Three factors explain this differential by sickness level. First, when pooling, sicker small group households lose access to the most comprehensive platinum plans, which are not available in the individual market. Second, when in a small employer group, sicker households benefit because their premiums reflect, in part, the health of their co-workers; in the individual market, their premiums are decoupled from group members and reflect only indirectly the broader population in the market. Finally, sicker households are more likely to choose gold plans, which face adverse selection in the individual market. As shown in Table 8, the percentage decline in premiums for gold plans is smaller than the decline for other metal tiers.

Surplus changes correlate with income in a similar way. The highest income households in the individual market tend to choose more comprehensive plans and receive no government subsidies, and thus benefit more from the reduction in equilibrium premiums after pooling. In the small group market, the richest consumers choose more comprehensive plans and thus benefit the least from pooling.

Finally, we compare surplus levels across household types in an environment in which individual market insurers respond to pooling by adding a markup of 25%. As we report in the third and fourth columns of Appendix Table [A8](#), the sickest small group consumers and those in the highest income brackets now lose surplus when the markets are pooled. For the more comprehensive plans these households choose, the pooled premiums approach the premium levels the households faced in the pre-pooling small group market. Combining this smaller savings in premiums with both the loss of platinum plans and the decoupling from healthier households within an employer group, the sicker households and those households seeking more comprehensive plans lose surplus.

9 When is pooling beneficial?

Our counterfactual simulations assess the impact of pooling markets using the preferences of consumers in Oregon in 2015-16 and the insurer conduct we observe in the same period. In this section we develop a broader policy analysis, investigating the effect of (a) changing the characteristics of the consumer population—including their moral hazard, risk aversion and sickness severity—and (b) changing insurer conduct. While we stay within the neighborhood of the data so that our modeling assumptions likely remain valid, our goal is to assess the robustness of our results to changes in market conditions and to quantify the trade-offs from pooling in different environments.

9.1 Varying consumer characteristics

We start by considering how the value of pooling may depend on household characteristics. Panels A and B of Figure [5](#) illustrate how consumer surplus among individual market and small group households, respectively, change with the underlying health state of the two populations. On each

figure, horizontal and vertical lines denote the average level of non-discretionary spending in our baseline population, with their intersection reflecting the baseline surplus gains from pooling.

Perhaps not surprisingly, we see individual market households gain more from pooling when the small group population is healthier and/or the individual market population is sicker. From the perspective of small group households, the gain in surplus is higher when the small group is healthier. In addition, in the counterfactuals we show a point at which individual market and small group households have the same spending needs. We emphasize that at that point in the figure, the consumer surplus gains for both populations remain positive and economically significant. This finding highlights the importance of modeling demand, including both moral hazard and risk aversion, rather than looking only at spending. Small group households prefer more comprehensive plans even when their non-discretionary spending needs are equal to or below the level of individual market households. As a consequence, pooling similarly sick or even slightly sicker small group households with individual market consumers could lower plan-level adverse selection in comprehensive plans. We find that individual market households begin to lose surplus under pooling only when small group households have average spending that is 30% higher than the level of individual market households.

Figure 6 illustrates the role of moral hazard in the change in surplus from pooling. We see that the change in individual market household surplus from pooling grows when small group households have lower moral hazard. Fixing the small group moral hazard, individual market households gain more from pooling when they themselves exhibit higher moral hazard: pooling with lower spending small group households brings down equilibrium premiums for plans that the higher spending individual market households select. For small group households, their gains from pooling are higher when individual market households exhibit lower moral hazard. They also benefit less from pooling when their own moral hazard is higher, because pooling will not generate as substantial a saving in premiums.

On risk aversion, we find relatively smaller effects on surplus. All else equal, increasing the small group consumers risk aversion, within the range of estimates found in the empirical literature (Marone and Sabety, 2022), generates larger individual market consumer surplus gains. Changing

the small group risk aversion in a small range from a gamble of \$99.6 to a more risk averse gamble of \$97.4, individual market households gain \$0.11/month more from pooling. The added gains arise because healthier households choose more comprehensive plans when risk averse, leading to lower premiums for these plans in equilibrium.

9.2 Varying insurer conduct

Next, we consider a scenario in which insurers respond to pooling by charging positive markups in the pooled environment. 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) regulation⁴¹. We conduct the exercise under our third counterfactual in which small employers contribute to ICHRA accounts and employees must purchase at least bronze-level coverage.

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.⁴² In addition, our assumption 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 of 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.

Changes in surplus for individual market consumers The effect of markups alone is simple: when premiums increase due to markups, the level of consumer surplus for individual market households falls. To quantify the overall surplus change from both pooling and markups, we design

⁴¹Under 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.

⁴²An alternative to our approach would be to apply the best-response iteration algorithm of Hastings, Hortaçsu, and Syverson (2017). 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.

our measurement exercise to compare the surplus gains from pooling with healthier enrollees to the losses 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: when small employers offer ICHRA policies with a coverage mandate, the standardized premiums for all plan types fall with pooling, generating an average gain of \$44 in surplus for individual market households. When we allow positive markups in the post-integration environment, equilibrium premiums increase relative to the competitive pooled environment. However, the pooled premiums with markups may still be lower than in the baseline environment with neither pooling nor markups. In fact we find that bronze and silver plan premiums exceed the initial equilibrium level under 25% markups but not 10% markups. Gold plan premiums, in contrast, remain lower under any markup.

We translate these changes in equilibrium premiums into changes in average consumer surplus. Relative to the surplus gain of \$44 for individual market households in the competitive setting, adding 10% markups lowers the gain from pooling to \$32 in surplus. With 25% markups, individual market households gain an average of \$18 from pooling. That is, the cost of even the maximal markup fails to outweigh the benefit of adding healthier enrollees in our sample. We illustrate this result graphically in Figure 7. Under any post-pooling markup shown on the x-axis, the surplus change remains strictly positive.

How much might these results change if we considered markets or time periods where insurer markups in the individual market exceed 0% before pooling? This can only *increase* the benefits to individual market enrollees from pooling, whatever our assumption about insurer conduct after pooling. Figure 7 plots a line reflecting larger surplus changes in an individual market that has 10% markups before pooling. As in the main case above, after pooling with a healthier population, equilibrium premiums continue to fall enough to compensate for even the maximal post-pooling markup.

Changes in surplus for small group consumers For households formerly insured directly in the small group market, switching to an ICHRA account with a mandate to buy an individual plan generated a \$58 gain for the average household with no markups in the pooled environment. Ex-

amining the effect of positive post-pooling markups helps us understand the different sources of the original welfare gains. First, small group households no longer face ‘tiered composite’ community rated premiums in the pooled market. Eliminating composite rating can explain \$8 of the difference in surplus. Second, premiums in the small group market exceed those in the individual market for the same coverage level, reflecting higher markups over medical costs, either because of weaker competitive effects or because of higher administrative costs, notably in the form of brokers’ fees.

The Kaiser Family Foundation (2020) reports a \$9.88 per-member per-month differential between broker fees in the small group and individual insurance markets in Oregon in 2016.⁴³

With ICHRA accounts, we find small group households gain surplus from pooling even when there are strictly positive markups in the pooled individual market. Relative to the baseline gain of \$58 in surplus, if we layer on markups of 10% and 25% in the individual market, the gain in surplus to small group consumers falls to \$46 and \$29, respectively. The gains remain in part because lower administrative costs in the individual market can translate into lower premiums. Given the average household size of 1.54 members in the small group in 2016, the differential in average broker fees alone can explain \$15.20 of the surplus gains.

Government Expenditures Finally, in Table 9 we also explore how the ICHRA policy change may affect government expenditures under various post-pooling markups. Under this policy, pooling small group enrollees leads to lower premiums in all metal tiers in the individual market, and thus lower subsidy levels. Even with the maximum post-integration markup allowed, government expenditure on the individual market is still below baseline by \$8 per month or 12%. Government tax expenditure on small group enrollees falls from \$124 to \$106 with the maximal markup, as realized premiums on average still fall below the pre-pooling small group market environment. As a result, fewer employees pay premiums above their employer’s contribution, reducing tax expenditures.

⁴³The 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.

10 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 sickness levels 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.

Our counterfactual analyses identify particular market pooling policies that can improve welfare for households in both market segments. 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 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 feasible markups.

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 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. 2019. “Plan Choice and Affordability in the Individual and Small-Group Markets: Policy and Performance—Past and Present.” *Health Affairs* 38 (4):675–683.
- Agency for Healthcare Research and Quality. 2016. “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.”
- Atal, Juan Pablo, Hanming Fang, Martin Karlsson, and Nicolas R. Ziebarth. 2020. “Long-Term Health Insurance: Theory Meets Evidence.” Working Paper 26870, National Bureau of Economic Research.
- Azevedo, Eduardo M. and Daniel Gottlieb. 2017. “Perfect Competition in Markets With Adverse Selection.” *Econometrica* 85 (1):67–105.
- Blumberg, Linda J., Shanna Rifkin, Sabrina Corlette, and Sarah J. Dash. 2013. “Stabilizing Premiums Under the Affordable Care Act: State Efforts to Reduce Adverse Selection.” Tech. rep., The Urban Institute.
- Brot-Goldberg, Zarek, Amitabh Chandra, Benjamin Handel, and Jonathan Kolstad. 2017. “What does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics.” *The Quarterly Journal of Economics* 132 (3):1261–1318.
- Bundorf, M. Kate, Jonathan Levin, and Neale Mahoney. 2012. “Pricing and Welfare in Health Plan Choice.” *American Economic Review* 102 (7):3214–48.
- Claxton, Gary, Matthew Rae, Michelle Long, Nirmita Panchal, Anthony Damico, Kevin Kenward, and Heidi Whitmore. 2015. “Employer Health Benefits: 2015 Annual Survey.” Tech. rep., The Kaiser Family Foundation.
- Cutler, David M. and Sarah J. Reber. 1998. “Paying for Health Insurance: the Trade-off Between Competition and Adverse Selection.” *The Quarterly Journal of Economics* 113 (2):433–466.

- Diamond, Rebecca, Michael J. Dickstein, Timothy McQuade, and Petra Persson. 2021. “Insurance without Commitment: Evidence from the ACA Marketplaces.” Working Paper 24668, National Bureau of Economic Research.
- Einav, Liran, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen. 2013. “Selection on Moral Hazard in Health Insurance.” *American Economic Review* 103 (1):178–219.
- Einav, Liran, Amy Finkelstein, and Pietro Tebaldi. 2019. “Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies.” Manuscript, Department of Economics, Stanford University.
- Fleitas, Sebastián, Caitlyn Fleming, Gautam Gowrisankaran, and Anthony T. Lo Sasso. 2021. “Selection, Subsidies, and Welfare in Health Insurance: Employer-Sponsored Health Insurance Versus the ACA Marketplaces.” Manuscript, Department of Economics, University of Leuven.
- Fleitas, Sebastián, Gautam Gowrisankaran, and Anthony Lo Sasso. 2020. “Reclassification Risk in the Small Group Health Insurance Market.” Working Paper 24663, National Bureau of Economic Research.
- Geruso, Michael, Timothy Layton, and Daniel Prinz. 2019. “Screening in Contract Design: Evidence from the ACA Health Insurance Exchanges.” *American Economic Journal: Economic Policy* 11 (2):64–107.
- Hall, Mark A. and Michael J. McCue. 2021. “Expanding Consumer Health Insurance Options by Easing the Boundaries Between Individual and Small-Group Markets.” Tech. rep., Commonwealth Fund.
- Handel, Benjamin. 2013. “Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts.” *American Economic Review* 103 (7):2643–2682.
- Handel, Benjamin, Igal Hendel, and Michael D. Whinston. 2015. “Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk.” *Econometrica* 83 (4):1261–1313.

- Handel, Benjamin and Jonathan Kolstad. 2015. “Health Insurance for ‘Humans’: Information Frictions, Plan Choice and Consumer Welfare.” *American Economic Review* 105 (8):2449–2500.
- Hastings, Justine, Ali Hortaçsu, and Chad Syverson. 2017. “Sales Force and Competition in Financial Product Markets: The Case of Mexico’s Social Security Privatization.” *Econometrica* 85 (6):1723–1761.
- Ho, Kate and Robin S. Lee. 2020. “Health Insurance Menu Design for Large Employers.” Working Paper 27868, National Bureau of Economic Research.
- Jaffe, Sonia and Mark Shepard. 2020. “Price-linked Subsidies and Imperfect Competition in Health Insurance.” *American Economic Journal: Economic Policy* 12 (3):279–311.
- Marone, Victoria R. and Adrienne Sabety. 2022. “When Should There Be Vertical Choice in Health Insurance Markets?” *American Economic Review* 112 (1):304–42.
- Polyakova, Maria and Stephen P. Ryan. 2019. “Subsidy Targeting with Market Power.” Working Paper 26367, National Bureau of Economic Research.
- Rae, Matthew, Cynthia Cox, Krutika Amin, and Tricia Neuman. 2021. “How Lowering the Medicare Eligibility Age Might Affect Employer-Sponsored Insurance Costs.” Tech. rep., The Kaiser Family Foundation.
- Tebaldi, Pietro. 2022. “Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA.” Working Paper 29869, National Bureau of Economic Research.
- Tebaldi, Pietro, Alexander Torgovitsky, and Hanbin Yang. 2023. “Nonparametric Estimates of Demand in the California Health Insurance Exchange.” *Econometrica* 91 (1):107–146.
- The Kaiser Family Foundation. 2020. “State Health Facts: Broker Compensation by Health Insurance Market.”
- Tikkanen, Roosa, Robin Osborn, Elias Mossialos, Ana Djordjevic, and George Wharton. 2020. “International Profiles of Health Care Systems.” Tech. rep., The Commonwealth Fund.

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

11 Tables and Figures

Table 1: Summary statistics on demographic variables

Variable	Individual market			Small group market		
	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 severity score	1.39	2.44	0.62	1.25	2.20	0.60
Income (share of 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 summary statistics on the demographics of households in Oregon who choose insurance plans in both the individual and small group markets in years 2014-2016. We compute the household health severity score as the sum of health status scores for members of a household, where we predict each member's health 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. We divide income by the federal poverty line (FPL) in a given year.

Table 2: Summary statistics on insurance variables

Variable	Individual market			Small group Market		
	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 summary statistics on the characteristics of individual and small group market insurance plans that households in Oregon choose 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 severity score [0/1]		-0.156	-0.190	0.208	0.185
		(0.002)	(0.002)	(0.016)	(0.016)
HH health severity 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		-2.186	-1.504	-1.505	-1.154
		(0.010)	(0.006)	(0.069)	(0.053)
ψ					
Constant		-2.685	-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}					
Number of		\$20	\$50	\$20	\$50
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) $\underline{c} = 50$. Low household health severity score [0/1] is an indicator equal to one if a household’s health severity score is below the 30th percentile of the distribution of scores. We compute the household severity score as the sum of health status scores for members of a household, where we predict each member’s score using the 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

Household type	Individual market			Small group market		
	Mean	S.D.	Share	Mean	S.D.	Share
$E[\lambda_i]$						
Overall	5.60	13.94		4.19	11.73	
No dependent, under-50	2.40	7.77	0.46	1.76	5.70	0.59
With dependent(s), under-50	6.18	13.68	0.12	7.04	15.00	0.13
No dependent, over-50	8.32	17.20	0.38	6.60	15.03	0.24
With dependent(s), over-50	14.83	22.72	0.04	15.97	23.93	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

Household type	Main		Switchers	
	(1)	(2)	(1)	(2)
ω_i				
Constant	0.112 (0.001)	0.222 (0.001)	0.222 (0.015)	0.315 (0.017)
ψ_i				
Constant	0.068 (0.001)	0.077 (0.001)	0.267 (0.007)	0.276 (0.007)
Loss interpretation of ψ_i				
Constant	99.322 (0.008)	99.232 (0.008)	97.402 (0.063)	97.316 (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 (0.026)	0.737 (0.027)	0.733 (0.024)	0.547 (0.054)
Administrative costs (t-1)				0.618 (0.166)
Year FEs				✓
Payer-Year FEs		✓	✓	
N	240	238	238	186
ψ				
1 st -stage F-stat			18.328	4.947
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. Appendix [B](#) describes the construction of a plan's predicted administrative costs.

Table 7: Counterfactual results: outcomes

Panel A: Individual market (N = 178,157)				
Counterfactual:	Base	Close small group	Extended ICHRA	Mandated insurance
<i>Welfare</i>				
Consumer surplus, change from base	0 (0)	18 (1)	28 (2)	44 (2)
Government expenditure	67 (0)	60 (0)	57 (0)	52 (0)
<i>Market shares</i>				
Uninsurance	56.5 (0.1)	54.4 (0.2)	53.1 (0.2)	50.7 (0.3)
Bronze	16.6 (0.0)	16.1 (0.1)	15.9 (0.1)	15.6 (0.1)
Silver	20.1 (0.1)	20.4 (0.1)	19.6 (0.1)	19.9 (0.1)
Gold	6.8 (0.1)	9.0 (0.2)	11.4 (0.3)	13.7 (0.2)
Panel B: Small group market (Base: N = 55,374; Merged: N = 125,527)				
<i>Welfare</i>				
Consumer surplus, change from base	0 (0)	218 (13)	264 (11)	58 (4)
Government expenditure	124 (0)	71 (1)	89 (1)	87 (1)
<i>Market shares</i>				
Uninsurance	0.0 (0.0)	61.2 (0.5)	53.9 (1.1)	0.0 (0.0)
Bronze	13.7 (0.0)	9.7 (0.2)	13.8 (0.3)	28.4 (0.4)
Silver	48.8 (0.0)	17.9 (0.4)	17.0 (0.4)	35.4 (0.1)
Gold/Platinum	24.8 (0.0)	11.2 (0.2)	15.3 (0.5)	36.1 (0.4)

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 at the monthly level. We calculate small group “base” numbers using observed choices and premiums. All other numbers are predicted using the counterfactual algorithm. To account for outliers, we report consumer surplus 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”. We compute standard errors using a bootstrap procedure described in Appendix [E.3](#).

Table 8: Counterfactual results: premiums

Panel A: Individual market (N = 178,157)					
	Counterfactual:	Base	Close small group	Extended ICHRA	Mandated insurance
Standardized gross premiums					
	Bronze	177 (1)	175 (1)	172 (1)	163 (2)
	Silver	219 (1)	205 (2)	208 (1)	195 (2)
	Gold	685 (10)	458 (13)	354 (10)	276 (6)
Population net premiums					
	Bronze	226 (2)	227 (2)	224 (2)	215 (2)
	Silver	294 (2)	277 (2)	287 (2)	275 (3)
	Gold	1,379 (24)	839 (27)	632 (19)	486 (12)
Panel B: Small group market (Base: N = 55,374; Merged: N = 125,527)					
Standardized gross premiums					
	Bronze	250 (0)	173 (2)	170 (1)	161 (2)
	Silver	312 (0)	199 (2)	200 (1)	186 (2)
	Gold	392 (0)	442 (15)	320 (11)	245 (6)
Population net premiums					
	Bronze	111 (0)	162 (2)	1 (0)	0 (0)
	Silver	136 (0)	197 (3)	8 (0)	3 (0)
	Gold	172 (0)	784 (33)	162 (18)	49 (6)

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 at the monthly level. We calculate small group “base” numbers 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. We average over these standardized premiums weighting by plan choice probabilities. Population net premiums reflect the average premium households face, accounting for age adjustments and subsidies and weighting by plan choice probabilities. We compute standard errors using a bootstrap procedure described in Appendix [E.3](#).

Table 9: Counterfactual results: varying markups

Panel A: Individual market (N = 178,157)				
Counterfactual:	Base	Mandated insurance		
Markup (in %):	0	0	10	25
<i>Welfare</i>				
Consumer surplus, change from base	0	44	32	18
Government expenditure	67	52	55	59
<i>Standardized gross premiums</i>				
Bronze	177	163	171	181
Silver	219	195	211	229
Gold	685	276	329	412
Panel B: Small group market (Base: N = 55,374; Merged: N = 125,527)				
<i>Welfare</i>				
Consumer surplus, change from base	0	58	46	29
Government expenditure	124	87	94	106
<i>Standardized gross premiums</i>				
Bronze	250	161	168	176
Silver	312	186	197	210
Gold	392	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 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, we report consumer surplus 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.

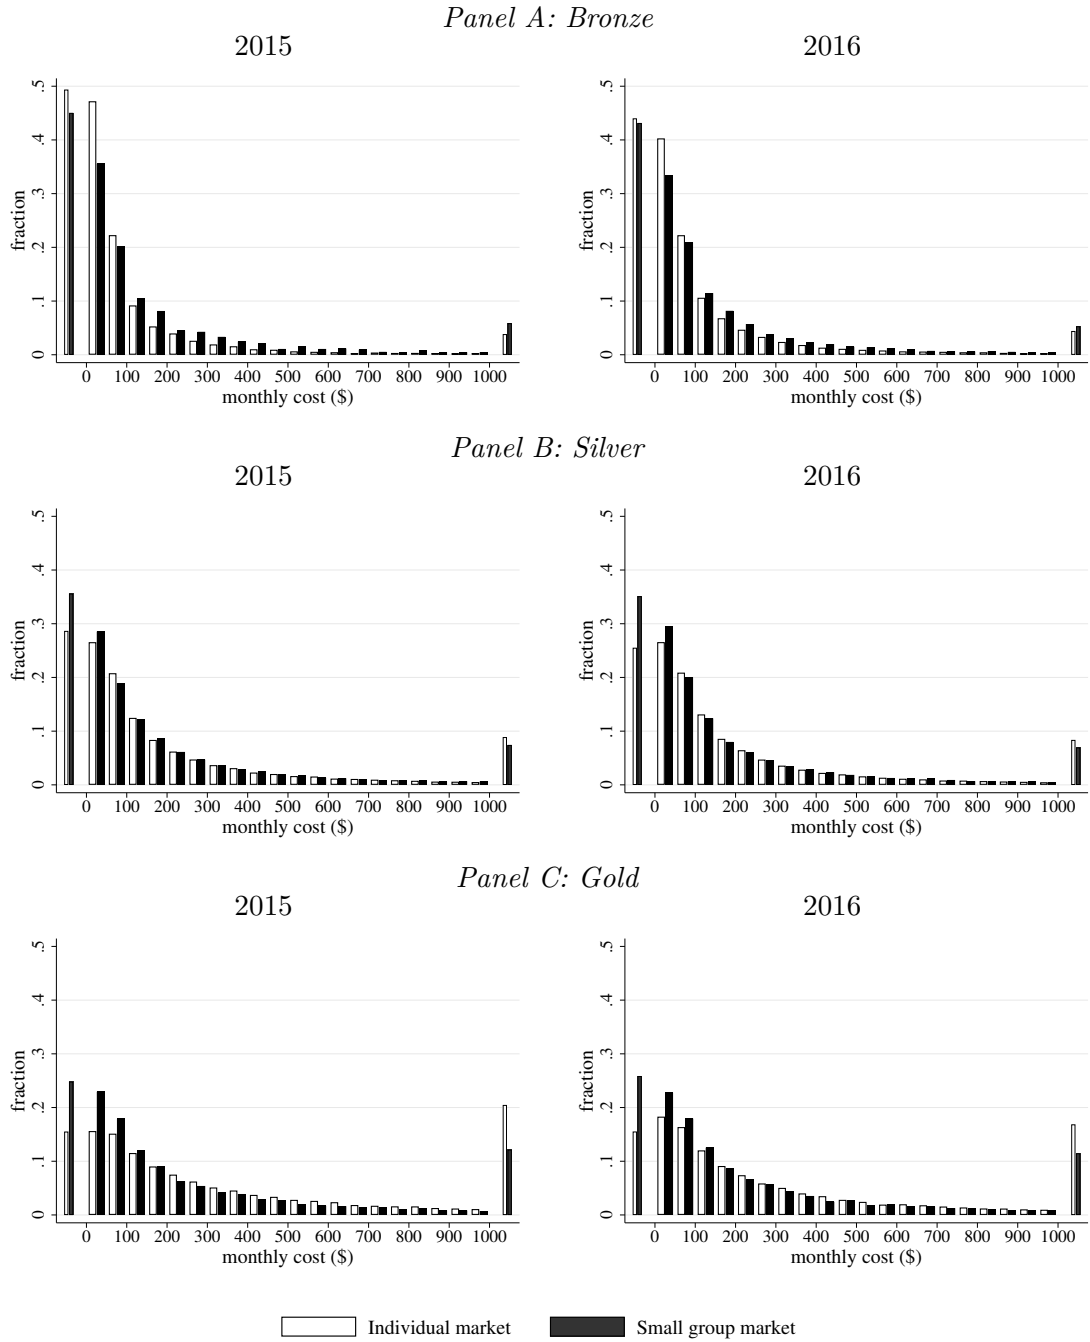


Figure 1: Distribution of monthly medical costs

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 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.




Figure 3: Medical markups

Note: This figure depicts the distribution of medical markups across payers, where we define 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. 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.

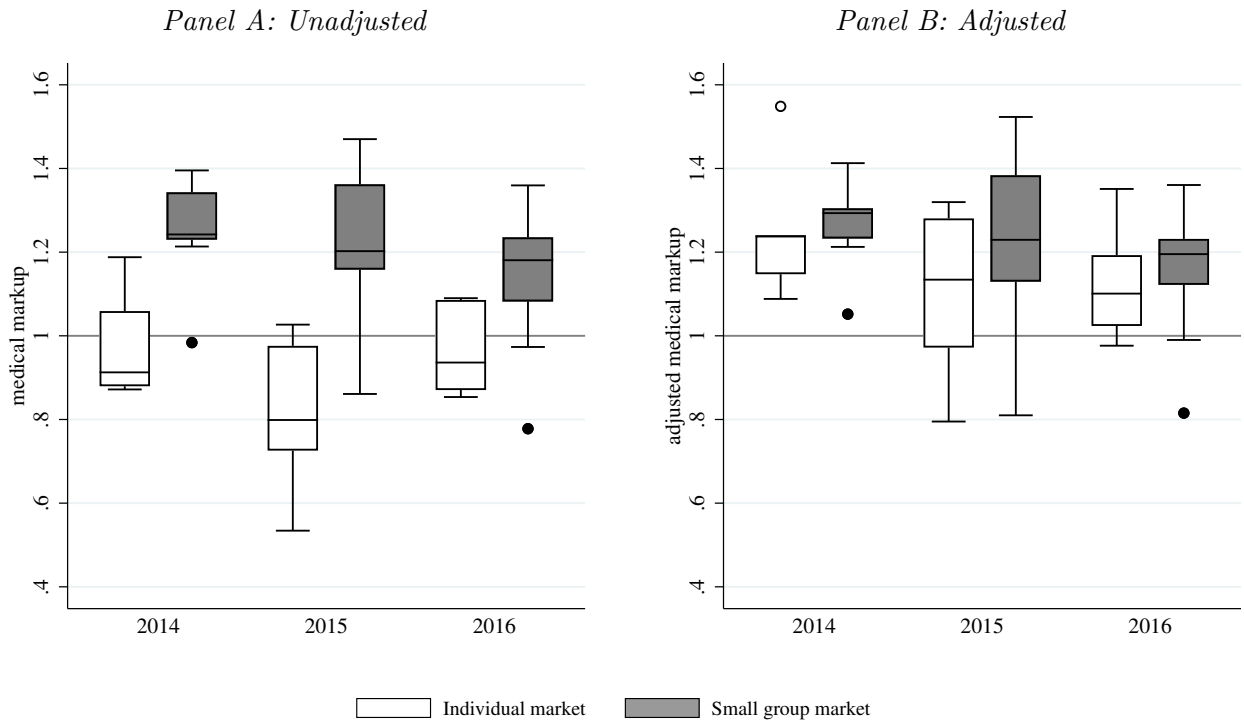


Figure 4: Insurer-level medical markup adjusted for government transfers

Note: This figure depicts the distribution of medical markups across payers. We calculate the medical markup as the ratio of a payer's total premium revenue divided by the total medical cost the payer incurs. In panel A, the total premium revenue is not adjusted. In panel B, we adjust the total premium revenue to account for risk adjustment, risk corridor, and reinsurance payments to the insurer by year and market.

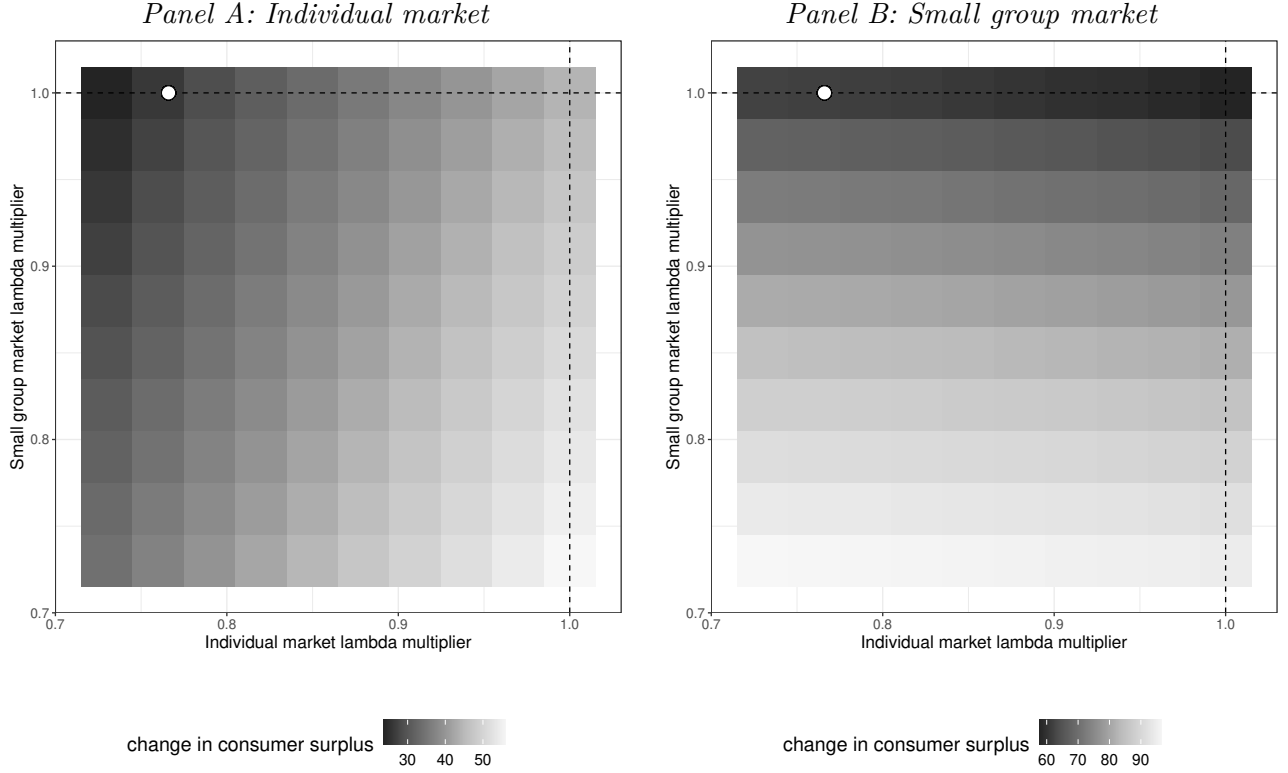


Figure 5: Change in consumer surplus with expected non-discretionary spending

Note: This figure depicts the change in consumer surplus from pooling relative to the baseline when varying average non-discretionary medical spending. The figure considers the counterfactual outcome of having extended ICHRA coverage and mandated insurance. Panel A shows the change in surplus for individual market households while Panel B depicts the same change for small group households. The change in consumer surplus is averaged over households and reported in dollars per month. The dotted lines represent the baseline multiplier values for the small group on the y-axis and the individual market on the x-axis (both of which equal one). The dot represents the point where expected non-discretionary medical spending is equivalent in the individual and small group markets. To account for outliers, we report consumer surplus for the set of households where, at baseline, $\alpha - \psi > 0.05$.

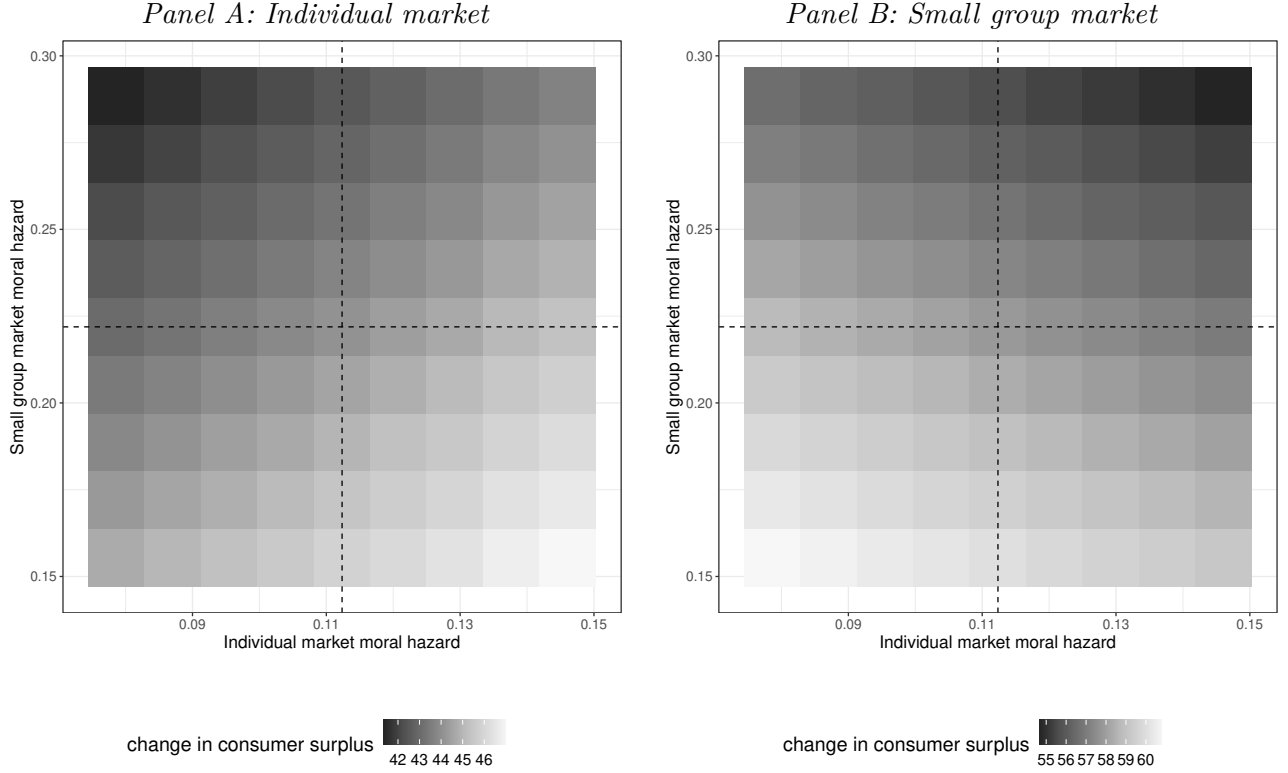


Figure 6: Change in consumer surplus with moral hazard

Note: This figure depicts the change in consumer surplus from pooling relative to the baseline when we vary the moral hazard parameter in the small group and individual market populations. The figure considers the counterfactual outcome of having extended ICHRA coverage and mandated insurance. Panel A shows the change in surplus for individual market households while Panel B depicts the same change for small group households. The change in consumer surplus is averaged over households and reported in dollars per month. The dotted lines represent the baseline multiplicative moral hazard parameter values for the small group on the y-axis and the individual market on the x-axis. To account for outliers, we report consumer surplus for the set of households where $\alpha - \psi > 0.05$.

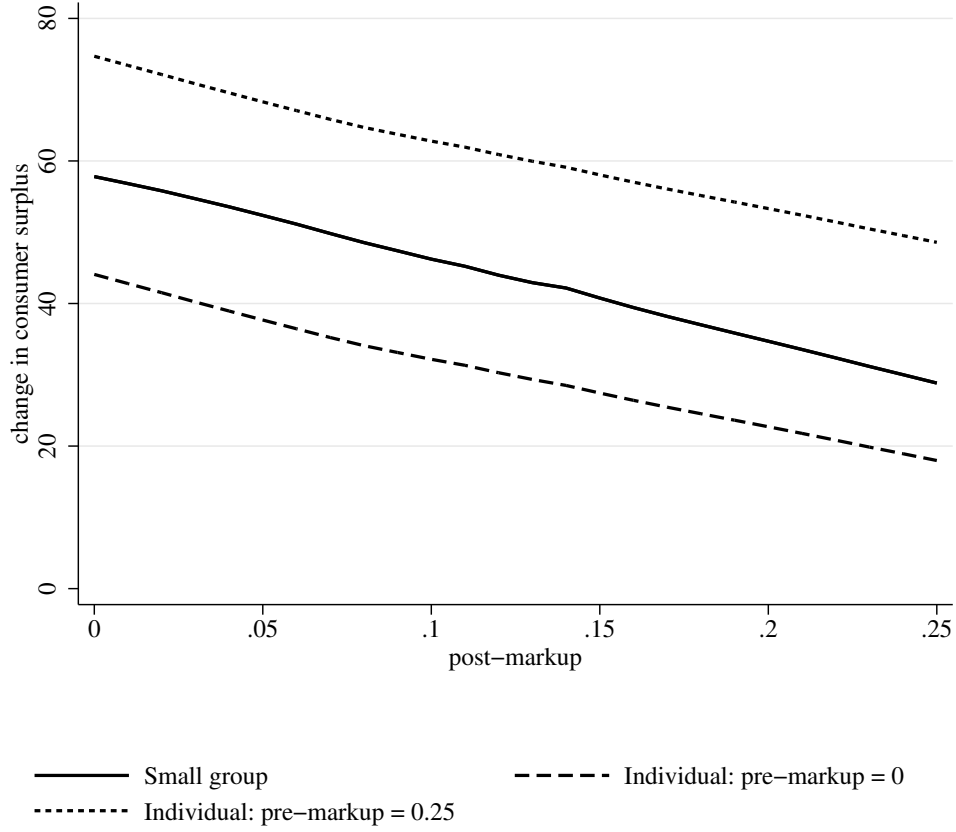


Figure 7: Change in consumer surplus and markups

Note: This figure depicts the effects of varying a fixed markup parameter on the counterfactual change in surplus under pooling. Change in consumer surplus is averaged over households and reported in dollars per month. The figure considers the counterfactual outcome of having extended ICHRA coverage and mandated insurance. We include the cases where the individual market has no baseline markup and a baseline markup of 25%. To account for outliers, we report consumer surplus for the set of households where $\alpha - \psi > 0.05$. Small group baseline numbers are calculated using observed choices and premiums. All other numbers are predicted using the counterfactual algorithm. In the small group market baseline, we omit households who choose grandfathered or other plans for which we do not observe premiums or metal tier in the data. In the counterfactual, we can simulate premiums and metal tiers for all households, including those omitted in the baseline.