Community Rating and Vertical Price Distortions in Insurance Menus*

Honglin Li Chenyuan Liu Justin Sydnor

February 2026

Abstract

How does community rating distort relative prices for vertically differentiated insurance plans? We show both theoretically and empirically that common community-rating approaches can distort price differentials in the opposite direction as distortions in overall price levels. Empirically, we show that partial age-based community rating in U.S. private health insurance exchanges generates marginal prices for generous coverage for older individuals significantly above marginal cost, often creating dominated options, while younger enrollees are subsidized on the margin. We present a theory on how these pricing distortions impact the efficiency of sorting across plans and empirical evidence of their distributional impacts. *JEL Codes*: G22, G52, I13.

*Contact information: Honglin Li (hli434@wisc.edu), Chenyuan Liu (liuchy3@sem.tsinghua.edu.cn, corresponding author), Justin Sydnor (justin.sydnor@wisc.edu). We thank Stuart Craig, Jonathan Gruber, Tim Layton, Anran Li, Dan Sacks, Mark Shepard, the workshop participants at the University of Wisconsin, Madison, the Munich Risk and Insurance Center at LMU, the 23rd Annual International Industrial Organization Conference, and the 4th Annual Risk Management & Insurance Research Symposium at Florida State University for their helpful comments. Chenyuan Liu receives financial support from the National Natural Science Foundation of China (project number 72203120). The funding organization had no role in the study's design, data collection, analysis, publication decision, or manuscript preparation. The authors declare that all findings and interpretations are entirely their own and have not been influenced by the funding entity.

1 Introduction

Community rating is a common feature in many insurance markets. Motivated sometimes by fairness concerns, desires for redistribution, and to mitigate reclassification risk, community-rating regulations limit the extent to which insurance prices reflect variation in individuals' expected risk. As a consequence, community rating generally lowers the price of insurance contracts for those with higher risk and raises the price for those with lower risk.

How does community rating affect insurance markets with multiple differentiated products? Prior literature has highlighted that community rating leads to inefficiencies in sorting into vertically differentiated plan options because it leads to price differentials between plans that do not reflect individual marginal costs for additional coverage (see e.g., Einav, Finkelstein and Cullen 2010; Bundorf, Levin and Mahoney 2012). Prior studies have focused on situations where the higher-risk types who are subsidized by community rating also face prices for additional coverage that are below their marginal cost, while the opposite is true for lower-risk types. Many studies have examined how these distortions can fuel adverse-selection dynamics (e.g., Cutler and Reber, 1998; Einav, Finkelstein and Cullen, 2010; Handel, 2013).

In this paper, we highlight a novel insight that it is possible for community-rating distortions in price *differentials* between product options to run in the opposite direction as the distortions in the overall *level* of prices. In other words, although community rating creates subsidies so that high-risk types face price levels that are below their expected cost (and vice-versa for low-risk types), it is possible that the premium gap between the more and less generous plans is higher (lower) than marginal cost for the high-risk (low-risk) types. We illustrate the point by developing a simple and general theoretical framework for analyzing community rating in a setting with two vertically differentiated plans. The framework focuses on multiplicative community rating, in which plan prices are regulated by a risk-rating factor that controls the ratio of prices that can be charged to higher-risk versus lower-risk types. We consider both the standard case of full community rating and the less-commonly studied situation of partial community rating. The framework can accommodate the effects of moral hazard responses and can be applied with various pricing institutions.[1]

The intuition for why community rating can generate price differentials that vary more by risk type than marginal costs lies in the interaction of community rating with nonlinear plan designs. With linear plans (e.g., constant coinsurance), marginal cost differences across risk

[1] We consider both highly regulated markets with risk-adjustment schemes and single-risk-pool pricing requirements (like ACA health insurance exchanges), and less-regulated competitive markets where prices adjust based on the average cost of enrollees. Our core insights are robust to these different pricing institutions.

types scale proportionally with total costs. With nonlinear designs—especially those featuring maximum out-of-pockets (MOOPs) —this proportionality breaks down. MOOPs mute the impact of higher losses on the marginal cost of the more generous plan. Since higher-risk individuals are more likely to have losses that are above the MOOPs across plans, their marginal costs will be affected more by this effect. As a result, marginal cost differentials may vary relatively little with risk. Community rating, by contrast, applies uniform pricing rules across plans, setting prices so that higher-risk types face a fixed multiple of lower-risk prices. Even when this multiple is smaller than the ratio of expected costs, it can exceed the ratio of marginal cost differentials, leading prices to vary excessively relative to marginal costs across risk types. While this pattern can occur at any level of community rating, it is more likely to occur with partial than full community rating.[2]

We show that these pricing distortions can have significant impacts on the nature of inefficiencies in sorting across plans. We adopt the common graphical framework from Einav, Finkelstein and Cullen (2010) (EFC) and revisit the analysis by Marone and Sabety (2022), extending the analysis to consider pricing distortions under partial community rating.[3] Marone and Sabety (2022) highlight that the efficiency of community rating depends on whether demand is steeper than marginal costs in the EFC framework. When demand is steeper than marginal costs, it is efficient for higher-risk types to sort into more generous coverage because their willingness to pay is above their marginal cost. When marginal cost is steeper than demand, in contrast, lower-risk types should efficiently sort into more generous coverage. Marone and Sabety (2022) show that full community rating can achieve efficient sorting when it is efficient for higher-risk types to purchase more generous coverage, but not when lower-risk types should efficiently sort into more generous coverage.

We find novel results when we extend this analysis to partial community rating. When partial community rating allows for only modest risk-rating across types, it generally has small but beneficial impacts on the efficiency of sorting because it moves prices closer to marginal costs. However, if our novel pattern of distortions occurs and partial community rating leads to price gradients across risk types that substantially exceed the slope of marginal costs, there can be significant impacts on the efficiency of sorting. In the case where higher-risk types should efficiently sort into more generous coverage, price gradients steeper than marginal costs will lead to highly inefficient patterns of sorting. High marginal prices for the higher-risk types dissuade

[2]This pattern happens under full community rating when the marginal costs for additional coverage are actually lower for higher-risk types than lower-risk types.

[3]Einav, Finkelstein and Cullen (2010) consider market equilibrium and efficiency under full community rating with no risk adjustment, while Geruso and Layton (2017) consider full community rating with perfect risk adjustment. Our contribution is to extend both cases by allowing partial community rating.

them from purchasing more generous coverage that they would otherwise value above marginal cost, and vice versa for lower-risk types. In contrast, when it is efficient for lower-risk types to purchase more generous coverage, these high price gradients support efficient sorting and can dramatically improve efficiency relative to full community rating.

We demonstrate the empirical relevance of these underappreciated pricing distortions in the setting of private health insurance plans in the United States. Since the Affordable Care Act (ACA) of 2010, private health insurance markets in the U.S. feature a combination of strong pricing regulations and multiple product options. This combination is also prevalent in health insurance markets in other countries, such as Switzerland and the Netherlands. The ACA established regulations that require community rating for plans but allow for prices to vary by geography, age, and, in some states, by smoking status. We focus attention on the pricing by age. Federal regulations establish age-rating factors that cap age-based price increases at a 3:1 ratio for 64-year-olds relative to 21-year-olds. When these regulations were passed in 2010, it was recognized that expected costs tended to be about five times higher for 64-year-olds relative to 21-year-olds, so that the 3:1 age-rating limit implied partial community rating that created a transfer from younger to older enrollees (Fontana, Murawski and Hilton, 2017).

We analyze the price differences that enrollees at different ages face for health plans with different levels of coverage. Insurers in ACA markets are required to offer plans in multiple coverage tiers that vary in the generosity of their coverage. Bronze plans have the highest levels of cost-sharing for consumers (e.g., higher deductibles) and must target an actuarial value (AV) covering approximately 60 percent of expected medical expenses for an average population established by the Center for Medicare and Medicaid Services (CMS). Silver, Gold, and Platinum plans target 70, 80, and 90 percent AVs, respectively. For our primary analysis, we use publicly available data from the California health insurance exchange, Covered California, that includes information on plan prices and individual-level plan enrollment decisions.[4] We combine these data with national-level medical-utilization data through the NBER's license with Truven Marketscan. We construct age-specific medical expenditure distributions and inflate them to match the Covered California observed premiums.[5] We then estimate the expected value of plan options at different ages using these distributions.

We show that while older enrollees face price levels below their expected cost under the

[4]While Covered California offers an attractive market for our primary analysis due to the availability of pricing and enrollment data at the individual level and standardized plan options, the basic findings are likely to hold in many other private health insurance exchanges. We provide evidence that the patterns in CA on dominated options discussed below appear to hold broadly outside of California as well.

[5]We use tier-specific adjustments to reflect moral hazard.

community-rating rules, they face price *differentials* (e.g., Bronze to Gold) that substantially exceed their expected marginal costs for more generous plans. For example, after adjusting the Truven Marketscan data to match the average cost of plans in California, we estimate that in 2017, if prices were based on age-level risk, Gold-tier plans in California should have been priced on average approximately $3,700 more per year than Bronze-tier plans for 64-year-olds.[6] In reality, the average observed annual price differential for 64-year-olds for Gold vs. Bronze plans was $4,300. For 30-year-olds, on the other hand, the additional price for Gold vs. Bronze-tier plans was about $300 less than the estimated fair premium differential for that age.

In situations where price levels are higher, the distortions are large enough that Gold-tier plans are likely *dominated* by Bronze-tier plans for many older enrollees. This tends to occur in geographies where price levels are higher or for plans offered by insurers with more costly plan designs (e.g., offering broad provider networks). We estimate that most enrollees 55 years and older in California have at least one set of plans in their rating area where the Gold-tier plan was dominated by a Bronze-tier plan offered by the same insurer with the same network coverage. Perhaps due to these high price differentials, older individuals are modestly less likely to enroll in Gold plans than individuals in the middle of the age spectrum. Yet, we estimate that among the roughly 10 percent of older enrollees who select Gold-tier plans, most of them select a plan that is dominated by a lower-tier option offered by the same insurer.

The surprising patterns in the ACA marketplace arise exactly due to the interaction between the community rating rules and the non-linearity of health insurance contracts. We estimate that for 25-year-olds, the Gold-tier plan covers approximately 80% of expected medical expenses, while the high-deductible health plan in the Bronze-tier covers only 54%. The reason is that these younger individuals are much more likely to have relatively low spending levels for which the Bronze-tier plan provides no coverage. In contrast, 64-year-olds are more likely to have high expenditure levels and hit the MOOP with both plan designs. We estimate much more similar AV levels of 85% for the Gold-tier and 71% for the Bronze-tier. Due to these differences, the expected coverage levels differ much more by age for Bronze-tier plans than they do for Gold-tier plans. However, the 3:1 age-factor regulation is applied in the same way to all plans. As such, the Bronze-tier plans involve significantly larger cross-age transfers than the Gold-tier plans. The result is that the marginal price for higher coverage tiers is inflated for older individuals

[6] We focus some of our analysis on the 2017 plan year since that was the most recent year before federal regulations governing the payment for "cost-sharing subsidies" (CSRs) were changed. Starting in 2018, the federal government stopped paying insurers for the additional cost of CSR plans insurers are required to offer to lower-income enrollees. Some states counteracted this policy by allowing insurers to adjust the price of their Silver-tier plans to accommodate. While this does not present a conceptual challenge to our analyses, it complicates the display of cross-tier price comparisons.

and compressed for younger individuals relative to the marginal costs across age groups.

Our theoretical framework highlights that the impact of these pricing distortions on the efficiency of plan selection depends on which types should efficiently sort into more generous plans if price differentials reflected marginal costs. It is difficult in this setting, though, to conclusively measure how much people in general, and specifically at different ages, value additional coverage. Choices are complex and likely influenced by significant choice frictions that are outside the conceptual framework, as evidenced in part by the fact that a nontrivial share of enrollees violate standard notions of dominance.

Given that reality, we turn our attention toward quantifying the potential market-level and distributional effects of the current community rating practice relative to two potential alternative approaches. We consider both risk-based pricing and full community rating as alternative approaches to regulating the relative prices of plan options. In both alternative scenarios, we consider a set of age-based taxes and subsidies that would maintain the overall degree of cross-subsidization of the system. In this way, we hold fixed the degree of cross-subsidization across ages on the *level* of overall insurance prices and focus on the impact on the price *differentials* across vertical plan tiers. We use a discrete-choice demand estimation model to examine how plan choices respond to different pricing-regulation schemes and determine new equilibrium allocations that account for tax and subsidy adjustments. Overall, our simulations suggest that the current system is moderately more (less) favorable for younger (older) enrollees than a risk-based alternative and substantially more (less) favorable for younger (older) enrollees relative to a full community rating system.

Our work contributes to a few streams of prior literature. First, our results contribute to a rich literature exploring how community rating distorts incentives and selection in insurance markets. Many studies have noted the potential issues that can arise with adverse selection distortions across plan options under community rating (e.g., Cutler and Reber 1998; Einav, Finkelstein and Cullen 2010; Bundorf, Levin and Mahoney 2012; Handel 2013; Geruso and Layton 2017). Most of this work focuses on full community rating. The exception is Veiga (2023), who studies selection dynamics under a continuum from full community rating to risk-based pricing in a single-product setting. Our consideration of partial community rating aligns with Veiga's approach; however, we emphasize that with vertical choices and non-linear plans, partial community rating does not necessarily create price gradients that fall between zero and full price discrimination. Instead, we demonstrate that partial community rating can result in an overshooting of price gradients compared to risk-based pricing. We demonstrate that this overshooting of marginal-cost gradients also has implications within the theoretical framework

of Marone and Sabety (2022) for the potential efficiency of offering vertical choices.

In other related work, Geruso et al. (2023) highlight that in markets with vertical insurance plans, policies needed to combat intensive-margin selection issues, such as risk adjustment, may conflict with policies to support extensive margin selection. Our results suggest some nuance to this tradeoff for partial community rating using common multiplicative factors. We estimate that the current partial community-rating approach in Covered California modestly increases (reduces) extensive margin participation in the market by younger (older) individuals relative to full community rating. Yet the distortionary effects on vertical decisions in this setting are more nuanced, as partial community rating does not simply reduce these distortions but rather reverses their direction.

Our analysis also identifies a new way that pricing distortions, and particularly dominated options, can arise in insurance menus. Prior research has primarily focused on how dominated options can arise either through deliberate decisions by menu designers or as the consequence of average-cost pricing under adverse selection (Handel, 2013; Bhargava, Loewenstein and Sydnor, 2017; Liu and Sydnor, 2022). Intuition would suggest that markets, such as the ACA health exchanges, with strong regulations and risk adjustment, would not have dominated options. Rasmussen and Anderson (2021) provides evidence of the emergence of dominant options on private health insurance exchanges as a consequence of distortionary changes in pricing practices in response to a government-induced cost shock. Our work demonstrates that dominant options can emerge more generally for older individuals under the ACA as a natural consequence of otherwise well-intentioned core pricing regulations. These findings emphasize the need for regulators to understand the full impact of pricing regulations, especially in environments with non-linear insurance contracts.

Finally, our results provide additional evidence of the potential value of helping both consumers and policymakers more easily compare the value and prices of insurance options. A number of prior studies have highlighted that people find it difficult to make informed comparisons across health insurance options without decision aids (Heiss, McFadden and Winter, 2010; Abaluck and Gruber, 2011; Schram and Sonnemans, 2011; Johnson et al., 2013; Kairies-Schwarz et al., 2014; Bhargava, Loewenstein and Sydnor, 2017; Samek and Sydnor, 2025; Abaluck and Gruber, 2023; Heiss et al., 2021; Handel et al., 2024). The fact that coverage-tier selection in the California insurance market is only modestly related to age - despite substantial variation in the effective marginal subsidy/surcharge for more comprehensive coverage - is likely further evidence that people do not understand how to compare insurance prices. The pricing distortions we document may be especially problematic for individuals who rely on simple heuristics when

making choices: older individuals might naturally assume their higher risk means the most generous coverage is appropriate for them, and vice versa for the young. Our work highlights the importance of designing market structures and providing decision support to avoid financially disadvantaging less sophisticated consumers.

2 Conceptual Framework

2.1 Model Setup

In this section, we develop a conceptual framework to examine how community rating affects the pricing of insurance contracts. We consider a setting where there is a continuum of individuals who differ in a risk type that can be indexed by an observable factor, such as age. Let θ denote each type (age) and, with a slight abuse of notation, also their expected total losses. For simplicity, we assume there is no moral hazard; thus, the expected total loss is the same under different insurance contracts. We assume that $\theta \in [\underline{\theta}, \bar{\theta}]$ with the CDF of $F(\cdot)$. We consider the simplest case for multiproduct markets, where there are two insurance plans, H and L, which are exogenously determined. Let m denote the plan index with $m = H, L$. The expected cost to the insurer of providing plan m to type θ is $\lambda_m(\theta)$. We assume that $\lambda_H(\theta) > \lambda_L(\theta),\ \forall\theta$, which implies that plan H has the higher expected cost for all types. In the context of health insurance with plans vertically differentiated by coverage level, this assumption implies that the expected covered expenditures are higher under plan H for all types, or, equivalently, that plan H has higher "generosity" for all types.

The actuarially fair premium of plan m for type θ would charge the type-specific expected cost, i.e.,

$$p_m^{fair}(\theta) = \lambda_m(\theta). \tag{1}$$

Now we consider the implementation of community rating. In line with empirical realities in many markets, we assume that community rating is implemented by controlling the extent to which individual prices vary around the full-population expected cost for each plan. Specifically, insurers follow a multiplicative risk-rating rule when setting prices, where the premium faced by consumers is determined by a base premium multiplied by a type-specific risk-rating factor. Let $a(\theta)$ denote the multiplicative risk-rating factor, which represents the allowed ratio of premium charged to type θ versus the type with the lowest expected costs, $\underline{\theta}$. Full community rating, in which all types pay the same price, occurs when $a(\theta) = 1, \forall\theta$. Cases in which $1 < a(\theta) < \theta/\underline{\theta}$ would typically be called "partial" community rating, reflecting the fact that some pricing

variation is allowed but not as much as the ratio of total expected losses.

Let b_m denote the base premium for plan m, which is the community-rated premium for type $\underline{\theta}$. Then the community-rated premiums of contract m for type θ is:

$$p_m^{cr}(\theta) = b_m a(\theta). \tag{2}$$

The value of b_m is determined in equilibrium, which depends on the regulatory environment. For example, under perfect risk adjustment and a single-risk-pool pricing requirement like what exists in the ACA private health exchanges, plans are priced based on the average cost of providing that plan to the entire population of enrollees (regardless of which enrollees actually enroll in that plan). In this case, the base premium for each plan is that population average cost divided by the average community-rating multiplier:

$$b_m = \frac{\int_\Theta \lambda_m(\theta) dF(\theta)}{\int_\Theta a(\theta) dF(\theta)}, \tag{3}$$

where Θ is the set of all individuals who enroll in coverage with some plan in equilibrium.[7]

2.2 Fair Premium Differentials under Different Plan Designs

We are interested in the premium *differentials* between the two plan options that each type observes under community rating. Recalling the assumption that plan H has the higher expected cost for all types, we can define the actuarially fair relative premiums (a positive value) for each risk type as the marginal cost of additional coverage going from L to H:

$$MC(\theta) \equiv p_H^{fair}(\theta) - p_L^{fair}(\theta) = \lambda_H(\theta) - \lambda_L(\theta). \tag{6}$$

[7]The system might not break even with these base prices under the risk adjustment and single-risk-pool requirement, potentially requiring a government transfer or a market-wide adjustment factor. The base premium b_m in a competitive equilibrium with no risk adjustment depends on pricing institutions. Weyl and Veiga (2017) summarizes two pricing institutions: total pricing (TP) and incremental pricing (IP). Under TP, insurers are responsible for the costs of both plans, and each plan breaks even in equilibrium. The base premium in this case is:

$$b_m = \frac{\int_{\Theta_m} \lambda_m(\theta) dF(\theta)}{\int_{\Theta_m} a(\theta) dF(\theta)} \tag{4}$$

where Θ_m is the set of types who choose plan m in equilibrium. Under IP, the government provides a baseline level of insurance (L) with base premium b_L while insurers are responsible for the incremental cost of H. In this case,

$$b_H = b_L + \frac{\int_{\Theta_H} \lambda_H(\theta) - \lambda_L(\theta) dF(\theta)}{\int_{\Theta_H} a(\theta) dF(\theta)}, \tag{5}$$

where Θ_H is the set of types who choose plan H in equilibrium.

The value of $MC(\theta)$ depends on the plan designs of H and L. At some level, that point is obvious, as the difference in coverage generosity between H and L clearly affects the marginal cost of additional coverage. A crucial insight for this paper, though, is that the details of plan design beyond just the expected level of coverage matter for how marginal costs *vary* by risk type.

When both plans are linear, such that $\lambda_m(\theta) = c_m\theta, \forall m, \forall\theta$, where $c_m \in [0, 1]$ is a constant coinsurance rate of a plan, $MC(\theta) = (c_H - c_L)\theta$. In other words, under linear plan designs, the marginal costs are always proportional to the total expected expenditure, and the proportion, $c_H - c_L$, is the same for all types. As such, the ratio of marginal costs between higher risk types and the lowest risk type ($MC(\theta)/MC(\underline{\theta})$) will be the same as the ratio of total expenditures ($\theta/\underline{\theta}$) when considering combinations of H and L with linear coverage.

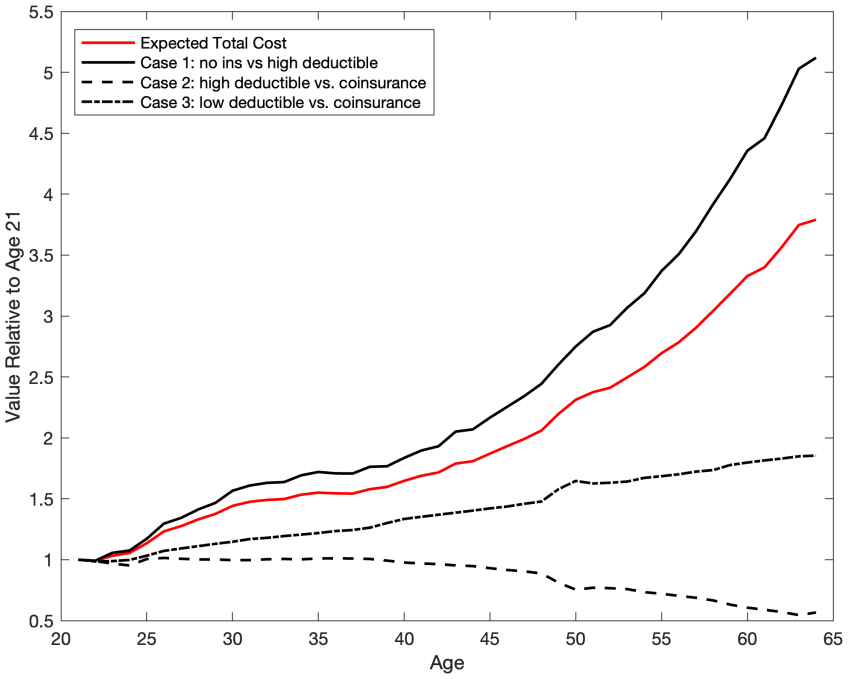
In contrast, when H and L have non-linear designs, as is common in health insurance and other markets, the marginal cost for additional coverage can vary proportionally more or less by risk type than the total expected costs do. The relationship between marginal cost and risk type depends on whether the additional coverage in the H plan primarily affects larger losses, which are relatively more common for higher-risk types, or smaller losses, which are relatively more common for lower-risk types.

We provide a simple simulation to illustrate this point. We construct age-specific total medical expenditure distributions using the Truven MarketScan data, where ages range from 21 to 64. For each age (θ), we compute $MC(\theta)$ under different non-linear plan designs for H and L. We discuss the details of the computations in Section 3.3. In Figure 1, we plot the ratio of marginal costs to those of the lowest risk type ($MC(\theta)/MC(\underline{\theta})$) for three different combinations of insurance coverage. We benchmark those ratios to the ratio of total expected costs, shown in red, which is also the ratio for any comparison of linear plan designs.

The first comparison (Case 1) in the solid black line is between L of no insurance and H of a straight high-deductible plan with a deductible and a MOOP set at $\$5,000$. In this case, the ratio of marginal costs rises much faster than total costs because the additional coverage only kicks in for relatively high losses above $\$5,000$. Those losses are common for higher-risk types, and hence the marginal cost is sizable for them, but are quite rare for low-risk types, for whom the marginal cost is quite low.

In contrast, the second comparison (Case 2) considers a case where L provides coverage for large losses with the straight deductible of $\$5,000$ and H provides additional coverage for lower losses by keeping the same MOOP at $\$5,000$ but setting a deductible of \$0 and a coinsurance

Figure 1: Ratios of Marginal Costs under Non-Linear Plan Designs by Age



Notes: The figure plots the fair premium differentials under different plan designs. The red line represents the expected total cost for each type, relative to the age of 21. "Case 1: no ins vs high deductible" represents the premium differentials between L of no insurance and H of a straight high-deductible plan with a deductible and MOOP set at $\$5,000$. Case 2 and Case 3 represent the premium differentials between a plan H with a zero deductible, a 30% coinsurance rate, and an MOOP of d, compared to a plan L with a straight-deductible plan with a deductible (and thus MOOP) of d. $d = \$500$ in Case 2 and $\$5000$ in Case 3.

rate for losses below $\$5,000$ of 30%. In this case, since the additional coverage targets lower and moderate-sized losses, the marginal cost is relatively high for lower-risk types. However, since large losses are covered under either plan, the marginal cost of additional coverage is partly muted for higher-risk types. As a consequence, as the dotted-dashed black line for Case 2 shows, the marginal costs rise in risk type at a rate that is lower than the proportion of total expected costs.

Finally, in Case 3, we consider the same dynamic but from a starting position where L already has generous coverage with a straight deductible of $\$500$ and H again offers 30% coinsurance for losses below that limit. In this situation, nearly all of the losses for higher-risk types are covered under either plan L or H, and as a consequence, marginal costs are very low for higher-risk types. Now we see that the dashed black line for Case 3 actually shows a ratio of marginal costs that is falling and below 1 for much of the range, reflecting that proportional marginal costs are rising not only less quickly in risk type than total expected costs, but are actually falling in absolute terms going from low to high risk types.

These three cases highlight that the way marginal costs vary by risk type depends on the details of the plan designs for the different vertical plan offerings (H and L). In the next subsection, we tie this back to a comparison of marginal costs to price differentials for plans under community rating. The core insight is that while marginal costs differentials depend on plan designs, community-rated price differentials are typically proportional to ratios of total expected costs and insensitive to plan designs. As a result, price differentials vary by risk type in surprising ways.

2.3 Relative Price Distortion under Community Rating

Our key question of interest is how the relative premium differentials under community rating compare to the fair premium differentials. In particular, under what conditions will the community-rated premium differential for the high-risk types be less than, equal to, or greater than the fair premium differential? To see this, we note that the question of interest is:

$$\Delta P(\theta) \equiv p_H^{cr}(\theta) - p_L^{cr}(\theta) = a(\theta)(b_H - b_L) \lesseqgtr MC(\theta), \tag{7}$$

and this, in turn, is equivalent to the following condition:

$$a(\theta) \lesseqgtr \frac{MC(\theta)}{b_H - b_L}. \tag{8}$$

The left-hand side is the risk-rating factor. The right-hand side is the ratio of the marginal costs for that risk type to the difference in base premiums between plans. It can be helpful to express the difference in base premiums ($b_H - b_L$) in terms of expected marginal costs. For the baseline case with perfect risk adjustment and a single-risk-pool pricing scheme considered in Equation (3), the base premium differential can be expressed as $\frac{E[MC(\theta)]}{E[\alpha(\theta)]}$. The differentials in base premiums are the average marginal costs across the full population divided by the average risk-rating factor across the population.[8] We can then rewrite the condition for whether the community-rated price differentials will exceed or fall below marginal costs as:

$$\frac{a(\theta)}{E[a(\theta)]} \lesseqgtr \frac{MC(\theta)}{E[MC(\theta)]}. \tag{9}$$

The right-hand side of this expression is the ratio of the type-specific marginal cost to the average marginal cost. In most situations, the ratio of marginal cost to average marginal cost

[8] Similarly, with no risk adjustment and the incremental pricing institution, we get the same formula for the base-premium differentials except that the expectations are over the population selecting plan H rather than the entire enrollee population.

will be above 1 for higher risk types and below 1 for lower risk types since higher risk types typically have higher marginal costs for additional coverage. This was true, for example, in Cases 1 and 2 considered in Figure 1. The exception will be cases, like Case 3 considered in Figure 1, where both plans are highly generous and have similar marginal costs for the higher-risk types.

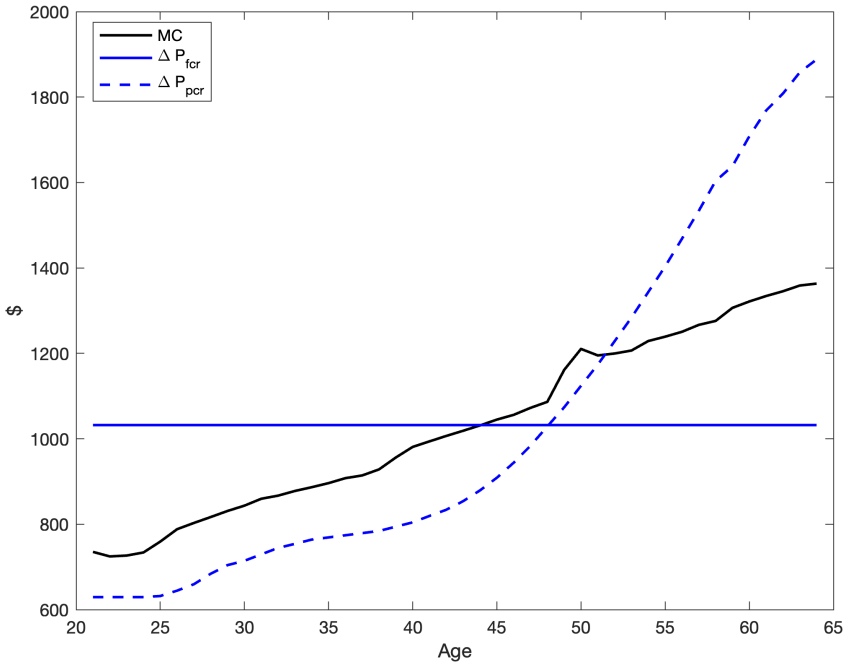
The types for whom community rating will generate price differentials that are higher or lower than marginal cost will depend on the risk-rating rule. Under full community rating, $a(\theta) = 1$ and the left-hand side is one for all types. This means that in most situations, the expected result will hold: price differentials will be above marginal cost for lower-risk types and below marginal cost for higher-risk types. This is the typical pricing distortion considered in the literature on vertical choice under community rating.[9]

Our novel theoretical result comes when we consider partial community rating that allows prices to vary somewhat by risk type. The naive intuition is that partial community rating will bring prices more in line with marginal costs. However, this intuition does not hold generally. Price differentials will be fair only when $\frac{a(\theta)}{E[a(\theta)]} = \frac{MC(\theta)}{E[MC(\theta)]}$, which happens when the proportional risk-rating factor varies with risk type exactly as the ratio of marginal costs to average marginal costs does. Unless partial community rating factors are selected with a consideration of how marginal costs vary by risk type, there is no guarantee that moving from full community rating to partial community rating will reduce premium distortions relative to marginal costs.

In particular, the key new insight of this paper is that under partial community rating, vertical price differentials can be significantly distorted from marginal costs in the *opposite* direction as the effect of community rating on the level of prices. This happens if the risk-rating factors rise for higher risk types faster than the proportional marginal costs do. As we noted in the prior section, with non-linear plan designs and at least moderate coverage, marginal costs may not rise much with risk type. This means that the right-hand side of Equation (9) will not deviate that much above (below) one for the higher (lower) risk types. However, if risk-rating factors are set to approach the ratios of total expected costs, the left-hand side of Equation (9) can be significantly above (below) one for the higher (lower) risk types. The implication is that price differentials can exceed marginal costs for higher risk types and be below marginal costs for lower risk types. Partial community rating does not simply reduce the price distortions relative to marginal costs that exist under full community rating; rather, it can lead to an "overcorrection," resulting in opposite price distortions.

[9]The exception under full community rating will be cases, like Case 3 considered in the prior subsection, where marginal costs are actually lower for higher risk types than lower risk types. In those cases, full community rating will lead to the reverse pattern.

Figure 2: Fair Premium Differentials and Community Rating Premium Differentials



Note: The figure plots an example of fair premium differentials and community-rated premium differentials under the assumption of perfect risk adjustment. We also assume the individual type (age) follows a uniform distribution. The black line, MC, represents the fair premium differentials under a plan H with a zero deductible, a 30% coinsurance rate, and an MOOP of $d = \$5,000$, relative to L with a straight-deductible plan with a deductible (and thus MOOP) of $d = \$5,000$. The solid blue line, ΔP_{fcr}, represents the full community-rated premium differentials, which are constant. The dashed blue line, ΔP_{pcr}, represents the partial community rated premium differentials where the age rating factor follows the design in the ACA marketplace.

Returning to our simple simulation from the prior subsection helps to illustrate this finding and its potential magnitudes. Recall from Figure 1 that we observed total expected medical costs that were about 3.8 times higher for the oldest relative to the youngest types. Suppose that partial community rating is implemented that allows prices to rise with age such that the oldest types pay 3 times what the youngest types do, consistent with the age-rating curve used for ACA health exchanges. For simplicity, assume a baseline scenario in which there is perfect risk adjustment and a single-risk-pool price, and that the population has a uniform age distribution from 21 to 64. How will relative prices compare to marginal costs by type? Consider our Case 2 from Figure 1, where the L plan has a straight deductible of $\$5,000$ and the H plan offers 30% coninsurance below that maximum limit. For this case, we observed in Figure 1 that marginal costs for the oldest types were a little less than 2 times those of the youngest.

Figure 2 shows how the price differentials compare to marginal costs for this situation. Marginal costs rise from a little less than \$800 for the youngest types to about $\$1,400$ for the oldest types. Under full community rating, all types face price differentials of a little more

than $\$1,000$, leading to the typical direction of price distortions across types. Under the 3:1 age-rating structure, however, price differentials are below marginal costs at young ages and dramatically above marginal costs at the oldest ages. In this example, despite being subsidized on the level of their insurance premiums, a 64-year-old considering the H plan faces a price differential that is about \$500 higher than their marginal cost.

Note that this result does not fundamentally depend on the assumption of perfect risk adjustment and single-risk pool pricing. The same basic dynamics exist if partial community rating is enforced in an environment where equilibrium prices reflect enrollment patterns.

2.4 Implications for Efficient Allocation

The preceding sections have documented how community rating regulations lead to distortions in the price differentials across insurance plans. In this section, we examine the impacts these price differentials have on the efficiency of sorting across vertical plan options. We adopt the graphical framework from Einav, Finkelstein and Cullen (2010) (EFC), which has been used extensively to examine the efficiency of sorting under full community rating. Our focus, and what is new here, is to consider how the premium distortions we highlighted under partial community rating impact sorting efficiency, especially in the case when the gradient of price differentials by risk type "overshoots" the gradient in marginal costs. We continue with two plan options H and L and ignore the outside option. On the supply side, we assume insurers are perfectly competitive and symmetric and thus earn zero profits in equilibrium. To simplify the figures in this section, we analyze the equilibrium under perfect risk adjustment and the single risk pool requirement, following Equation (3). In Appendix A, we discuss the competitive equilibrium where there is no risk adjustment and find similar insights.

Let $D(\theta)$ denote the willingness to pay for the extra coverage of H relative to L for each type, and $MC(\theta)$ denote the marginal cost of the extra coverage. We assume the simplest situation where both D and MC are linear, increasing in θ, and there is no additional heterogeneity in risk within the θ risk types. As in the standard EFC framework, we graphically sort by decreasing θ so that D and MC are downward sloping. The efficient quantity of H is determined at the point where D crosses MC, such that individuals with $D > MC$ should choose H. We use AC to denote the average cost curve. The full community rating price, ΔP_{fcr}, is a horizontal line with the level being the average value of the marginal cost across the full population, which equates to the AC when the full population is enrolled in H. The partial community rating price curve, ΔP_{pcr}, rotates around the pivot point where MC intersects P_{fcr}, and the slope is

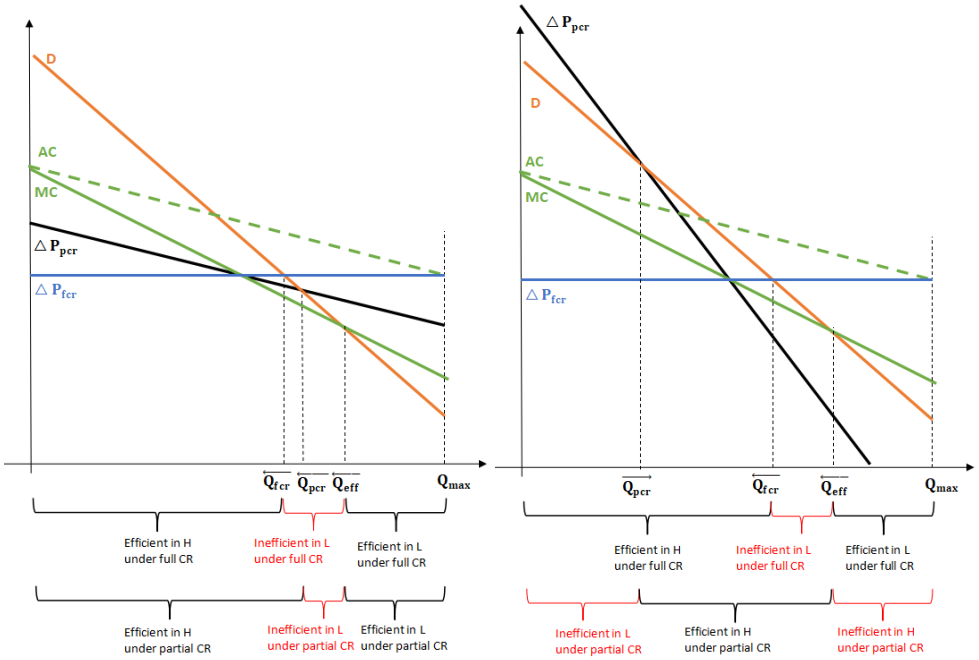
such that $\Delta P_{pcr}(\bar{\theta})/\Delta P_{pcr}(\underline{\theta}) = a$, thus the curve become steeper when a increases.

As Marone and Sabety (2022) highlight, the implications of community rating for efficient sorting in vertical-choice environments depend on how the slope of willingness to pay compares to the slope of marginal cost. The first case, denoted Case A, which is the standard case as in Einav, Finkelstein and Cullen (2010), occurs where D is steeper than MC, implying that it is efficient for higher risk types to sort into generous coverage and lower risk types to sort into less generous coverage. Marone and Sabety (2022) illustrate that in this situation, full community rating can achieve efficiency if the marginal price for additional coverage is set at the right level. More generally, as many papers, including Einav, Finkelstein and Cullen (2010) have discussed, the price under full community rating can be too high to achieve full efficiency, and this adverse-selection dynamic will cause some moderate-risk types to inefficiently sort into the L plan over the H plan. The second case considered by Marone and Sabety (2022) and denoted Case B occurs when MC is steeper than D. In this situation, it is efficient for lower-risk types to sort into H and higher-risk types to sort into L. However, as Marone and Sabety (2022) demonstrate, under full community rating, efficiency cannot be achieved because the highest (lowest) risk types will (will not) sort into H given a single marginal price that induces some demand for H.

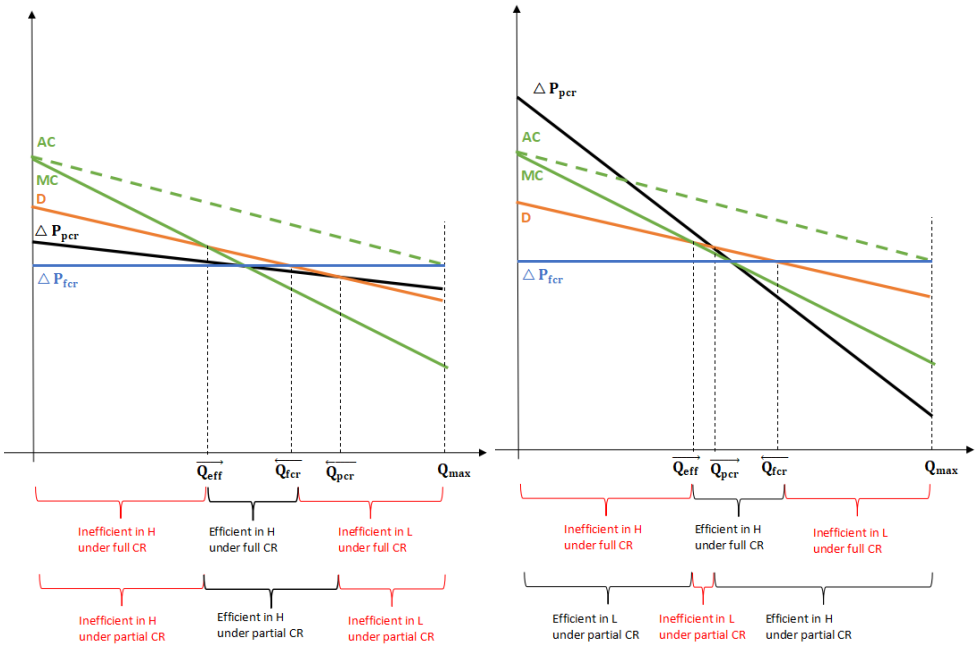
In Figure 3, we analyze the efficiency of sorting for these two cases, introducing partial community rating. The top two panels (Figures 3a and 3b) analyze Case A from Marone and Sabety (2022), where it is efficient for higher risk types to sort into H. The bottom two panels (Figures 3c and 3d) analyze Case B, where it is efficient for lower risk types to sort into L. For each of these cases, we consider two scenarios for partial community rating. In the first scenario, partial community rating increases the price differentials, represented by the curve ΔP_{pcr} in the direction of the marginal cost curve MC without overshooting the slope. This scenario is considered in the left-hand side panels (Figures 3a and 3c). The other scenario is the novel situation we highlight in this paper, where partial community rating causes the slope of the price differentials to overshoot the slope of the marginal costs, causing the ΔP_{pcr} to be above (below) MC for high (low) risk types. We illustrate a strong version of this dynamic where the ΔP_{pcr} is steep enough that it causes the marginal price of H to exceed the willingness to pay D for the highest risk types. This situation is analyzed in the right-hand side panels (Figures 3b and 3d).

Consider first Case A. The demand, marginal cost, average cost and full community rated price (ΔP_{fcr}) are the same across Figures 3a and 3b, and the only difference between the panels is the slope of the price differentials under partial community rating ΔP_{pcr}. Both figures illus-

Figure 3: Equilibrium Under Community Rating and Perfect Risk Adjustment



(a) $D > MC$ and $D > \Delta P_{pcr}$ for the old (b) $D > MC$ and $D < \Delta P_{pcr}$ for the old



(c) $D < MC$ and $D > \Delta P_{pcr}$ for the old (d) $D < MC$ and $D < \Delta P_{pcr}$ for the old

Note: The figures plot the equilibrium under perfect risk adjustment, where Q_{eff}, Q_{fcr}, and Q_{fcr} indicate the quantity under efficient equilibrium, full community rating equilibrium, and partial community rating equilibrium, respectively. The arrow points to the individuals who should select H under the corresponding equilibrium.

trate that the efficient quantity involves higher risk types sorting into H. Under full community rating, we observe some modest inefficiency in the standard direction: there is too little enrollment in H as some moderate-risk individuals sort instead into L. How does partial community

rating impact sorting? In Figure 3a, we see that when partial community rating generates a price schedule ΔP_{pcr} that goes toward MC without overshooting it, there is a modest improvement in efficiency as some additional moderate-risk types now efficiently sort into H instead of L because the price has dropped for them. We observe in Figure 3b that things change and become inefficient in very different ways when the pricing schedule ΔP_{pcr} is much steeper. In this situation, higher risk types are now priced out of H and inefficiently sort into L, while lower risk types now see a price differential substantially below MC and are induced to inefficiently sort into L.

The implications of partial community rating are quite different in Case B. In Figures 3a and 3b, we see that full community rating leads to substantial inefficiency because the single price causes higher risk types to inefficiently sort into H and lower risk types to inefficiently sort into L. In Figure 3c, we again see that a modest degree of partial community rating that generates a pricing schedule that is not steeper than MC leads to some improvement in efficient sorting. In this specific illustration, we see that the share of moderately lower risk types that are efficiently sorted into H rises. In the situation where partial community rating leads to a very steep pricing gradient, we again see a big change in the nature of sorting. Now, because the high-risk types face very high prices, they are induced to efficiently sort into L while the below-marginal-cost prices for the lower-risk types get them to efficiently sort into H. We observe in this illustration only a narrow range of inefficient sorting for middle-risk types under this steep pricing gradient.

The general intuition that emerges from this analysis is that when partial community rating causes price gradients that are much steeper than marginal cost gradients (the possibility we highlight in this paper), it can generate problematic patterns of inefficient sorting if it is efficient for higher risk types to sort into the more generous plan. On the other hand, if it is efficient for the lower risk, rather than higher risk, types to sort into the more generous plan, then steep price gradients under partial community rating may help promote efficient sorting. Importantly, in both situations, the impact of allowing more risk-based pricing in the community rating scheme does not have monotonic effects on the efficiency of sorting. Instead, the nature of the sorting patterns changes if the gradient of pricing exceeds "tipping points" compared to the slopes of MC and D.

3 Data and Background on Pricing in Private Health Insurance Exchanges

This section presents empirical evidence on price distortions under community rating rules. We examine the issue within the context of the Covered California health insurance market and supplement the analysis with data from the Federal ACA exchange. Covered California is particularly suitable to study the issue because it implements community rating rules and features a standardized choice menu with multiple products.

3.1 Institutional Background

The Covered California Health Exchange is California's health insurance marketplace, offering ACA-compliant plans to individuals and small businesses. The market features vertical choices, where plans are ranked by AV in four metal tiers: Bronze (60%), Silver (70%), Gold (80%), and Platinum (90%).[10] Unlike the ACA markets in many other states, where insurers offer a wide range of plan designs (Liu, 2025), Covered California standardizes cost-sharing designs within a metal tier and year. There are up to two allowed cost-sharing designs in each metal tier. In the Gold and Silver tiers, the cost-sharing designs are standardized, except for a few benefits that have copays and coinsurance options. For Bronze plans, there are two designs in a specific year, one with a health savings account and the other without. All insurers can only offer the standardized plan designs. Our analysis focuses on the Bronze, Silver, and Gold plans because most enrollees (more than 96%) are in these metal tiers.

There are several pricing regulations in this market. First, the single risk pool requirement requires insurers to set premiums based on all enrollees choosing their plans, rather than the pool selecting into a specific plan. The single risk pool requirement is consistent with the premium setting equations laid out in Section 2.1. Second, the market implements risk adjustment by estimating the risk scores of enrollees based on individual demographics and pre-existing conditions and transferring funds among insurers based on their enrollees' risk scores. These requirements aim to eliminate insurers' incentive to select healthier individuals.

Finally, there is community rating in the market. Plan premiums are only allowed to vary by geographic regions (rating areas) and by a pre-specified age curve. Specifically, the regulator constructs an age rating factor that specifies the premium ratio for each age group relative to that of a 21-year-old. Insurers set a base premium for the 21-year-old for a plan, and premiums

[10] There are also catastrophic plans with minimum coverage. However, these plans are only available to enrollees under age 30 or with certain exemptions.

for all other ages are generated based on the age rating curve. Most importantly, the age rating curve is the same for all plans. Our empirical analysis focuses on the age-based partial community rating rules and examines how the pricing rule distorts the relative premiums of Bronze, Silver, and Gold plans at different ages. The combination of multiplicative (partial) community rating with the single risk pool requirement means that the ACA marketplaces line up with the conceptual framework presented in the prior section.

The Covered California health insurance marketplace is a large market, enrolling nearly 2 million enrollees in 2025. Moreover, the market features many commonalities with health and drug insurance plans in other states and countries, including vertical choices and (partial) community rating with risk adjustment. The market also has unique features; for example, the cost-sharing designs are standardized within a metal tier. The standardized plan menu helps to simplify our calculations. To illustrate the generality of the results, we also replicate some of the analysis using the ACA Federal Exchange, which has the same regulations as the Covered California marketplace but does not require standardized plan designs.

3.2 Data

Covered California plan information. We collect the cost-sharing designs (deductibles, MOOPs, coinsurance rates, etc.) for plans in the market from the official Covered California website.[11] Our analysis focuses on plans offered in the individual market from 2015 to 2020. We drop 2014 from the analysis because some key plan attributes (e.g., AV, the fraction of losses covered for the average population) are not reported for this year. We supplement this data with the HIX Compare Datasets compiled by the Robert Wood Johnson Foundation.[12] The datasets contain information on other plan attributes, including carrier names, plans available region, premiums, and plan network types (HMO, PPO, or EPO).

Covered California enrollment data. The 2015 - 2020 Covered California enrollment data were acquired via a public records request and have been used in previous studies (e.g. Saltzman 2019; Tebaldi 2025). The data include enrollment information for all individuals who choose a plan via Covered California. In the data, we observe each enrollee's chosen product, age, rating area, and net premiums paid. We extract carrier name, metal tier, and plan type information from the plan name variable. We then determine the exact cost-sharing design each enrollee

[11] https://hbex.coveredca.com/stakeholders/plan-management/
[12] https://hix-compare.org/

chooses by matching to the plan information described above.[13] We follow Tebaldi (2025)'s method to infer each household's Federal Poverty Level (FPL) based on regulations on premiums and subsidies.

In our subsequent analysis, we examine counterfactual enrollment patterns under alternative pricing schemes. To assess how these changes affect the share of potential enrollees opting into the market, we need to estimate the size of the population of potential buyers. We use the American Community Survey (ACS) Public Use File, accessed via IPUMS (Ruggles et al. 2025), and follow the methodology of Finkelstein, Hendren and Shepard (2019), Tebaldi, Torgovitsky and Yang (2023), and Tebaldi (2025). Specifically, we restrict the sample to individuals who are either uninsured or privately insured and smooth the estimates of potential buyers using a flexible regression across year, rating area, age, and income groups.

Medical expenditure claims data. Our calculation of relative premium distortions requires information about the underlying medical expenditure distribution. Unfortunately, we do not have the claim data for the Covered California population. Instead, we construct estimates of the distributions by age using the 2013 Truven MarketScan data. The Truven MarketScan data is a national health insurance claim database based on the employer-sponsored market and has been used in the early years of the ACA to determine plans' AVs. We keep individuals with enrollment for all 12 months in 2013, aged between 21 and 64. We aggregate each individual's total medical expenditure from all claims (inpatient, outpatient, drugs, etc.) in 2013. These data are then used to estimate the medical expenditure distribution for each age. To reflect medical-cost inflation and average-cost differences in California relative to the national market, we apply an inflation adjustment to the estimated expenditure distributions that we describe below in Section 3.3 and Appendix A.

Plan information of the ACA Federal Exchange. We replicate some of the key results using the ACA Federal Exchange market. For this analysis, we focus on 2017 plan information since policy changes in 2018 relating to the payment of cost-sharing subsidies led to distortions in the price of Silver tier plans that are an unnecessary complication for our analysis. We collect the 2017 plan information from the Health Insurance Exchange Public Use Files published by the Center for Medicare and Medicaid Services.[14] We collect each plan's attributes from these files, including the cost-sharing design, insurer, plan type, network ID, formulary ID, launched

[13] Less than 1% of enrollees are enrolled in products not identified in the plan information data. For example, the recorded product name does not indicate the plan type, while the carrier offers multiple plan types in that metal tier and rating area. We drop these individuals from the analysis.

[14] https://www.cms.gov/marketplace/resources/data/public-use-files

county, and premiums.

3.3 Adjustments to Medical Expenditure and Plan Data

Our goal in the analysis that follows is to compare, for each age and plan, the premiums under risk-based pricing (i.e., the expected covered expenditure) and community rating. Let λ_m^i indicate the expected covered expenditure for age i under plan m, where m denotes Bronze, Silver, or Gold. We map the analysis in Section 2 into the ACA setting and have

$$p_{im}^{fair} = \lambda_m^i, \tag{10}$$

$$p_{im}^{cr} = a_i \frac{\sum_{s=21}^{64} \gamma_s \lambda_m^s}{\sum_{s=21}^{64} \gamma_s a_s} + \tau, \ \forall i, m, \tag{11}$$

where γ_i is the population weights of age i (among those who choose a plan), a_i is the age-rating factor, and τ is a level shift adjustment to all premiums to make sure the system breaks even.[15] The other key component of the calculation is λ_m^i, which is a function of the total uncovered medical expenditure distribution, $x^i \sim F^i$, and the plan cost-sharing design, $oop_m(\cdot)$: $\lambda_m^i = \mathbb{E}[oop_m(x^i)]$. Below we specify how we calculate $x^i \sim F^i$, and $oop_m(\cdot)$ respectively.

To obtain the distribution of medical costs F^i, we estimate the total medical expenditure distribution for each age using the 2013 Truven MarketScan data. We first group individuals by age. For each age, we create total medical expenditure bins (0, \$0-100, \$100-200, etc.), then calculate the average expenditure and share of individuals in that expenditure bin. We use these discretized distributions for each age as the basis for all calculations below. We then use two terms to adjust these distributions to account for the fact that health expenditure levels rise over time, differ between California and the national Truven data, and that plan prices will include a loading factor by insurers (e.g., for profits, administrative costs, and anticipated moral hazard factors at different metal tiers). The first adjustment term is a proportional inflation factor that is applied separately for each metal tier and each year. This term allows us to inflate the Truven data to match the observed average costs in the California market. The tier-specific factor accounts for the fact that different metal tiers may induce different moral hazard responses. The second term is a common additive adjustment term applied to all plans, guaranteeing that the system of premiums is expected to break even, accounting for the observed enrollment shares by age and metal tier in the California health insurance exchange.

[15] We estimate γ_s using the California enrollment data. To calculate τ , we first set $\tau = 0$, solve p_{im}^{cr} and p_{im}^{fair}. We then calculate the system-level budget using the observed 2017 California age by metal tier enrollment share, ω_{im}, where $\omega_{im} = \frac{\#\ \text{of enrollees of age i and choosing metal tier m}}{\#\ \text{of enrollees in Covered California}}$. $\tau = \sum_i \sum_m \omega_{im}(p_{im}^{fair} - p_{im}^{cr})$.

We calibrate these terms by solving for the set of proportional adjustment terms (i.e., the first term) that minimizes the squared difference between the predicted and average observed premium for each metal tier in California in the relevant year. The details of this procedure are in Appendix B. The key assumption here is that while the level of costs may differ in California, the distribution of costs by age is proportional in California to what is observed in the national data. In section 4.4.2, we conduct robustness checks to shrink or enlarge the cost differentials between ages and find similar results.[16]

The second step in our calculations is to create simplified plan designs to represent the cost-sharing of plans in California, i.e., $oop_m(\cdot)$. We consider individual coverage for the in-network services of all Covered California plan cost-sharing designs. In reality, the cost-sharing rules of a plan can be complicated, as different services may have varying coinsurance rates or copayments applied. Following Liu and Sydnor (2022), we convert all plan designs into a simplified three-arm design: The first arm is the deductible range, where no services are covered; the second arm is a coinsurance range, and the third arm is the MOOP range, under which consumers' out-of-pocket spending is capped. We convert using the CMS AV calculator for the ACA market. The calculator estimates the AV of each plan for the average population in the ACA market. We fix the deductible and MOOP of each plan and then estimate a single coinsurance rate applied to all services between the deductible and MOOP such that the plan's AV is the same as the reported AV.

The underlying assumption for this simplification approach is that the service mix is relatively stable for different ages, conditional on the total expenditure level. We examine this using the Truven MarketScan data by calculating the fraction of losses from inpatient, outpatient, and drug expenditures for three age groups, conditional on the total expenditure level. Appendix Figure C1 shows that the shares are close for these age groups for a wide range of total expenditure levels. We also examine the robustness of the method by comparing the designs of the two Bronze plans in California. The Bronze plan with a health savings account (HSA) features a 40% coinsurance rate for all services in the second arm, making it a simplified three-arm design; therefore, there is no need to convert. We compare the expected covered expenditure for each age for the two Bronze plans. If the simplification is correct, the two Bronze plans should have

[16]One piece of evidence to support this assumption comes from a consulting report by Milliman for Covered California in 2013 at the outset of the launch of the Covered California exchange. In that report, Milliman analyzed the 3:1 age-rating factor rule and reported that, based on their calculations, the actual covered cost differentials for 64- to 21-year-olds should be 4.5:1 (see page 9 of Cosway and Abbott 2013). This is very similar to both the national estimates available at that time and to our estimates. At the very least, it suggests that the type of age differentials we use in this paper is consistent with the analysis that Covered California was receiving from professional consultants.

similar λ_m^i. Indeed, Appendix Figure C2 shows that the two λ_m^i lines are extremely close to each other.

Table 1: Covered California 2017 Simplified Plan Designs

Plan	Deductible ($)	MOOP ($)	Consumers' Coinsurance rate
Gold	0	6,750	0.28
Standard Silver	2,750	6,800	0.09
Bronze	6,800	6,800	0.92
Bronze - HSA	4,800	6,550	0.40

Note: This table shows the deductibles and MOOPs in the standardized plan designs for California in 2017. The simplified co-insurance rate in the final column is our calculation. It provides a coinsurance rate that would result in a simplified three-arm plan design with a similar AV as the actual standardized plan designs, which incorporate somewhat more complex rules (including co-payments for different services). The simplified plan designs are calculated using the 2017 Actuarial Value Calculator published by CMS. All designs are for individual, first-tier in-network coverage.

Table 1 shows the simplified plan design of Gold, standard Silver, and Bronze plans for the year 2017. For all the following analyses, we use 2017 as an illustrative example. We focus on 2017 because it was the last year before the federal government stopped funding mandatory cost-sharing reductions for lower-income enrollees, which in turn led to a distortion in Silver-tier premiums in many states. While this change does not affect our primary analysis or key message of the paper, it complicates some of the comparisons. Since the two Bronze plans result in similar expected covered losses for all ages, we use the plan with no HSA as the benchmark.[17] Using the simplified plan designs and applying the adjusted Truven distributions described above, we calculate the actuarially fair and community-rated premiums using equations (10) and (11).

4 Empirical Results

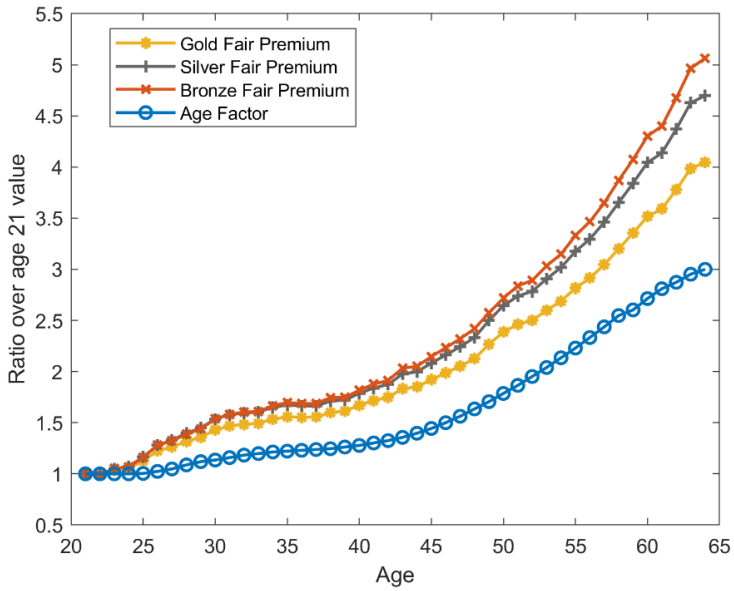
4.1 Price Distortions in Covered California

We first illustrate how the current age-based community rules distort the plan premiums relative to risk-based pricing. Figure 4 shows the age rating curve (the blue circles) and the fair premiums for Gold, Silver, and Bronze plans. We divide the premium by the level for the 21-year-olds, so the number represents the premium ratios for an age relative to the 21-year-

[17] We use a slightly different procedure and have a slightly different interpretation of the Simplified Co-insurance rate for the Bronze plan. The Bronze plan has a deductible of 6800, the same as the MOOP, while some services are covered before hitting the deductible level. We convert it into the simplified three-arm design by setting a zero deductible and finding the equivalent coinsurance rate after 0, resulting in the same AV. The resulting 0.92 factor implies that the simplified design for the Bronze plan involves 8% coverage for all losses up to the MOOP of $6,800.

olds. The age factors set a value of one for the 21-year-olds. The age rating factor suppresses the premium levels for older ages in all metal tiers, as the fair-premium curve is steeper than the age rating factors for all three plans. Moreover, the fair premium increases differently by age: the 64 to 21 fair premium ratio is about 4:1 for the Gold plan, while the ratio is 5:1 for the Bronze plan.

Figure 4: ACA Age Factors and Fair 2017 California Premiums



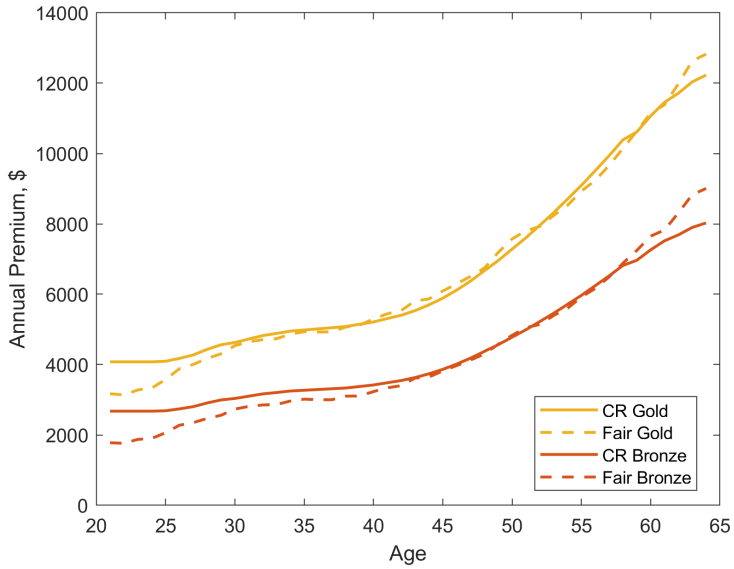
Notes: We use the 2017 California plan design to calculate the figure. In the calculation, we inflate the Truven distribution such that the simulated premiums match the actual observed 2017 premiums. All premiums are annual values for individual coverage.

Figure 5 illustrates the distortions in premium levels for Gold and Bronze plans by age. For both plans, the community-rated premiums are higher for the younger population and lower for the older population, representing a transfer from the young to the old. This is consistent with the fact that the allowed age-rating variation is lower than the underlying variation in expected costs (as shown in Figure 4).

However, the relative premium distortion is different by age. Figure 6 panel (a) shows the premium difference between Gold and Bronze plans for each age. The red line with crosses represents the risk-based actuarially fair premium differences, while the blue circle represents the community-rated premium differences. We find that the age-rating rules result in a larger premium gap than the fair level for the older population and a smaller gap for the younger population. We observe a similar pattern between Gold and Silver plans (panel (b)) and a relatively minor distortion in the premium between Silver and Bronze plans (panel (c)).

In summary, we find that the age-based community rating rules increase the premium gap

Figure 5: Premium Distortion under Community Rating: California 2017



Notes: We use the 2017 California plan design to calculate the figure. "CR" indicates community rated premiums, while "Fair" indicates premium level under risk-based pricing. The y-axis indicates the annual premiums for individual coverage.

between Gold and Bronze plans for the older population in the Covered California market.

4.2 Dominated Gold Plans in Covered California

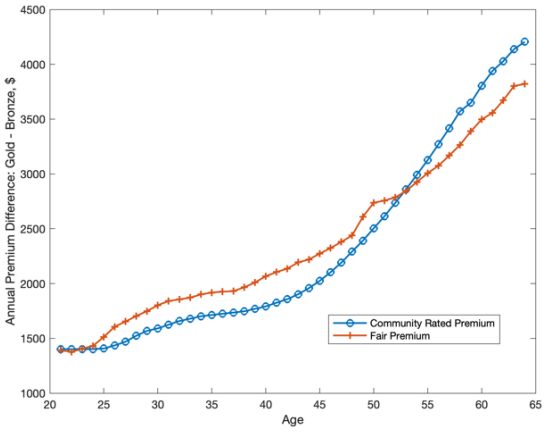
Community rating makes the Gold plan relatively more expensive for older people; in the extreme, if the premium gap between the Gold and Bronze plans is large enough, the Gold plan might be financially dominated by the Bronze plan.[18] We examine whether the premium-gap distortion is large enough to create a menu with dominated options in the Covered California market.

First, we create pairs of otherwise similar Gold-Bronze plans. We require the two plans to have the same insurer and plan type (e.g., HMO vs. PPO) and be launched in the same rating area. For example, Anthem HMO Gold and Anthem HMO Bronze plans are offered in rating area one. Some of the Bronze plans may have an HSA – we still pair them with Gold plans without an HSA because the exercise determines whether the Bronze plan dominates the Gold plan, and the availability of an HSA is an additional benefit of the Bronze plan. Sometimes, a Gold plan may not have a respective Bronze plan. For example, the insurer may offer one PPO Gold and one HMO Bronze. We count these Gold plans as not dominated by the Bronze plan

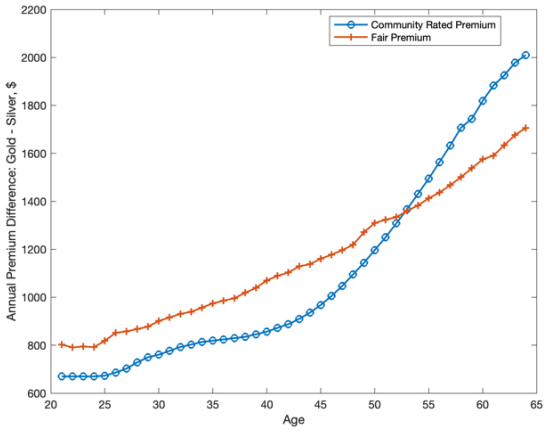
[18]In principle, the Silver plans could also dominate the Gold plans. We focus here, though, on the Gold vs. Bronze comparison since the pricing distortions are slightly larger for Gold vs. Bronze.

Figure 6: Premium Differences Across Metal Tiers: California 2017

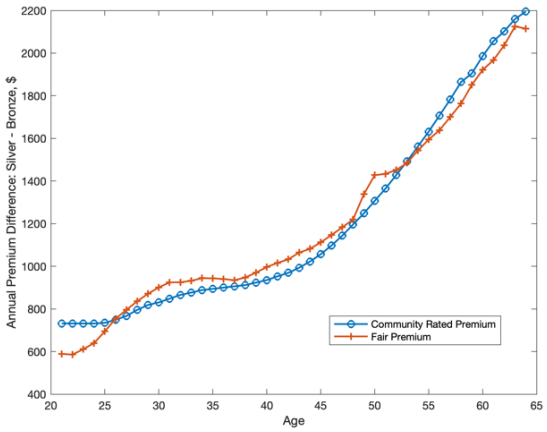
(a) Gold - Bronze



(b) Gold - Silver



(c) Silver - Bronze



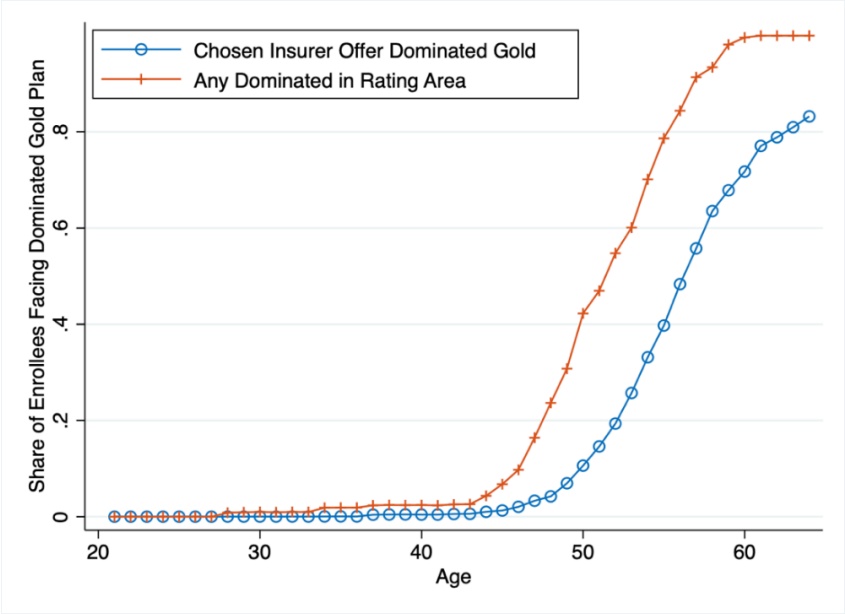
Note: We use the 2017 California plan design to calculate the figure. The y-axis indicates the annual premium differences between metal tiers for individual coverage.

in our analysis.

Second, for each pair, we calculate the threshold age above which the Bronze plan will dominate the Gold plan. To do so, we use the simplified plan design of the two plans and calculate the largest out-of-pocket spending difference between the two plans across all possible spending levels. For example, this maximum OOP difference is around $4,800 in 2017. In any situation (age, rating area, insurer, plan type, etc.) where the premium difference between the Gold and Bronze plan exceeds this value, the Bronze plan dominates the Gold plan. We then divide the value by the premium gap for the 21-year-olds to find the age factor above which Gold will be dominated, and then map it to the age. Finally, we calculate the availability of these dominated Gold plans for individuals not eligible for the cost-sharing reduction (CSR)

variation plans.[19]

Figure 7: Share Facing Dominated Gold Plan by Age: California 2015 - 2020



Notes: Data from enrollment information through Covered California 2015-2020. The sample includes individuals not eligible for CSR-variation plans. We construct Gold-Bronze pairs by requiring that they are offered by the same insurer, in the same rating area, and with the same plan type (HMO, PPO, or EPO). For each individual, we calculate whether there exists at least one Gold plan by their chosen insurer in the area they live, or among all insurers in their rating area which is dominated by the respected Bronze plan.

Figure 7 shows the share of non-CSR eligible enrollees facing a dominated Gold plan in their choice set by age. The red-cross line indicates the share of enrollees having at least one dominated Gold plan in their choice set. We find that about 40% of age 50 enrollees face at least one dominated Gold plan, and almost all age 60 enrollees in Covered California face one dominated Gold plan. If we restrict to their chosen insurer (the blue circle line), about 80% of age 60 enrollees face a dominated Gold plan offered by the insurer they eventually choose. The dominance is more likely to happen among more expensive plans (like PPOs), more expensive insurers, and more expensive rating areas.

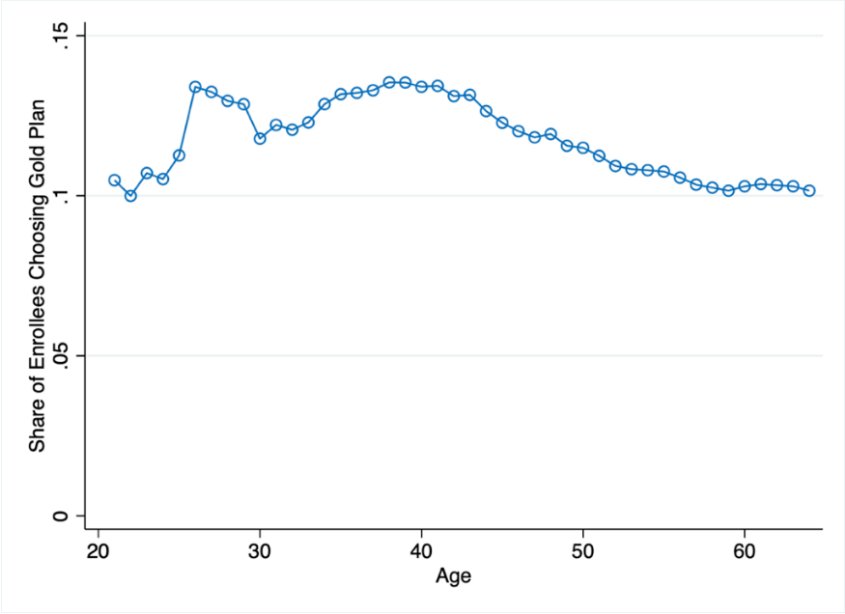
In summary, the analysis above demonstrates that the premium gap distortion resulting from the community rating rule is substantial enough to make dominated Gold plans prevalent among older enrollees in Covered California. The dominated plan creates room for making mistakes in choosing the plan option, which we examine in the next section.

[19]We exclude the CSR-eligible population, because these individuals can choose from more generous plans with the Silver premium, making Gold plans possibly dominated by these CSR-variation Silver plans. Such dominance has been documented in the literature (DeLeire et al., 2017). However, it is not the focus of our paper.

4.3 Gold Enrollment Share by Age in Covered California

First, we examine how the Gold enrollment share relates to the relative premium distortion by age. We restrict the analysis to enrollees not eligible for CSR. Figure 8 shows the share of individuals enrolled in Gold plans by age. The curve is a reverse U-shape: the Gold enrollment share increases with age up to age 40 and then decreases for enrollees 40 and above. While there may be several reasons for these age-related enrollment patterns, we note that they are at least consistent with the distortions in relative premiums we document above, which make Gold-tier plans relatively expensive for the oldest enrollees.

Figure 8: Share Choosing Gold Plan by Age: California 2015 - 2020

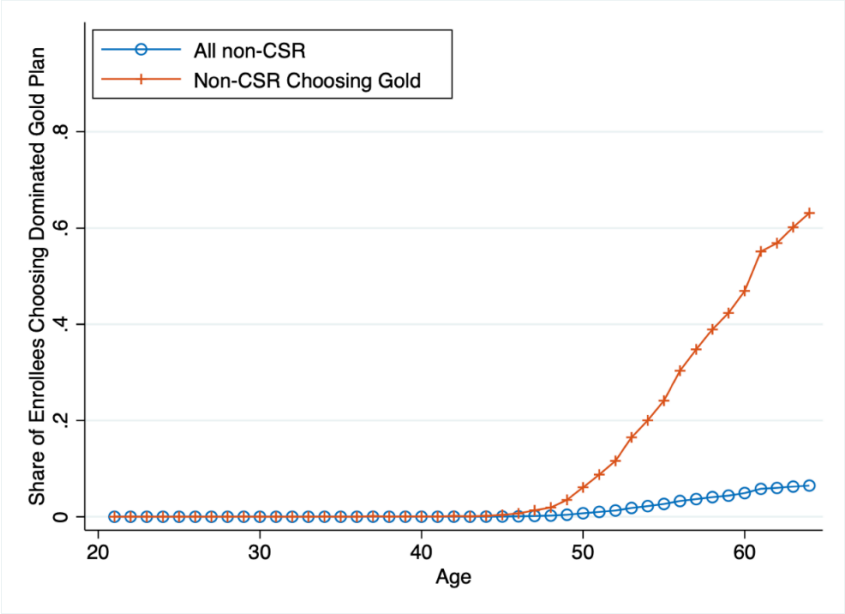


Notes: Data from enrollment information through Covered California 2015-2020. We drop the 2014 plan year because essential plan attributes (AV) are missing for this year. The sample includes those who are not eligible for CSR variations.

Second, we find that the existence of dominated Gold plans creates room for mistakes in plan selection. Among enrollees not eligible for CSR, about 2% of them choose a dominated Gold plan. Figure 9 shows the share by age. The share increases with age because older people are more likely to face a dominated option: about 5% of individuals aged 55 and above choose a dominated Gold plan. We also find that among those who choose a Gold plan, 16% of them select a dominated Gold plan, and more than 46% of enrollees aged above 54 who choose a Gold plan choose a dominated Gold plan.

Finally, the Gold enrollment share decreases more among enrollees choosing a PPO plan (see Appendix Figure C3). Conventionally, we would expect that older (and higher expenditure)

Figure 9: Share Choosing Dominated Gold Plan by Age: California 2015 - 2020



Notes: Data from enrollment information through Covered California 2015-2020. The sample includes individuals not eligible for CSR-variation plans. We calculate the share of enrollees choosing a dominated Gold plan among all non-CSR eligible enrollees and non-CSR eligible enrollees who select a Gold plan at that age.

enrollees would prefer more coverage (thus a Gold plan) and a broader network (thus a PPO plan). However, we find that the drop in Gold enrollment share by age is steeper among enrollees choosing a PPO plan. The pattern is consistent with the premium distortion: PPO plans often have higher base premiums, making the Gold-Bronze premium gap larger than HMO plans. The result suggests that the current community rating rules may distort preferences for other plan attributes.

4.4 Robustness

In this section, we examine the robustness of our main results in other markets and under alternative assumptions.

4.4.1 Dominated Gold Plans in the Federal Exchange

One might wonder whether the above patterns are a special case because California standardizes plan designs and creates menus where Bronze and Gold plans have similar MOOPs. In fact, the price distortion is a result of the non-linearity of plan design, as illustrated in the conceptual framework in Section 2. If we replace the Covered California non-linear Bronze plan with a constant coinsurance plan with the same AV, then the premium distortion is much

smaller (see Appendix Figure C4).

In this subsection, we explore the robustness of our findings in other markets. The Federal ACA exchange does not standardize plan designs, resulting in insurers creating a wide range of cost-sharing designs in the market (Liu, 2025). We replicate the dominance calculation using the 2017 Federal exchange plans. The sample includes states using Healthcare.gov in 2017 for their ACA health insurance marketplace. We construct otherwise similar Bronze-Gold pairs by requiring plans to have the same insurer, plan type, network ID, drug formulary ID, and national network status, and be launched in the same county. The requirement that the two plans have the same drug formulary ID is quite restrictive, with only 28% of counties having at least one comparable Bronze-Gold pair. As a result, we also relax this requirement (while keeping the others) and have more than 84% of counties with at least one comparable Bronze-Gold pair. We create simplified three-arm designs of these plans, and then use the same method as in Section 4.2 to calculate, for each county, the minimum age at which there is at least one dominated Gold plan.

Appendix Figure C5 shows the share of counties with at least one dominated Gold plan for that age. We find more than 50% of counties with a comparable Bronze-Gold pair have at least one dominated Gold plan for enrollees older than 58 if we use the most restrictive definition of plan attributes. If we allow Gold and Bronze plans to have possibly different drug formularies, then about 50% of counties have at least one dominated Gold plan for enrollees aged 50 and older.

4.4.2 Alternative Age-Based Medical Cost Distributions

In the baseline analysis, we rely on age distributions estimated using the Truven MarketScan data. Although we inflate the distributions to match the overall levels of Covered California, it is possible that the cost distribution differentials across ages may not be accurately estimated. In this section, we examine the robustness of the results by shrinking or enlarging the cost distribution gaps between the old and young ages.

First, we consider that there could be a bigger difference in medical spending by age than our baseline estimates. Tebaldi (2025) uses the Medical Expenditure Survey Data and estimates a gap in the medical expenditure ratio for 64-to-26 year olds that is about 13% larger than our baseline estimates using the Truven data. We adjust all distributions to match the Tebaldi (2025) estimates and the overall Covered California level.[20] Appendix Figure C6 panel (a)

[20]Specifically, we apply an age-specific multiplicative deflation factor to the baseline distributions, with the factor set as 1 for 64-year-olds, 0.89 for age 26, and a convex combination of the two for other ages where the

shows the Gold-Bronze premium gap under risk-based pricing and current community-rating pricing. We find that the distortion in relative premiums is similar to the baseline level.

Second, we also examine the robustness of the relative premium distortion when there is a smaller age gap than the Truven estimates.[21] Again, we find that the distortion in relative premiums between the Gold and Bronze plans is similar to the baseline level, as shown in Appendix Figure C6 panel (b). These findings suggest that the premium distortions we find are robust to other loss distributions.

4.4.3 Selection Across Metal Tiers and Pass Through to Premiums

In our baseline analysis, we assume that the current ACA pricing regulations, specifically the single risk pool requirement, are fully effective. That is, firms would not factor the cost differentials of enrollees selecting into the Gold and Bronze plans, if any, into the respective premiums. As a result, we interpret the tier-specific inflation factors (used to match the Truven distribution and the observed premiums in Covered California) as reflecting the differences in moral hazard responses induced by different plan tiers. In this section, we provide further justification for this assumption and discuss the robustness and implications of the main results if this assumption is violated.

First, we calculate the Gold-Bronze premium ratios of all comparable pairs (offered by the same insurer in the same rating area and with the same network type), and examine the distributions of these ratios. We find that different insurers set different ratios, which could be consistent with insurers differing in their estimates of the degree of moral hazard induced by different coverage levels. However, within an insurer, the ratio of Gold to Bronze premiums is similar, and often identical, across different rating areas and plan network types (see Appendix Figure C7.) Since selection patterns would be likely to vary across rating areas and network types, this is suggestive evidence that the tier differentials are not reflecting plan-level selection patterns.

We also find evidence that there appears to be limited selection by age into Gold or Bronze plans: in Appendix Figure C8, we plot, for each 2017 insurer, the share of enrollees at each age who chose that insurer's Bronze plan and Gold plan. We find that the two age shares are close to each other. Moreover, we find little correlation between an insurer's Gold-Bronze

weights are the age gap from 64 and 26. We then adjust all distributions by applying the same factor that matches the overall Covered California level.

[21] We hold fixed the baseline distribution for enrollees 40 or older and inflate the distribution for younger enrollees. We inflate the average spending by 10% for enrollees at age 21 and then linearly decrease the inflation from 10% to 0 as we go from age 21 to age 40.

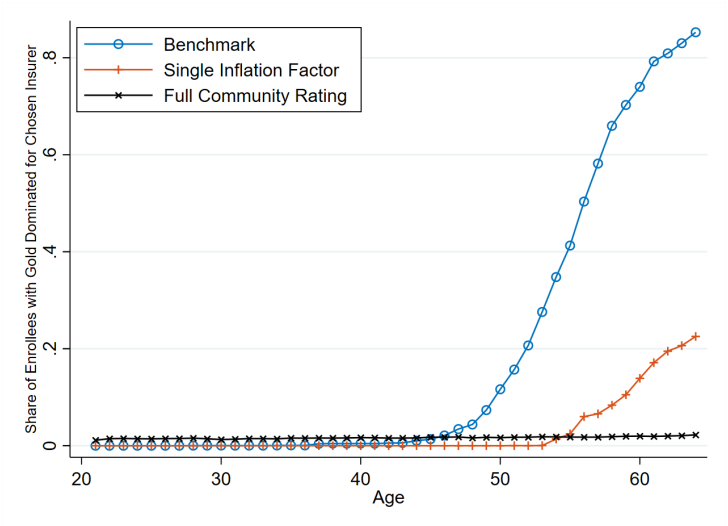
premium ratio and the Gold-Bronze expected cost difference calculated based on the age shares of enrolled individuals in the insurer and the tier (see Appendix Figure C9). These patterns suggest that there is little age-based selection across metal tiers, and if there is any, insurers do not appear to pass the cost differentials on to premiums.

Nevertheless, we can consider the implications for our analysis if some of the price gaps between metal tiers partly reflect selection. In this case, our inflation of the distributions of gold-tier enrollees to reflect moral hazard is incorrect, and the true underlying medical usage for a given individual across tiers is more similar than we are estimating. This, in turn, means that the actual fair pricing differential between metal tiers is lower than what we are currently estimating. We can bound the potential impact of this possibility by making an assumption that there is no moral hazard response to different tier enrollments. In this case, we use a single inflation factor on the Truven medical-spending distributions for all tiers to match the expected spending to the observed overall Covered California premiums. This change results in a smaller premium difference between Gold and Bronze plans, as expected. However, the premium distortions, calculated as the difference between the Gold-Bronze premium gap under the current partial community rating and the risk-based pricing, are very similar in this alternative calculation. We find that the distortion is slightly larger for the oldest and youngest enrollees under a single inflation factor, as shown in Appendix Figure C10.

Finally, we can also examine how much of the prevalence of dominated Gold-plan options can be attributed solely to the partial-community rating distortions. The actual prices for different tiers reflect the "moral hazard adjustment" factors across tiers that we build into our tier-specific inflation factors in our baseline analysis. Those increased differentials across tiers, whether they reflect selection patterns or actual moral hazard differentials, can themselves be a potential source of dominance because they raise the price differentials across tiers. To gauge how often dominance would occur without this additional consideration, we use a simple procedure to create hypothetical alternative premiums. For each insurer, we use the Silver-tier premium for their plans in a region but then set the Bronze and Gold plan premiums using the regulated AV differences (e.g., 80% for Gold vs. 70% for Silver).[22] This ensures that the premium differentials only reflect the regulated AV differences and no other adjustments. We then apply the age-

[22]To calculate the counterfactual Bronze and Gold premiums under the single inflation factor assumption, we first find the Silver plan with the same insurer-rating area-network type for each Gold-Bronze pair and take its premium p_{silver}. If there are multiple such Silver plans, we take the average premiums of these Silver plans. In our sample, all Gold-Bronze pairs have at least one respective Silver plan. We then calculate the Gold premium as $p_{silver}/0.7 \times 0.8$, and the Bronze premium as $p_{silver}/0.7 \times 0.6$. The formula is derived based on the assumption that for each Gold-Silver-Bronze pair, the inflation factor for all three plans is the same, and thus could be canceled when taking the ratio.

Figure 10: Robustness Check with Single Inflation Factor: Dominated Gold



Notes: Data from enrollment information through Covered California 2015-2020. The sample includes individuals not eligible for CSR-variation plans. We construct Gold-Bronze pairs by requiring that they are offered by the same insurer, in the same rating area, and with the same plan network type. For each individual, we calculate whether there exists at least one Gold plan by their chosen insurer in the area they live in, which is dominated by the respected Bronze plan, under three pricing schemes: the current partial community rating with tier-specific inflation factor ("Benchmark"), the current partial community rating with single inflation factor ("Single Inflation Factor"), and full community rating with age-specific inflation factor ("Full Community Rating").

rating factors to these premiums to calculate premiums at different ages. Figure 10 shows the comparison, providing both our baseline estimates of the share facing a dominated option from their chosen insurer and the share who would under this alternative approach. In this alternative, the rate of dominance is significantly lower, but still meaningful, peaking at about 21% for those in the oldest age range. At the same time, we can also consider whether the "moral hazard adjustment" factors generate patterns of dominated options if there were no age-rating differentials. For this counterfactual, we generate premiums using the baseline inflation factors that include the "moral hazard adjustment" factors, but set a fully community-rated premium that is the same for all ages. This is shown as the black line in Figure 10. We see that almost no one faces a dominated option in the absence of the age-rating differentials. Overall, this exercise highlights that high-cost areas and insurers setting larger moral hazard adjustments interact with partial community rating by age to result in the prevalence of dominated Gold plans for older enrollees.[23]

[23]This analysis takes as given the moral hazard adjustments observed in ACA pricing, which are a proportion that does not vary by age. It is also possible that the moral hazard impacts are larger for older individuals. This, in turn, would increase fair premium differentials for older individuals relative to what we calculate in our baseline analysis and imply that a greater share of the observed dominance could be attributed directly to moral-hazard adjustments in pricing. Incorporating age-specific moral hazard adjustments requires finer data

5 Counterfactual Analysis

What are the consequences of these counterintuitive pricing distortions, and how do outcomes under the partial community-rating approach in the ACA markets compare to alternative approaches to regulation? In this section, we focus on quantifying the potential market-level and distributional effects of the current community rating practice compared to other potential pricing schemes, using plan menus and enrollment data from 2017 Covered California.

5.1 Counterfactual Pricing Schemes

We simulate counterfactuals under two alternative community rating schemes: risk-based pricing and full community rating pricing. In our calculation, we hold the age-specific transfers fixed at the same level as the current partial community rating schemes. Thus, for each age, the average price level across tiers is the same under all pricing schemes. The comparison highlights how enrollment and net premium payments would change when the relative premiums across vertical options are affected by different approaches to community rating.

We first calculate the amount of the average transfer received or given at each age in the current system. We calculate these values by taking the difference between the observed premium under the current partial community rating for each age (p_{im}^{cr}) and the estimated fair premium level they would face under risk-based pricing by age (p_{im}^{fair}). We average these values for each age i across plans using the observed enrollment share. Let T_i denote the transfer level for each age, then we have:

$$T_i^{cr} = \sum_m \omega'_{im}(p_{im}^{fair} - p_{im}^{cr}), \tag{12}$$

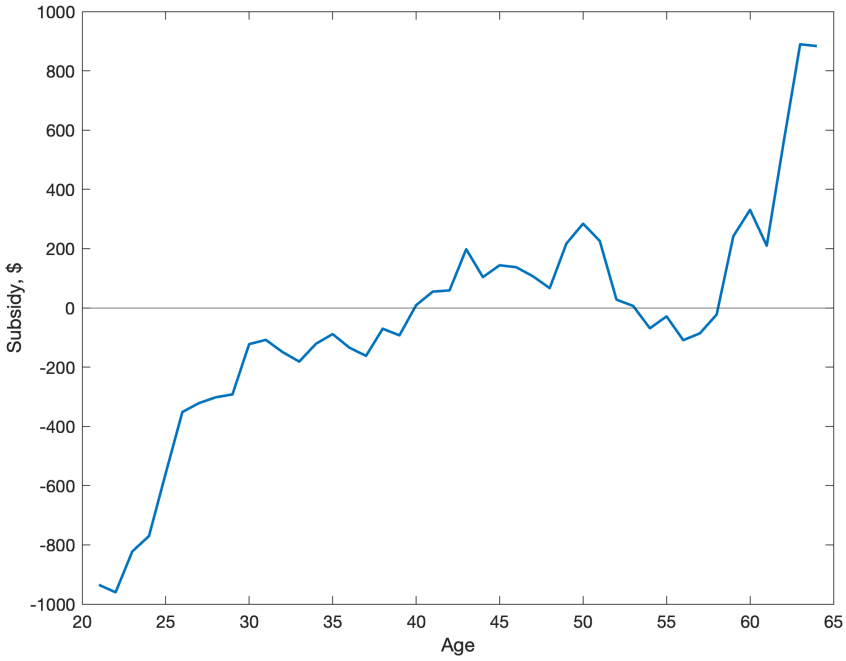
where ω'_{im} is the enrollment share of metal tier m among enrollees of age i in 2017 Covered California.[24] Figure 11 shows the results. Those above 55 years old receive positive subsidies, maxing out at transfers of almost $1,000 annually for the oldest enrollees. Those under 30, in contrast, are effectively paying average surcharges, peaking at around $1,000 for the youngest enrollees.

Next, we design a risk-based pricing system that incorporates community rating to achieve transfers in the level of premiums but allows for price differentials between plan options in the menu that reflect expected costs. Specifically, we calculate the counterfactual premium under

than we currently have and is left for future work.

[24] Specifically, $\omega'_{im} = \frac{\#\ of\ enrollees\ of\ age\ i\ and\ choose\ metal\ tier\ m}{\#\ of\ Covered\ California\ enrollees\ of\ age\ i}$. When calculating ω'_{im}, we condition on enrollees making a plan choice, because under ACA, only individuals who buy an insurance plan are eligible for age-specific transfers.

Figure 11: Implied Annual Subsidy by Age: California 2017



Notes: We use the 2017 California plan design to calculate the figure. We first calculate the subsidy level for each age and metal tier by subtracting the annual fair premium from the community-rated premiums. We then calculate the average subsidy level for a particular age by weighting the subsidy for Gold, Silver, and Bronze of that age by the California enrollment share in these three tiers at that age. The enrollment share is based on the non-CSR eligible enrollees.

risk-based pricing for age i and metal tier m, p_{im}^{rb} as:

$$p_{im}^{rb} = p_{im}^{fair} - T_i^{cr}. \tag{13}$$

The risk-based pricing with age-specific side transfers could be achieved by modifying the pricing regulations. For example, the age-rating curve could be applied to Bronze plans, but then insurers could be required to price higher tiers based on expected cost differentials by age. Another approach could be to allow for risk-based pricing of plans, but then to impose a series of surcharges and subsidies based on observable risk (but not plan selection).

Third, we design a full community rating pricing system with side transfers across ages. For each plan, we first calculate the full community-rated premiums, $\bar{p}_m$, as a weighted average of the fair premiums of the plans across ages: $\bar{p}_m = \sum_i \tilde{\omega}_i p_{im}^{fair}$, where $\tilde{\omega}_i$ is the share of enrollees of age i among those enrolled in marketplace coverage.[25] We then adjust the full community-rated premiums by adding age-specific transfers, such that transfer levels are the same as the baseline

[25] Specifically, $\tilde{\omega}_i = \frac{\#\,of\ enrollees\ of\ age\ i}{\#\,of\ enrollees\ in\ Covered\ California}$

level:

$$p_{im}^{full} = \bar{p}_m + \Delta T_i + \tau^{full}. \tag{14}$$

We calculate the age-specific transfer term as $\Delta T_i = T_i^{cr} - T_i^{full}$, where $T_i^{full} = \sum_m \omega'_{im}(p_{im}^{fair} - \bar{p}_{im})$. Thus, enrollees receive the same level of transfers as the baseline level. τ^{full} is a level shift adjustment to all plan premiums to make sure the system breaks even: $\tau^{full} = \sum_i \sum_m \omega_{im}(p_{im}^{fair} - p_{im}^{full})$.

5.2 Results with No Enrollment Change

We first consider the case where the enrollment shares are held fixed at the current partial community rating level (i.e., the same as the observed data). When this is true, there will only be distributional effects due to premium differences.

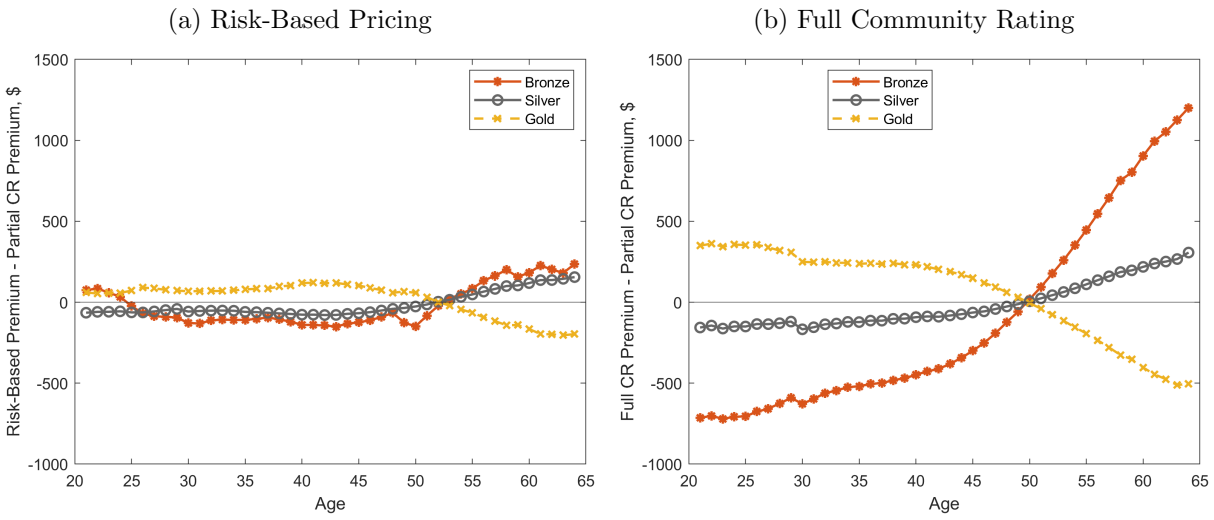
We calculate how the premiums would change for each metal tier and each age under the two counterfactual pricing schemes. The premium is defined as the annual premium for a single coverage. In the calculation, we do not consider the premium subsidy linked to income (the premium tax credit, APTC), so the results reflect the premium changes among those who are not eligible for the premium subsidy. Figure 12 panel (a) graphs these values by age for the Bronze, Silver, and Gold-tier plans. Under risk-based pricing, the biggest change in premiums would occur for Gold-tier plans for older individuals. Gold-tier prices would fall by as much as \$200 per year for the oldest enrollees. The biggest increases in Gold-tier premiums would occur for those around age 42, who would see increases of around \$100 per year. In contrast, Bronze and Silver-tier prices would rise for those above age 55 and fall modestly for those from age 25 to 55.

Figure 12 panel (b) shows the annual premium change under full community rating. We find a similar pattern to that in the risk-based pricing case, but the magnitude of impacts is much larger. For example, the Gold-tier prices would fall by \$500 per year for the oldest enrollees, while the Bronze-tier prices would increase by around \$1000 for individuals older than 55.

5.3 Results with Enrollment Change

We then calculate the counterfactual under the two pricing schemes, allowing for the possibility that enrollees would change plan choices in response to the premium changes. We first formulate and estimate a model of individual health insurance demand. We then use the demand

Figure 12: Annual Premium Change under Counterfactuals



Notes: We use the 2017 California plan design to calculate the figure. The y-axis indicates the annual premium difference (for individual coverage) between two counterfactual pricing rules and the current partial community rating pricing. In panel (a), the counterfactual pricing is the risk-based pricing with age-specific side transfers. In panel (b), the counterfactual pricing is the full community rating with age-specific side transfers. A positive number indicates that the counterfactual premiums are higher than the current premiums.

estimates to simulate the enrollment and premium changes under the three pricing schemes and compare the equilibrium outcomes.

Demand Estimation We estimate individual-level demand for health insurance plans offered through Covered California, the state's ACA marketplace. We consolidate plan options and let enrollees to choose from five option, four metal tiers and the outside option (not buying insurance). Following the literature (Saltzman, 2019; Saltzman, Swanson and Polsky, 2021; Tebaldi, Torgovitsky and Yang, 2023; Tebaldi, 2025), we model individuals' indirect utility for plan options as a function of the net premiums they pay, age-specific AV, and other observed plan attributes. We normalize the outside option, i.e., not insured, as having a value of zero. We then allow the premium and AV coefficients to flexibly vary with individual characteristics of income, age, and regions, to capture rich heterogeneity in preferences for plan attributes.

Our estimation focuses on single-person households with individuals aged 25 to 64. We construct the sample of potential buyers in the market (including those who choose no insurance) using data from the American Community Survey (ACS). We restrict attention to this group to avoid complications related to family-level plan selection and Medicare eligibility. Inertia is a common feature of health insurance markets, where consumers often stick with previous choices even when better options exist (Handel, 2013; Saltzman, Swanson and Polsky, 2021). To

abstract from inertia and focus on active decision-making, we use data from 2014—the first year Covered California was launched—when all enrollees were required to make an active choice. This also aligns with our counterfactual policy experiments, which consider a setting where pricing schemes were different from the outset of the program or all individuals were required to make active choices from the beginning. While modeling inertia is important in many settings, our focus here is on the demand response under active choice, and so 2014 provides a clean environment to estimate such behavior.

We describe the demand model and estimation strategy in detail in Appendix B. The estimated model fits the observed plan shares well and yields price sensitivity estimates that are broadly in line with prior work (Saltzman, 2019; Saltzman, Swanson and Polsky, 2021; Tebaldi, Torgovitsky and Yang, 2023; Tebaldi, 2025). The implied own-price elasticities are modest, consistent with the literature on health insurance demand in ACA exchanges. These estimates provide a sensible foundation for our counterfactual analyses, where relatively muted price responsiveness implies that changes in pricing rules will generate limited re-sorting across plans.

Counterfactual Calculation and Results We draw a 10% random sample from the 2017 Covered California enrollment data and use it in the counterfactual calculations. Thus, the counterfactual is based on the 2017 plan menu and population distribution, and comparable to the values in Section 5.2. We solve the equilibrium under three pricing schemes: the partial community rule with the 3-to-1 age rating factors, the risk-based pricing with age-specific side transfers, and the full community rating with age-specific side transfers. We solve the equilibrium by assuming a competitive supply side and numerically searching for the enrollment allocation and premium levels under which enrollees optimally choose a plan according to the estimated demand model. Insurers earn an average of zero profits across all plans. We specify the details of the estimation and equilibrium calculation in Appendix B.

We present the results in Table 2. We estimate that the current partial community-rating leads to plan allocations and government costs that are quite close to those that could be achieved under a risk-based pricing alternative. The primary difference between the current practice and the risk-based alternative is that under the current practice, there is a distributional impact primarily on enrollees who opt into more generous "Gold plans". In Figure 13 panel (a), we show that older individuals enrolled in Gold plans pay about $400 more per year than they would under risk-based pricing, while younger individuals enrolled in Gold plans pay about $200 less under the partial-community rating approach.

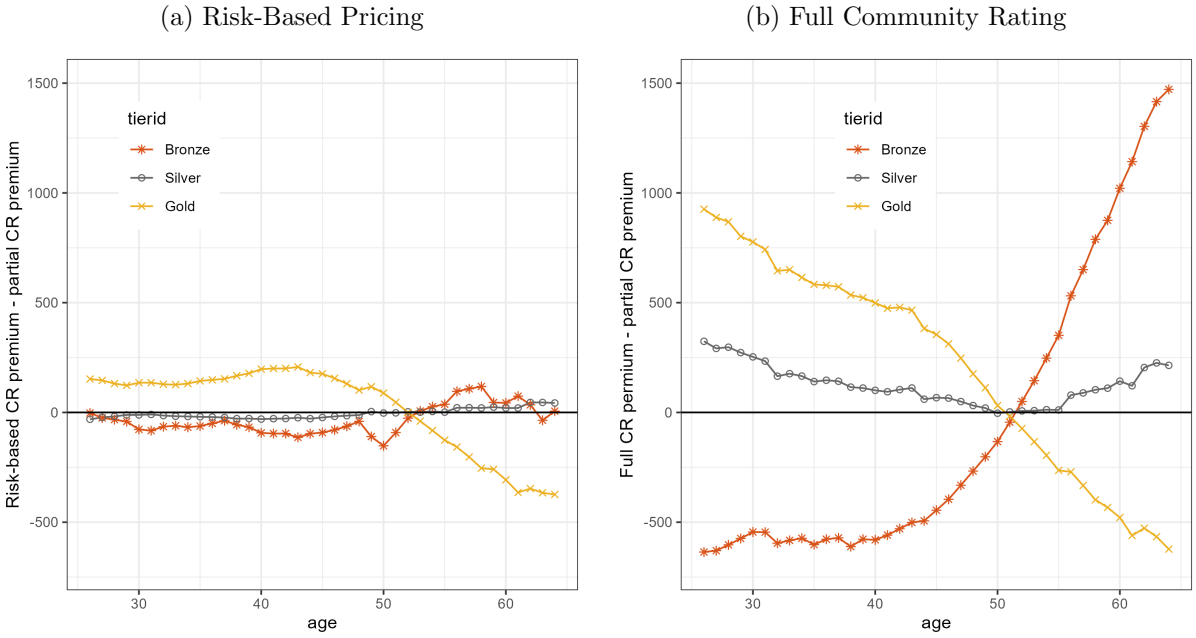
Table 2: Equilibrium Premium and Enrollment under Three Pricing Schemes

	Pricing schemes	all	young	middle	old
Panel A. Premiums ($1,000s)					
Total annual premium	risk-based	6.31	4.17	5.44	8.68
	partial CR	6.32	4.11	5.38	8.78
	full CR	6.44	4.51	5.52	8.65
Net annual premium	risk-based	3.99	2.83	3.73	5.15
	partial CR	4.00	2.78	3.65	5.24
	full CR	4.08	3.26	3.87	4.91
Government subsidy	risk-based	3.12	2.53	2.55	3.98
	partial CR	3.13	2.52	2.57	3.98
	full CR	3.16	2.44	2.49	4.19
Panel B. Enrollment					
No insurance rate	risk-based	0.75	0.74	0.78	0.74
	partial CR	0.75	0.74	0.78	0.74
	full CR	0.75	0.77	0.79	0.72
Bronze enrollment share	risk-based	0.19	0.18	0.19	0.20
	partial CR	0.19	0.17	0.18	0.20
	full CR	0.23	0.34	0.25	0.14
Silver enrollment share	risk-based	0.67	0.69	0.68	0.64
	partial CR	0.67	0.68	0.68	0.66
	full CR	0.63	0.59	0.64	0.66
Gold enrollment share	risk-based	0.11	0.10	0.10	0.13
	partial CR	0.12	0.12	0.12	0.11
	full CR	0.11	0.06	0.09	0.15

Note: In column "All", the values are averaged across all enrollees. In column "Young", the values are averaged across enrollees aged between 26 and 39. In column "Middle", the values are averaged across enrollees aged between 40 and 50. In column "Old", the values are averaged across enrollees aged between 51 and 64. Total annual premiums refer to the average premiums of plans that the enrollees face (other than the outside option). Net annual premiums are the total annual premiums minus the APTC payment (if eligible). Government Subsidy refers to the per-enrollee payment (including all enrollees) for the APTC and CSR payments. All three values are measured in $1,000s. Bronze, Silver, and Gold shares are also conditional when choosing a plan (The platinum plan is committed in the table).

When compared to an alternative of full community rating, the partial community rating generates bigger differences. Under full community rating, older individuals would be highly *subsidized* on the margin for additional coverage, while younger individuals would be charged substantially above their marginal cost. Our demand estimation model suggests overall only modest enrollment reactions to prices in these markets (likely reflecting substantial frictions in plan choice). Nonetheless, we estimate that switching to full community rating would reduce the fraction of eligible young individuals enrolling in the ACA by around three percentage points, while increasing enrollment at older ages by around two percentage points. We also estimate that full community rating would modestly increase the share of older individuals in Gold plans while roughly halving the fraction of young individuals with Gold plans. The distributional

Figure 13: Equilibrium Annual Premium Difference btw Counterfactual and Current Pricing



Notes: We use the 2017 California plan design to calculate the figure. The y-axis indicates the annual premium difference (for individual coverage) between two counterfactual pricing rules and the current partial community rating pricing. In panel (a), the counterfactual pricing is the risk-based pricing with age-specific side transfers. In panel (b), the counterfactual pricing is the full community rating with age-specific side transfers. A positive number indicates that the counterfactual premiums are higher than the current premiums.

impacts of switching to full community rating would also be large. As shown in Figure 13 panel (b), we estimate that older individuals who remain in Gold plans would save over $500 per year under full community rating, while older individuals enrolled in Bronze plans would pay an additional $1,500 per year. The impacts are reversed for younger individuals. Appendix figures C11-C15 further break down these results by income levels and show similar patterns.

5.4 Practical Alternative for the ACA

The risk-based pricing scheme with side transfers analyzed above can fix the distortions in relative premiums and the associated inefficiencies in vertical plan choices. However, this approach is a big departure from the current practice and is not clearly a practical policy recommendation. In this subsection, we briefly discuss an alternative tweak to existing community-rating approaches in the US private health exchanges that would not get the marginal prices exactly right, but would move them much closer.

The first step in this alternative would be to apply the current ACA pricing rules for Silver plans under the 3:1 age-rating curve. Second, for Gold and Bronze plans, calculate separately

for each age the ratio between the AV under Silver and the alternative tier (Gold or Bronze) plans. The calculation requires age-specific health spending distributions, as we do in the paper, and averaging across different plan designs within a tier available in the market. Alternatively, if a standardized plan design is in place, that plan design could be used as the one to calculate the age-specific AVs. Finally, compute the respective Bronze and Gold premiums by multiplying the Silver premium by the computed AV ratio for each age.[26]

The end result of this alternative method would be to bring the marginal prices across tiers more in line with the marginal costs across tiers at each age. The levels of prices would still be partially cross-subsidized from the young to the old. We can anticipate that this would lead to lower bronze and higher gold prices for the young, and higher bronze and lower gold prices for the old. It is worth noting that the 3-to-1 age rating would now hold for Silver plans but not for Bronze or Gold plans. The effective age-rating ratio for Bronze will end up higher than 3-to-1 and for Gold lower than 3-to-1, consistent with the fact that costs vary differently by age across the tiers.

6 Conclusion

We have examined the effects of community rating as it is commonly practiced on the differentials in prices for insurance options. Our conceptual framework provides a set of simple formulas for identifying the direction of these distortions. The key consideration is whether the allowed pricing variation across risk types exceeds the ratio of marginal costs between the high-risk and low-risk population. If it does, then the distortions in premium differentials will run counter to the direction of the distortions in premium levels. This is possible when insurance plans expected costs do not vary linearly with risk, which is typical for products like health insurance. Our empirical analysis shows that this dynamic occurs in the private health insurance exchanges in California, and likely in many other private health insurance markets throughout the United States. The effect is strong enough that the most generous plans are often dominated by lower-generosity options for the oldest enrollees in the private health insurance markets in the U.S.

In principle, insurers could correct the premium distortion by charging an age-rating factor smaller than the allowed 3-to-1 ratio (though they can't use larger age-rating factors). However, in practice, all insurers in Covered California use the 3-to-1 age-rating factors. It is unclear

[26]It would also be possible to retain current ACA practice and add an additional overall "moral-hazard adjustment factor" to different tiers that applies the same across all ages. Alternatively, the moral-hazard adjustment factor could be built into the AV calculations and be age-specific.

whether the fact that they all use the 3-to-1 is a conscious attempt to inflate the differentials, due, for example, to imperfections in the risk adjustment formula, or if they are just mechanically using the suggested ratio and not focusing on this differential pricing impact. Future work with access to micro-level data on enrollee costs could explore the competitive impacts of different potential age-rating strategies insurers could employ in these markets and their implications for our understanding of firm behavior and the impacts of these policies.

How do the considerations we raise in this paper affect the way economists should think about the welfare consequences of the interaction between community rating and multiple product options in insurance markets? Prior research has highlighted that inefficiencies can arise in multi-product insurance markets where price differentials between options do not align with expected costs (see e.g., Einav, Finkelstein and Cullen 2010; Ericson and Sydnor 2017; Marone and Sabety 2022). This work has, in turn, called into question whether vertical choice between coverage levels under community rating provides welfare improvements, often concluding that the inefficiencies from price distortions dramatically reduce or reverse the value of offering choice. Our paper does not challenge that existing logic, but highlights that the direction of the distortion in price differentials can sometimes run in a different direction than common intuition would suggest. The pattern we observe empirically – in which more generous plan options are especially costly for older individuals while effectively subsidized on the margin for younger individuals – could impact the system's efficiency, but the direction is ex-ante ambiguous. On one hand, if older individuals are more prone to inefficient moral hazard than younger individuals, these price distortions may improve the efficiency of the community-rated multiproduct market by pushing older individuals away from generous coverage. On the other hand, if older individuals are more risk averse than younger individuals, the effect would cut the other way. We lack the type of data that would be necessary to fully evaluate the welfare impacts in this market relative to particular alternative systems of interest, but the results highlight that evaluating welfare in multiproduct markets requires attention to how community rating interacts with insurance plan designs.

It seems likely that the counterintuitive pricing distortions we uncover interact with choice heuristics in ways that would harm the most naïve enrollees. When faced with a choice between health insurance plans that are vertically differentiated, more naïve individuals who cannot easily compare plans may fall back on heuristics. A natural heuristic in the face of vertical choice would be to assume that if you are higher-risk, you would benefit from more generous coverage and vice versa. In this setting, though, that heuristic would run counter to the marginal prices and would result in many cases in which more naïve older enrollees would violate dominance.

While the implications of these observations for the implementation of community rating regulations are unclear, these results clearly highlight a need for a better understanding of the impact of pricing regulations in insurance markets with multiple options. They also likely call for additional decision support for both consumers and policymakers, hoping to navigate the complexities of pricing for these complex financial products.

7 Bibliography

Abaluck, Jason, and Jonathan Gruber. 2011. "Choice inconsistencies among the elderly: Evidence from plan choice in the Medicare Part D program." *American Economic Review*, 101(4): 1180–1210.

Abaluck, Jason, and Jonathan Gruber. 2023. "When less is more: Improving choices in health insurance markets." *The Review of Economic Studies*, 90(3): 1011–1040.

Bhargava, Saurabh, George Loewenstein, and Justin Sydnor. 2017. "Choose to Lose: Health Plan Choices from a Menu with Dominated Option." *Quarterly Journal of Economics*, 132(3): 1319–1372.

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

Cosway, Robert, and Barbara Abbott. 2013. "Factors Affecting Individual Premium Rates in 2014 for California." Milliman Client Report.

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

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

Ericson, Keith M., and Justin Sydnor. 2017. "The Questionable Value of Having a Choice of Levels of Health Insurance Coverage." *Journal of Economic Perspectives*, 31(4): 51–72.

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2019. "Subsidizing health insurance for low-income adults: Evidence from Massachusetts." *American Economic Review*, 109(4): 1530–1567.

Fontana, Joanne, Thomas Murawski, and Sean Hilton. 2017. "Impact of Changing ACA Age Rating Structure: An Analysis of Premiums and Enrollment by Age Band." Milliman Research Report.

Geruso, Michael, and Timothy J. Layton. 2017. "Selection in Health Insurance Markets and Its Policy Remedies." *Journal of Economic Perspectives*, 31(4): 23–50.

Geruso, Michael, Timothy J. Layton, Grace McCormack, and Mark Shepard. 2023. "The Two-Margin Problem in Insurance Markets." *Review of Economics and Statistics*, 105(2): 237–257.

Handel, Benjamin, Jonathan Kolstad, Thomas Minten, and Johannes Spinnewijn. 2024. "The socioeconomic distribution of choice quality: evidence from health insurance in the Netherlands." *American Economic Review: Insights*, 6(3): 395–412.

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

Heiss, Florian, Daniel McFadden, and Joachim Winter. 2010. "Mind the Gap! Consumer Perceptions and Choices of Medicare Part D Prescription Drug Plans." In *Research Findings in the Economics of Aging*. University of Chicago Press.

Heiss, Florian, Daniel McFadden, Joachim Winter, Amelie Wuppermann, and Bo Zhou. 2021. "Inattention and Switching Costs as Sources of Inertia in Medicare Part D." *American Economic Review*, 111(9): 2737–2781.

Johnson, Eric J., Ran Hassin, Tom Baker, Allison T. Bajger, and Galen Treuer. 2013. "Can Consumers Make Affordable Care Affordable? The Value of Choice Architecture." *PLOS ONE*, 8(12): e81521.

Kairies-Schwarz, Nadja, Johanna Kokot, Markus Vomhof, and Jens Wessling. 2014. "How Do Consumers Choose Health Insurance?–An Experiment on Heterogeneity in Attribute Tastes and Risk Preferences." Ruhr Economic Paper 537.

Liu, Chenyuan. 2025. "Sorting on Plan Design: Theory and Evidence of the ACA." Working paper.

Liu, Chenyuan, and Justin Sydnor. 2022. "Dominated Options in Health Insurance Plans." *American Economic Journal: Economic Policy*, 14(1): 277–300.

Marone, Victoria R., and Adrienne Sabety. 2022. "When Should There Be Vertical Choice in Health Insurance Markets?" *American Economic Review*, 112(1): 304–342.

Rasmussen, Petra W., and David Anderson. 2021. "When All That Glitters Is Gold: Dominated Plan Choice on Covered California for the 2018 Plan Year." *Milbank Quarterly*, 99(4): 1059–1087.

Ruggles, Steven, Sarah Flood, Sarah Flood, Matthew Sobek, Daniel Backman, Grace Cooper, Julia A. Rivera Drew, Richards Stephanie, Renae Rodgers, Jonathan Schroeder, and Kari C.W. Williams. 2025. "IPUMS USA: Version 16.0 [dataset]." Minneapolis, MN: IPUMS.

Saltzman, Evan. 2019. "Demand for Health Insurance: Evidence from the California and Washington ACA Exchanges." *Journal of Health Economics*, 63: 197–222.

Saltzman, Evan, Ashley Swanson, and Daniel Polsky. 2021. "Inertia, market power, and adverse selection in health insurance: evidence from the ACA exchanges." National Bureau of Economic Research.

Samek, Anya, and Justin Sydnor. 2025. "The Impact of Consequence Information on Insurance Choice." *Review of Economics and Statistics*, 1–46.

Schram, Arthur, and Joep Sonnemans. 2011. "How Individuals Choose Health Insurance: An Experimental Analysis." *European Economic Review*, 55(6): 799–819.

Tebaldi, Pietro. 2025. "Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca." *Review of Economic Studies*, 92(1): 586–620.

Tebaldi, Pietro, Alexander Torgovitsky, and Hanbin Yang. 2023. "Nonparametric estimates of demand in the california health insurance exchange." *Econometrica*, 91(1): 107–146.

Veiga, Andre. 2023. "Price Discrimination in Selection Markets." *Review of Economics and Statistics*, 1–45.

Weyl, E. Glen, and André Veiga. 2017. "Pricing Institutions and the Welfare Cost of Adverse Selection." *American Economic Journal: Microeconomics*, 9(2): 139–148.

Appendix A Equilibrium under Community Rating and No Risk Adjustment

We discuss the equilibrium under community rating and no risk adjustment. As in the main text, we consider the choice between H and L and ignore the outside option. We assume insurers are responsible for the incremental cost of H relative to L, are perfectly competitive and symmetric, thus earn zero profits in equilibrium.[27]

Both $D(\theta)$ and $MC(\theta)$ are the same as in the main text. The efficient quantity of H is determined when D crosses MC, implying that individuals to the left of Q_{eff} should choose H. Under full community rating, all types face the same price, the equilibrium is determined at the point where the average cost curve (AC) crosses D, resulting in individuals to the left of Q_{fcr} choosing H. In this case, full community rating results in a group of middle-aged individuals who inefficiently sort in L, a result well documented in the literature.

Now consider the equilibrium under partial community rating in the graph. Suppose $a(\bar{\theta}) = a > 1$. Let $\Delta P_{pcr}(\theta)$ denote the price faced by each type. $\Delta P_{pcr}(\theta) = a(\theta)\Delta b$, where $\Delta b = b_H - b_L$. Thus, ΔP_{pcr} is downward-sloping in the graph. For any given ΔP_{pcr} curve, we could also draw the AP_{pcr} curve, indicating the average price among individuals who choose H. AP_{pcr} is downward-sloping, above ΔP_{pcr}, and crosses ΔP_{pcr} on the y-axis. Given any a, a change in Δb ($= \Delta P_{pcr}(\underline{\theta})$) means changes in both the level and slopes of ΔP_{pcr} and AP_{pcr} because we need to keep $\Delta P_{pcr}(\bar{\theta})/\Delta P_{pcr}(\underline{\theta}) = a$. Individuals with ΔP_{pcr} below D choose H. Equilibrium happens when there exists a θ such that $\Delta P_{pcr}(\theta) = D(\theta)$ and $AP_{pcr}(\theta) = AC(\theta)$. The condition $\Delta P_{pcr}(\theta) = D(\theta)$ determines the individuals who purchase H, and $AP_{pcr}(\theta) = AC(\theta)$ makes sure insurers break even for these individuals. In the graph, this condition is satisfied when we move and rotate ΔP_{pcr} and AP_{pcr} curves to a position such that the point where AC crosses AP_{pcr} is vertically aligned with the point where D crosses ΔP_{pcr}. In Figure A1a, we show one such case, where a is close to one and ΔP_{pcr} is relatively flat. The resulting equilibrium is that all individuals to the left of Q_{pcr} choose H. Note that in this case, partial community rating mitigates the inefficiency of full community rating, a result consistent with the conventional wisdom.

Our new insight is that under partial community rating, sometimes ΔP_{pcr} may overshoot and become larger than MC for the old individuals (on the left) and smaller than MC for the

[27] Weyl and Veiga (2017) highlights that under full community rating, the market equilibrium is different under incremental pricing than total pricing. The market is more likely to unravel under total pricing. Our welfare discussion below focuses on the incremental pricing institution for ease of exposition, but the key insights are similar under the total pricing institution.

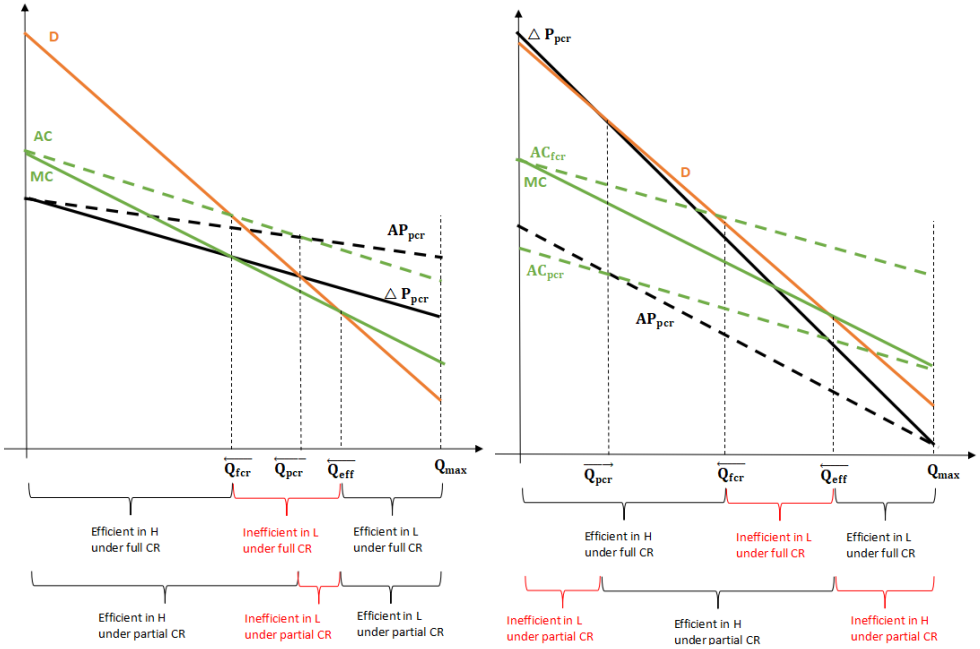
young individuals (on the right). Figure A1b depicts a case under which the price distortion of relative premium is so large that ΔP_{pcr} becomes steeper than D for the old individuals. Now, because ΔP_{pcr} is below D for the young individuals and above D for the old individuals, AP_{pcr} curve is below ΔP_{pcr} and crosses ΔP_{pcr} at Q_{max}. To find the equilibrium, we also draw AC_{pcr}, which is the average cost integrated from the right. The equilibrium is determined when D crosses ΔP_{pcr} and AC_{pcr} crosses AP_{pcr}. We find that in this case, individuals to the left of Q_{pcr} choose L and to the right demand H. The allocation results in large misallocation in the market, where the old individuals inefficiently sort in L and the young individuals inefficiently sort in H.

In Figures A1c and A1d, we plot the cases where MC is above D for the young individuals and below for the old individuals. Marone and Sabety (2022) points out that full community rating results in large inefficiency. In both figures, AC (AC_{fcr}) is above MC and crosses MC at Q_{fcr}. Under full community rating, most young individuals are inefficiently sorted in L, and the oldest individuals are inefficiently sorted in H.

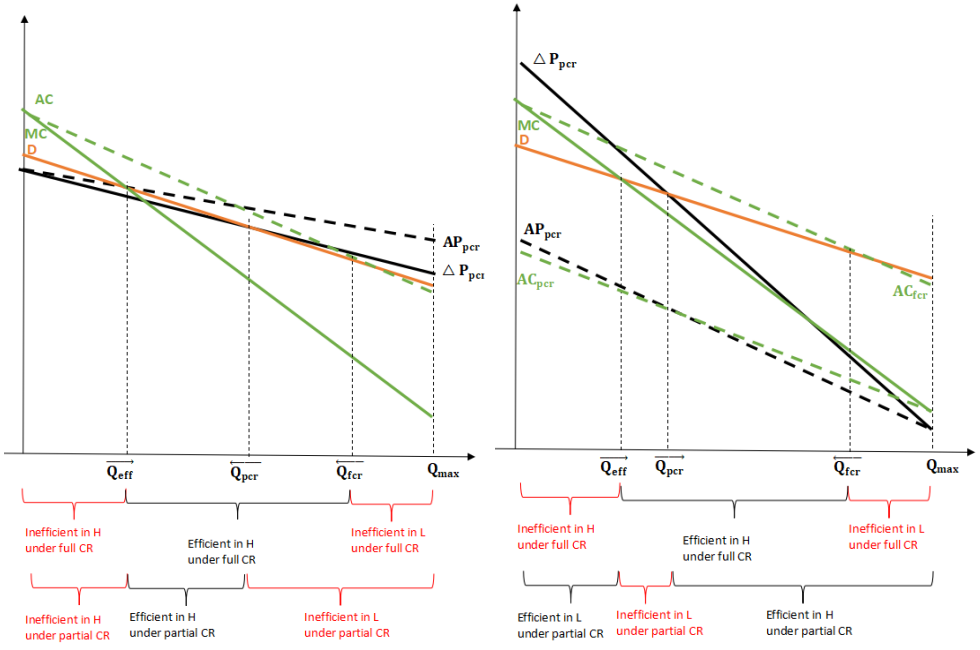
In Figure A1c, a is relatively small, resulting in $\Delta P_{pcr} < D$ for the old individuals. Partial community rating results in similar sizable misallocation, where there are both inefficient sorting by young individuals in L and inefficient sorting in H by the old individuals. The inefficiency has a similar direction as the full community rating case, while the size is larger. However, when a becomes larger, and ΔP_{pcr} becomes steeper than D, as shown in Figure A1d, partial community rating corrects the misallocation direction and Q_{pcr} becomes closer to Q_{eff}. The increase in relative premium differentials under partial community rating greatly mitigates the welfare loss from full community rating.

This section again shows that partial community rating may result in a more or less efficient allocation than full community rating, depending on the relative positions of the demand for the extra coverage and its marginal costs. Consistent with our primary findings in the main text, we find that the efficiency does not change monotonically with regard to the partial age rating factor, with efficient allocation achieved at a partial community rating level that falls between full community rating and the ratio of the total uncovered losses between types.

Figure A1: Equilibrium Under Community Rating and No Risk Adjustment



(a) $D > MC$ and $D > \Delta P_{pcr}$ for the old (b) $D > MC$ and $D < \Delta P_{pcr}$ for the old



(c) $D < MC$ and $D > \Delta P_{pcr}$ for the old (d) $D < MC$ and $D < \Delta P_{pcr}$ for the old

Note: The figures plot the equilibrium under no risk adjustment, where Q_{eff}, Q_{fcr}, and Q_{fcr} indicate the quantity under efficient equilibrium, full community rating equilibrium, and partial community rating equilibrium, respectively. The arrow points to the individuals who select H under the corresponding equilibrium.

Appendix B Estimating Fair Premiums

Construction of Medical Expenditure Distribution for All Ages Using Truven Data In the Truven MarketScan data, we take all individuals with full-year enrollment in 2013, and group them by age. For each age group, we create discretized medical expenditure distributions in the following way: we construct 84 expenditure bins as in the AV Calculator continuation table, 0, 0-100, 100-200, etc. For each bin s, we calculate their average medical expenditure x_{is}, and the probability of being in that bin, μ_{is}.

Adjust Truven Distributions to Match California Premiums We then mean-shifted the medical expenditure distributions to match the observed California premium levels using the following steps.

First, we calculate the 2017 average premiums for Bronze, Silver, and Gold plans in California. Specifically, among all CA enrollees, we calculate the fraction of individuals in each carrier-rating area and year. We use these as weights to calculate the weighted average observed premiums of plans in 2017, separately for each metal tier, for 21-year-olds. We then use this base premium and the published age rating factors to calculate the premium levels for all ages for that metal tier. Let p^o_{im} denote plan m's premium for age i in the data, where m is Bronze, Silver, and Gold.

Next, we assume that all ages have the same inflation factor h_m for plan m. Let x_i denote the uncovered expenditure level for age i. We multiply x_i by h_m, and then calculate the community rated premiums, p^s_{im} in 2017 as a function of h_m. Specifically, let λ_{im} denote the expected covered costs for age i under plan m. We calculate λ_{im} as

$$\lambda_{im}(h_m) = \sum_s oop_m(h_m x_{is})\mu_{is}, \tag{15}$$

where $oop_m(\cdot)$ is the cost-sharing design of plan m. Given $\lambda_{im}(h_m)$, we then calculate the community rated premiums as follows:

$$p^{cr}_{im}(h_m) = a_i \frac{\sum_{s=21}^{64} \gamma_s \lambda_{sm}(h_m)}{\sum_{s=21}^{64} \gamma_s a_s}, \ \forall i, m, \tag{16}$$

where γ_i is the population weights of age i, and a_i is the age-rating factor. The resulting p^{cr}_{im}may result in a non-neutral budget for the system. Let ω_{im} denote the fraction of individuals with age i choosing plan m (as a fraction of the total population.) Then we calculate the budget of

the system under p_{im}^{cr} as

$$b\left(h_m\right) = \sum_i \sum_m \omega_{im} \lambda_{im}\left(h_m\right) - \sum_i \sum_m \omega_{im} p_{im}^{cr}\left(h_m\right). \tag{17}$$

We adjust the levels of premiums of all plans by $b\left(h_m\right)$ such that the budget is neutral given observed enrollment shares:

$$p_{im}^{s}\left(h_m\right) = p_{im}^{cr}\left(h_m\right) + b\left(h_m\right). \tag{18}$$

Finally, we solve for h_m by minimizing $\sum_m \sum_i \left(p_{im}^{s} - p_{im}^{o}\right)^2$. To simplify, we use Bronze, standard Silver, and Gold plans for the calculation and ignore the Platinum plan due to its small share and uncommon presence in the market.[28] We find that the solved level adjustment b is close to zero, suggesting that even without the adjustment, the system is budget neutral (as it should be in reality.)

In Figure B1, we plot the estimated h_m against AV for Bronze, Standard Silver, and Gold plan designs in red dots. We find that h_ms for all plans are greater than one, suggesting that the California distribution has a larger expenditure level than the Truven distribution. We also find that the h_m increases with AV, which suggests that more generous plan cost-sharing designs induce more moral hazard responses. We then fit a linear line on these three points and extrapolate the value for all other plan designs.[29]

Calculating Fair Premiums For All Plans We calculate the market average fair premium for each metal tier as p_{im}^{fair} at the solved h_m. These values are used in Section 4, when we demonstrate the distortions in relative premiums, and in Section 5.2, when we calculate the counterfactual premiums under no enrollment change. In the counterfactual calculation with enrollment change, we need to construct the fair premium for each plan faced by enrollees, which may differ from the market average because plan premiums may vary by rating areas.[30] Let p_{imr}^{o} denote the observed premium for an age i of metal tier m in rating area r. We assume that the expected covered cost for age i under plan m in rating area r is $\lambda_{imr} = \theta_r \lambda_{im}$, a scaled version of the market average expected cost. According to equation (16), we have $p_{imr}^{o} = \theta_r p_{im}^{o}$. Thus, we calculate

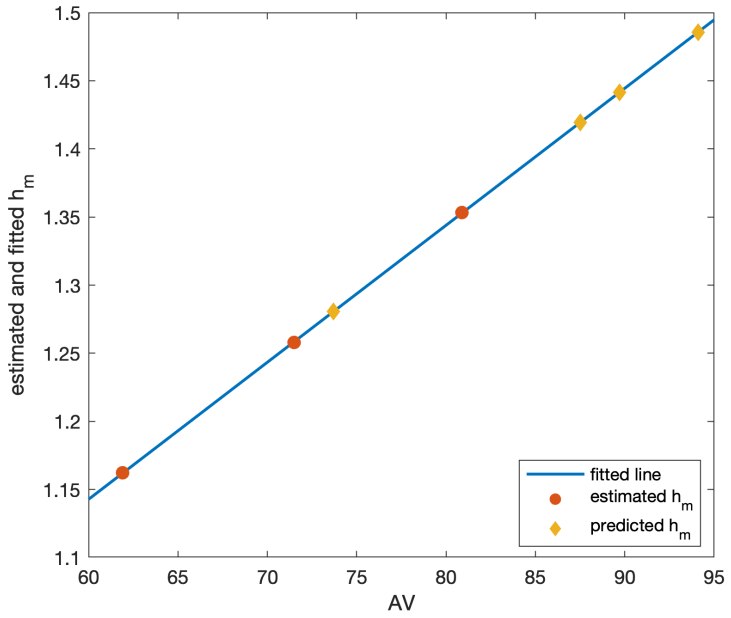
$$\lambda_{imr} = \frac{\lambda_{im}}{p_{im}^{o}} p_{imr}^{o}. \tag{19}$$

[28] Note that the extra costs of CSR Silver plans (relative to the standard Silver plan) are fully reimbursed by the government, thus should not be included in the calculation.

[29] The AVs of plans are reported in the official Covered California documents and might be different from the regulated AV due to the allowed error margin.

[30] Plan premiums may also vary by carriers, network types, and other plan attributes. However, we aggregate plan options into metal tiers; thus, we could ignore the premium variation along these dimensions.

Figure B1: Estimated and Predicted h_m



Notes: The dots from left to right represent: Bronze, Standard Silver, CSR73 Silver, Gold, CSR88 Silver, Platinum, and CSR94 Silver.

Appendix C Counterfactual Calculation Details

Demand Estimation To simplify the choice set, we aggregate all plans into four metal tiers - Bronze, Silver, Gold, and Platinum - to focus on vertical selection. Additionally, individuals can remain uninsured, which we consider the outside option. The final choice set thus includes five alternatives: the four tiered plan options and the outside option. For each plan option, we observe its base premium and non-price attributes, including the fraction of plans that are HMO, PPO (with the baseline being the other network types), and the fraction of plans offered by the three top insurers (Blue Shield, Kaiser, and Anthem).

Our estimation focuses on single-person households with individuals aged 25 to 64. We restrict attention to this group to avoid complications related to family-level plan selection and Medicare eligibility. Inertia is a common issue in the market. To rule out its impact and focus on the active choices made by enrollees, we use the 2014 data, the first year Covered California was launched.

Let i index individuals and j index plan alternatives. The indirect utility that individual i derives from choosing plan j is specified as follows:

$$u_{ij} = \mathbb{1}_{\{j\neq 0\}} \cdot (-\alpha_i p_{ij} + \beta_i AV_{ij} + X_j'\gamma^x + \delta_i) + \varepsilon_{ij}, \tag{20}$$

where $p_{ij} = P_j(\mathbf{b}, \mathbf{z}_i)$ is the net premium faced by individual i, which depend on the base premium, $\mathbf{b}$, and individual's age, income and rating regions: $\mathbf{z}_i = \{age_i, inc_i, region_i\}$. The function P captures partial community rating by age and income-linked premium subsidies by regulation:

$$p_{ij} = P_j(\mathbf{b}, \mathbf{z}_i) = \max\{a_i b_j - \max\{a_i b_{2ls,i} - cap_i\}, 0\}, \tag{21}$$

where a_i is the age rating factor, $b_{2ls,i}$ is the baseline premium of the second-lowest Silver plan in a rating area of individual i, and cap_i is the enrollee's income contribution limit. According to the regulations, premiums in Covered California can only vary across 19 rating areas and age. AV_{ij} is the AV of a plan faced by individuals, which varies by individual income for cost-sharing reduction plans. X_j represents other plan characteristics, including a dummy for the Silver-tier plan, the fraction of plans of a particular network type, and carriers. The term $\mathbb{1}_{\{j\neq 0\}}$ is an indicator for whether the individual chooses any plan rather than the outside option (denoted $j = 0$), and ε_{ij} captures unobserved preferences.

In the indirect utility function, δ_i is an individual-specific intercept, while α_i and β_i are individual-specific slope coefficients. We define coarse age bins as $\{26-31, 32-37, 38-43, 44-49, 50-55, 56-64, 62-64\}$, and coarse income bins (federal poverty line) as $\{100-137, 138-150, 151-200, 201-250, 251-400, 401+\}$. We allow both α_i and δ_i to vary by rating area, coarse age bins and coarse income bins: $\alpha_i = \alpha_i^{inc} + \alpha_i^{age} + \alpha_i^{region}$ and $\delta_i = \delta_i^{inc} + \delta_i^{age} + \delta_i^{region}$. We allow β_i to vary by coarse age groups: $\beta_i = \beta_i^{age}$. This specification enables us to capture heterogeneous preferences for plan premiums, generosity (AV), and the propensity to take up insurance, varying by coarse age brackets, income levels, and geographic regions.

We estimate the model using a random 5% sample drawn from the 2014 Covered California enrollment data. In estimation, we assume the error terms follow a generalized extreme value distribution, and the model is a logit model at the enrollee level. A common identification issue is the endogeneity of the premium variable. We identify the premium coefficient leveraging nonlinearities in the premium equation (21). For example, some enrollees are eligible for zero-premium Bronze plans when the premium subsidy exceeds the full premiums. Similarly, we identify the AV coefficient by leveraging the discontinuity in eligibility for CSR plans (Saltzman, Swanson and Polsky, 2021; Tebaldi, 2025): at FPL levels of 150, 200, and 250, individuals face a discontinuity in the AV of the Silver plans they can choose from.

Table C3 shows the estimation results. Column (1) displays the coefficients when we do not include the interactives. Column (2) shows the results when the interactive terms are included. We find that the AV coefficients increase with age, consistent with estimates in

previous literature (Tebaldi, 2025). Table C4 summarizes the extensive-margin semi-elasticity of demand—measured as the percentage drop in the likelihood of choosing marketplace coverage if all plans' premiums increase by $10 per month, and the average own-price elasticity of demand for Silver plans, equal to the percentage drop in the share of buyers selecting a Silver plan if the plan's premium increases by 1%. We find that individuals with higher incomes tend to be more price-sensitive. We also find that middle-aged individuals are the most price-sensitive. The estimates are close to what's estimated in the previous literature, for example, Tebaldi (2025) and Saltzman (2019).

Table C3: Demand Estimation Results

	(1)	(2)		(1)	(2)
Monthly Premium ($100)	−0.49***	−1.18***	inside	−2.96***	−13.61***
	(0.01)	(0.09)		(0.15)	(0.62)
prm X FPL 138		−0.40***	inside X Age 32		−0.60***
		(0.10)			(0.15)
prm X FPL 151		−0.56***	inside X Age 38		−1.13***
		(0.09)			(0.17)
prm X FPL 201		−0.57***	inside X Age 44		−1.16***
		(0.09)			(0.15)
prm X FPL 251		−0.36***	inside X Age 50		−1.61***
		(0.08)			(0.14)
prm X FPL 401		−0.07	inside X Age 56		−2.04***
		(0.08)			(0.14)
prm X Age 32		0.12***	inside X Age 62		−2.00***
		(0.03)			(0.17)
prm X Age 38		0.14***	inside X FPL 138		1.88***
		(0.03)			(0.07)
prm X Age 44		0.24***	inside X FPL 151		2.87***
		(0.03)			(0.06)
prm X Age 50		0.40***	inside X FPL 201		3.90***
		(0.02)			(0.07)
prm X Age 56		0.62***	inside X FPL 251		4.23***
		(0.02)			(0.08)
prm X Age 62		0.72***	inside X FPL 401		3.52***
		(0.02)			(0.13)
prm X Region 2		0.10***	inside X Region 2		2.61***
		(0.03)			(0.15)
prm X Region 3		0.12***	inside X Region 3		2.30***
		(0.03)			(0.21)
prm X Region 4		0.24***	inside X Region 4		2.06***
		(0.03)			(0.14)
prm X Region 5		0.16***	inside X Region 5		2.33***
		(0.03)			(0.24)
prm X Region 6		0.09**	inside X Region 6		0.12
		(0.03)			(0.08)
prm X Region 7		0.11***	inside X Region 7		3.16***
		(0.03)			(0.24)
prm X Region 8		0.27***	inside X Region 8		2.13***
		(0.03)			(0.24)
prm X Region 9		0.29***	inside X Region 9		2.78***
		(0.03)			(0.63)
prm X Region 10		0.12***	inside X Region 10		1.83***
		(0.03)			(0.26)
prm X Region 11		0.13**	inside X Region 11		1.00***
		(0.04)			(0.25)
prm X Region 12		0.00	inside X Region 12		0.77**
		(0.03)			(0.24)
prm X Region 13		−0.24	inside X Region 13		0.26
		(0.14)			(0.29)
prm X Region 14		0.11*	inside X Region 14		1.26***
		(0.05)			(0.27)
prm X Region 15		0.03	inside X Region 15		2.34***
		(0.03)			(0.14)
prm X Region 16		0.02	inside X Region 16		2.56***
		(0.03)			(0.14)
prm X Region 17		0.04	inside X Region 17		2.42***
		(0.03)			(0.21)
prm X Region 18		0.02	inside X Region 18		2.03***
		(0.03)			(0.11)
prm X Region 19		0.07*	inside X Region 19		2.99***
		(0.03)			(0.16)

	(1)	(2)
Silver Plan	1.51***	1.32***
	(0.01)	(0.02)
HMO	0.19	0.53
	(0.16)	(0.51)
PPO	−0.54***	−2.08**
	(0.07)	(0.67)
Blue Shield	2.58***	9.84***
	(0.43)	(1.54)
Kaiser	0.89***	8.03***
	(0.17)	(1.63)
Anthem	−1.68***	−0.35
	(0.17)	(0.42)
AV	0.78***	5.67***
	(0.06)	(0.16)
AV X Age 32		0.32
		(0.19)
AV X Age 38		1.00***
		(0.20)
AV X Age 44		1.17***
		(0.19)
AV X Age 50		1.93***
		(0.18)
AV X Age 56		2.53***
		(0.17)
AV X Age 62		2.42***
		(0.21)
Num.Obs.	644 490	644 490
AIC	175 434.5	163 447.5
RMSE	0.45	0.43

Table C4: Summary of Elasticity and Semi-Elasticity Estimates By Age And Income Groups

	All tiers	Silver
Overall	-2.14	-2.33
Age		
26-31	-1.92	-2.18
32-37	-2.24	-2.44
38-43	-2.37	-2.54
44-49	-2.40	-2.59
50-55	-2.26	-2.45
56-61	-1.90	-2.08
62-64	-1.76	-1.94
Income (FPL)		
100-137	-0.32	-0.27
138-150	-0.46	-0.50
151-200	-0.79	-0.95
201-250	-1.51	-1.94
251-400	-2.87	-3.16
400+	-3.37	-3.42

Note: The column "All tiers" shows the percentage change in enrollment in any plans if all plans' premiums increase by $10/month. The column "Silver" shows the percentage change in enrollment in the Silver plan if its premium increases by one percent.

We examine the out-of-sample model fit by drawing a 10% random sample from the 2014 data, which are not used in estimation. Table C5 shows how our model predicts the tier shares by age. Overall, our model predictions are close to the data.

Table C5: Model Fit

Age	Type	No Insurance	Bronze	Silver	Gold
26 - 31	data	0.76	0.07	0.14	0.01
	model	0.76	0.06	0.15	0.02
32 - 37	data	0.83	0.05	0.10	0.01
	model	0.83	0.04	0.11	0.01
38 - 43	data	0.84	0.04	0.10	0.01
	model	0.84	0.04	0.10	0.01
44 - 49	data	0.83	0.04	0.11	0.01
	model	0.83	0.04	0.11	0.01
50 - 55	data	0.80	0.05	0.13	0.01
	model	0.80	0.05	0.13	0.01
56 - 61	data	0.74	0.06	0.17	0.01
	model	0.74	0.06	0.16	0.02
62 - 64	data	0.73	0.07	0.17	0.01
	model	0.73	0.07	0.17	0.02

Calculating the Equilibrium Under Three Pricing Schemes We calculate the counterfactual equilibrium using the 2017 plan designs and enrollees. We draw a 10% random sample from the 2017 Covered California potential buyers. Let $\Phi = \{\varphi_i, i = 1, 2, ..., N\}$ denote the allocation of plan choices to each enrollee, where $\varphi_i = \{s_{im}, m = 0, 1, 2, 3, 4\}$ is a vector denoting individual i's probability of choosing each metal tier and the outside option. Let $\mathbf{p} = \{p_{im}, m = 0, 1, 2, 3, 4\}$ denote the premium of each plan m for individual i. Let a_i denote the age-rating factor under the current pricing schemes, which depends on individual i's age.

We define two functions mapping between enrollment allocation and market premiums. First, let $\mathcal{D}(\cdot)$ denote the mapping from premiums $\mathbf{p}$ to market shares. Function $\mathcal{D}$ is derived based on demand estimates, holding all other plan attributes and plan subsidy rules constant, including CSR eligibility and APTC calculation rules in the market.

Second, let $\mathcal{P}(\cdot; \iota)$ denote the mapping from the enrollment allocation to the premiums, under the pricing scheme ι and the assumption that insurers are perfectly competitive and break even. A competitive equilibrium is defined as a pair of $(\Phi, \mathbf{p})$ such that $\Phi = \mathcal{D}(\mathbf{p})$ and $\mathbf{p}(\iota) = \mathcal{P}(\Phi; \iota)$. In equilibrium, potential consumers optimize by maximizing their expected utility, and insurers earn zero profits.

We now define the function $\mathcal{P}(\cdot; \iota)$ for each pricing scheme. Fix λ_{im}, the fair premiums of each plan for each individual. The fair premiums are obtained using methods specified in Appendix A. Under partial community rating with age-rating factors a_i, $\mathcal{P}(\cdot)$ is calculated as follows. For the outside option, the premium is always zero: $p^{cr}_{i0} = 0, \forall i$. For plans other than the outside option, i.e., $m \neq 0$,

$$p^{cr}_{im} = \max\{\tilde{p}^{cr}_{r_i m} - \max\{a_i b_{2ls,i} - cap_i\}, 0\}, \tag{22}$$

$$\tilde{p}^{cr}_{im} = a_{r_i} \frac{\sum_{s=21}^{64} \gamma_s \lambda_{sm}}{\sum_{s=21}^{64} \gamma_s a_s} + \tau^{cr}, \tag{23}$$

where $b_{2ls,i} = b^o_{2ls,i} + \tau^{cr}$ is the base premium of the second-lowest Silver plan in the rating area where individual i lives, and $b^o_{2ls,i}$ is the observed value from data. $\tilde{p}^{cr}_{im}$ is the premium before the APTC payment. γ_s is the share of enrollees of the same age as individual i among those who choose a plan, implied by Φ. We apply the formula to standard plans and assign the values of the standard Silver plans to their respective CSR plans. The per-enrollee budget for the insurer is calculated as the difference between the cost and revenue: $\tau^{cr} = \frac{\sum_i \sum_{m \neq 0} s_{im}(\lambda_{im} - \tilde{p}^{cr}_{im})}{\sum_i (1 - s_{i0})}$.

We solve for the equilibrium by starting with an initial value of Φ (we use the observed data). Given the initial value, Φ_0, we calculate $\mathbf{p}^{cr}_0 = \mathcal{P}(\Phi_0; \iota_{cr})$, and update Φ_1 using $\mathcal{D}(\mathbf{p}_0)$.

We iterate the process until τ_{cr} is within the tolerance level. Once we find the solution, we then calculate the age-specific transfer, $T_i^{cr} = \frac{\sum_{r_i=r}\sum_m s_{im}(\lambda_{im}-\tilde{p}_{im}^{cr})}{\sum_{r_i=r}(1-s_{i0})}$. We hold fixed T_i^{cr} for the other two pricing schemes.

Under the risk-based pricing schemes ι^{rb}, $\mathcal{P}(\cdot;\iota^{rb})$ is defined as:

$$p_{im}^{rb} = \max\{\tilde{p}_{im}^{rb} - \max\{b_{2ls,i}^{rb} - cap_i\}, 0\}, \tag{24}$$

$$\tilde{p}_{im}^{rb} = \lambda_{im} - T_i^{cr} + \tau^{rb}, \tag{25}$$

where $b_{2ls,i}^{rb}$ is the risk-based premium for the second-lowest Silver plan, calculated as the fair premium for the plan plus τ^{rb}. The fair premium for the second-lowest Silver plan is calculated using equation (19). The per-enrollee budget for the insurer is calculated as the difference between the cost and revenue: $\tau^{rb} = \frac{\sum_i\sum_{m\neq 0} s_{im}(\lambda_{im}-\tilde{p}_{im}^{rb})}{\sum_i(1-s_{i0})}$. We solve for the equilibrium by starting with an initial value of Φ (we use the equilibrium allocation under the partial community rating calculated above). Given the initial value, Φ_0, we calculate $\mathbf{p}_0^{rb} = \mathcal{P}(\Phi_0;\iota^{rb})$, and update Φ_1 using $\mathcal{D}(\mathbf{p}_0)$. We iterate the process until τ^{rb} is within the tolerance level.

Under the full community rating schemes ι^{full}, $\mathcal{P}(\cdot;\iota^{full})$ is defined as:

$$p_{im}^{full} = \max\{\tilde{p}_{im}^{full} - \max\{b_{2ls,i}^{full} - cap_i\}, 0\}, \tag{26}$$

$$\tilde{p}_{im}^{full} = \bar{p}_m - \Delta T_i + \tau^{full}, \tag{27}$$

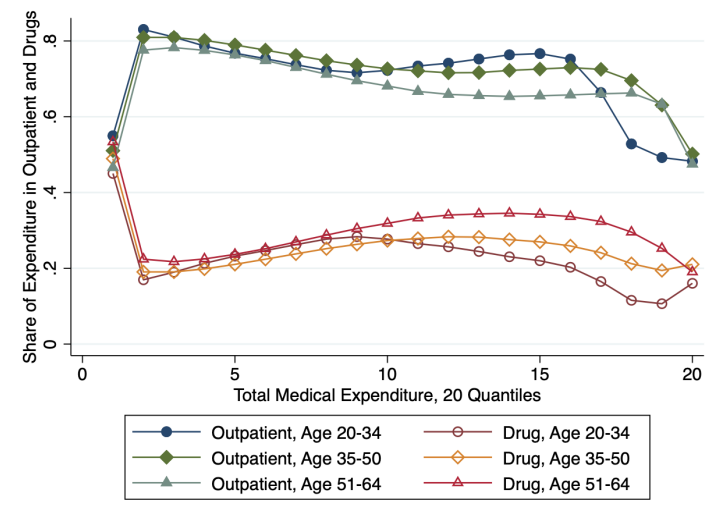
where $\bar{p}_m$ is the market average premium, calculated as the weighted average of plan m's premiums in the rating area where enrollee i lives for an age, where the weights are the age share among those who choose a plan in the whole market. $b_{2ls,i}^{full}$ is the full-community rated premium for the second-lowest Silver plan, calculated as the full-CR premium for the plan plus τ^{full}. We calculate the age-specific transfer term as $\Delta T_i = T_i^{cr} - T_i^{full}$, where T_i^{full} is age-specific, calculated as $\frac{\sum_i\sum_m s_{im}(\lambda_{im}-\tilde{p}_{im}^{full})}{\sum_i(1-s_{i0})}$ among enrollees of a particular age. The per-enrollee budget for the insurer is calculated as the difference between the cost and revenue: $\tau^{full} = \frac{\sum_i\sum_{m\neq 0} s_{im}(\lambda_{im}-\bar{p}_m-\tau^{full})}{\sum_i(1-s_{i0})}$.

We solve for the equilibrium by starting with an initial value of Φ (we use the equilibrium allocation under the partial community rating calculated above). Given the initial value, Φ_0, we calculate $\mathbf{p}_0^{full} = \mathcal{P}(\Phi_0;\iota^{full})$, and update Φ_1 using $\mathcal{D}(\mathbf{p}_0)$. We iterate the process until the process converges to a fixed point of $(\Phi, \mathbf{p})$.[31]

[31] We numerically find the fixed point by calculating the l^2-norm of the premiums between the two steps, and terminate once the distance is below the tolerance level. We try different starting points and ensure the program converges to the same fixed point.

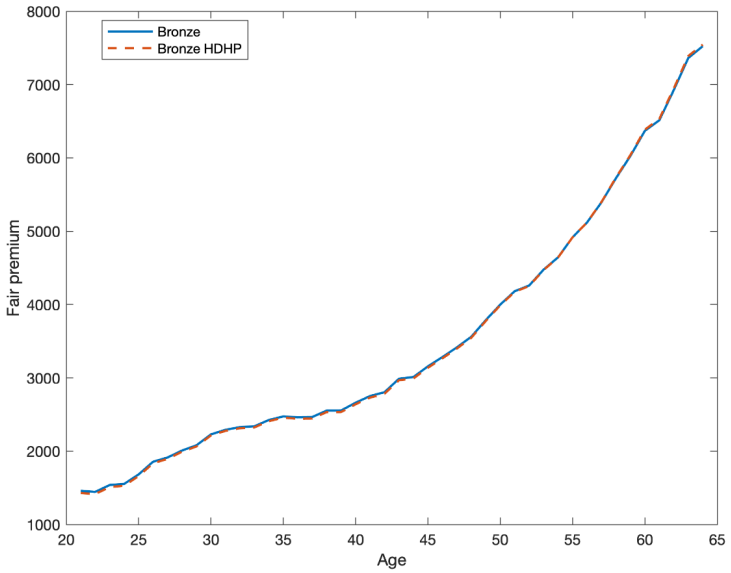
Appendix D More Figures for Empirical Analysis

Figure C1: Service Share and Total Medical Expenditure Level by Age



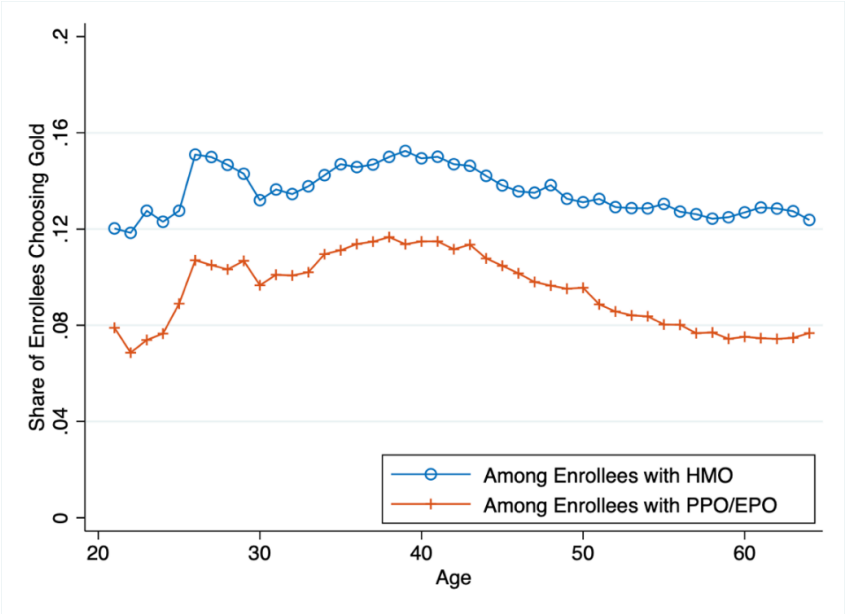
Notes: We use the 2013 Truven MarketScan data to plot the graph. We first group individuals with positive total medical expenditure into 20 quantiles. We then calculate the share of their expenditure from outpatient, drugs, and inpatient (ignored category). The y-axis shows the share by total expenditure quantiles, separately for three age groups: age 20-34, age 35-50, and age 51-64.

Figure C2: Simplifying Plan Designs: Robustness Checks



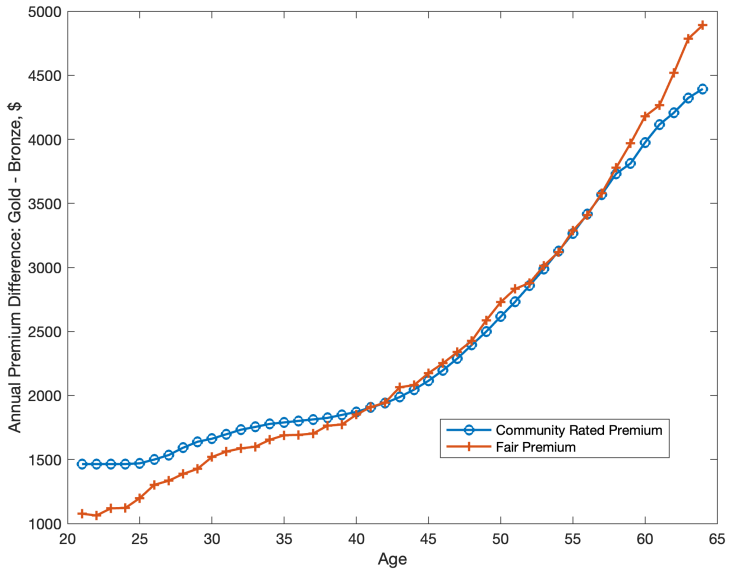
Notes: In California, and in some years, there might be two similar designs within a metal tier. For example, in 2017, there were two Bronze plans: one normal Bronze plan and another HDHP Bronze plan. The HDHP Bronze plan has a simplified three-arm cost-sharing rule by design. While the normal Bronze plan has more complicated designs, it is transformed into the three-arm design using the 2017 AV Calculator. We plot the premium levels for these two plans by age. All values are for annual individual coverage.

Figure C3: Share Choosing Gold by Plan Type and Age: California 2015-2020



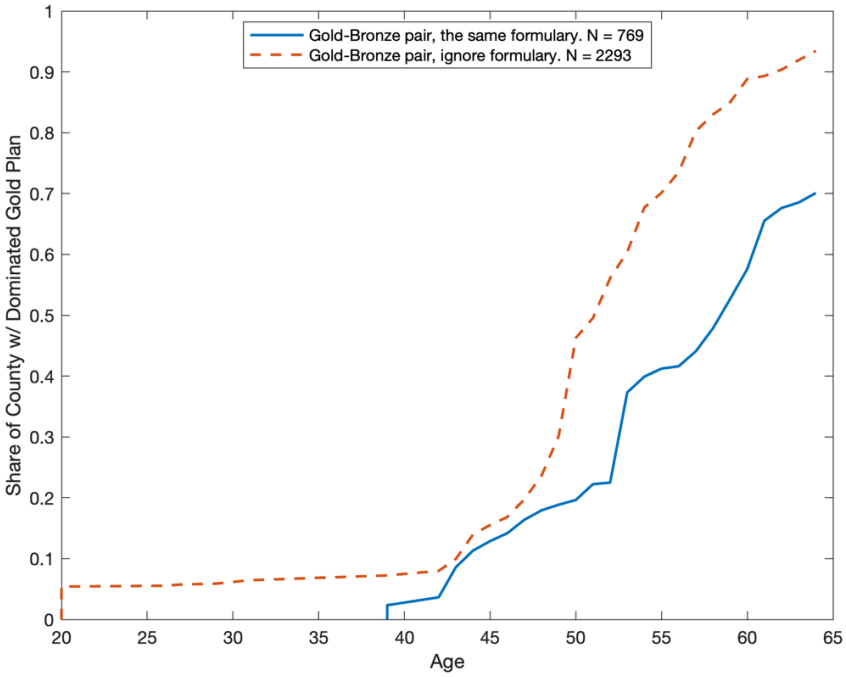
Notes: Data from enrollment information through Covered California 2015-2020. The sample includes individuals not eligible for CSR-variation plans. We calculate the share of enrollees choosing the Gold plan separately for two types of enrollees: those choosing an HMO, and those choosing a PPO.

Figure C4: Robustness Check: Linear Bronze Plan



Notes: We replace the 2017 Covered California Bronze plan by a constant coinsurance plan with the same AV (0.61). We then calculate the difference between the community-rated premiums and the fair premiums for the Gold plan and the counterfactual Bronze plan.

Figure C5: Share Facing Dominated Gold Plan by Age: Federal Exchange 2017

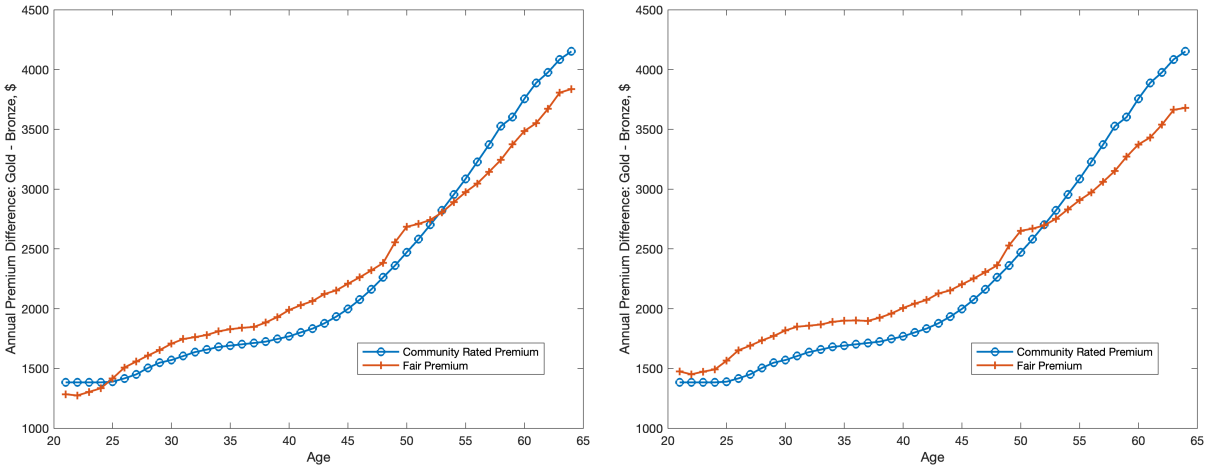


Notes: Data from states using Healthcare.gov in 2017. Plan designs are for individual, first-tier in-network coverage. The y-axis denotes the fraction of counties with at least one Gold plan dominated by a Bronze plan with the same other attributes for a particular age. We exclude counties without comparable Gold-Bronze pairs in each case, and the number of counties used in the calculation is labeled in the legend.

Figure C6: Robustness Check: Alternative Distributions

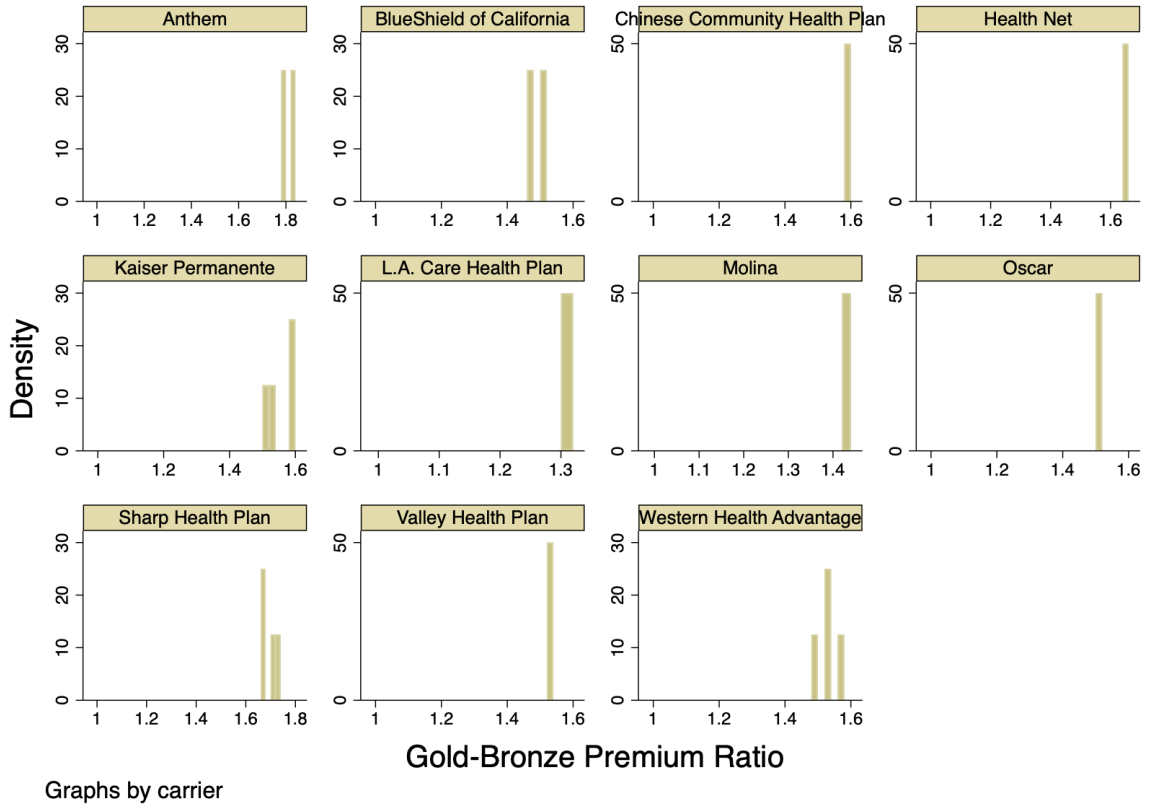
(a) Increasing Old-Young Gap

(b) Decreasing Old-Young Gap



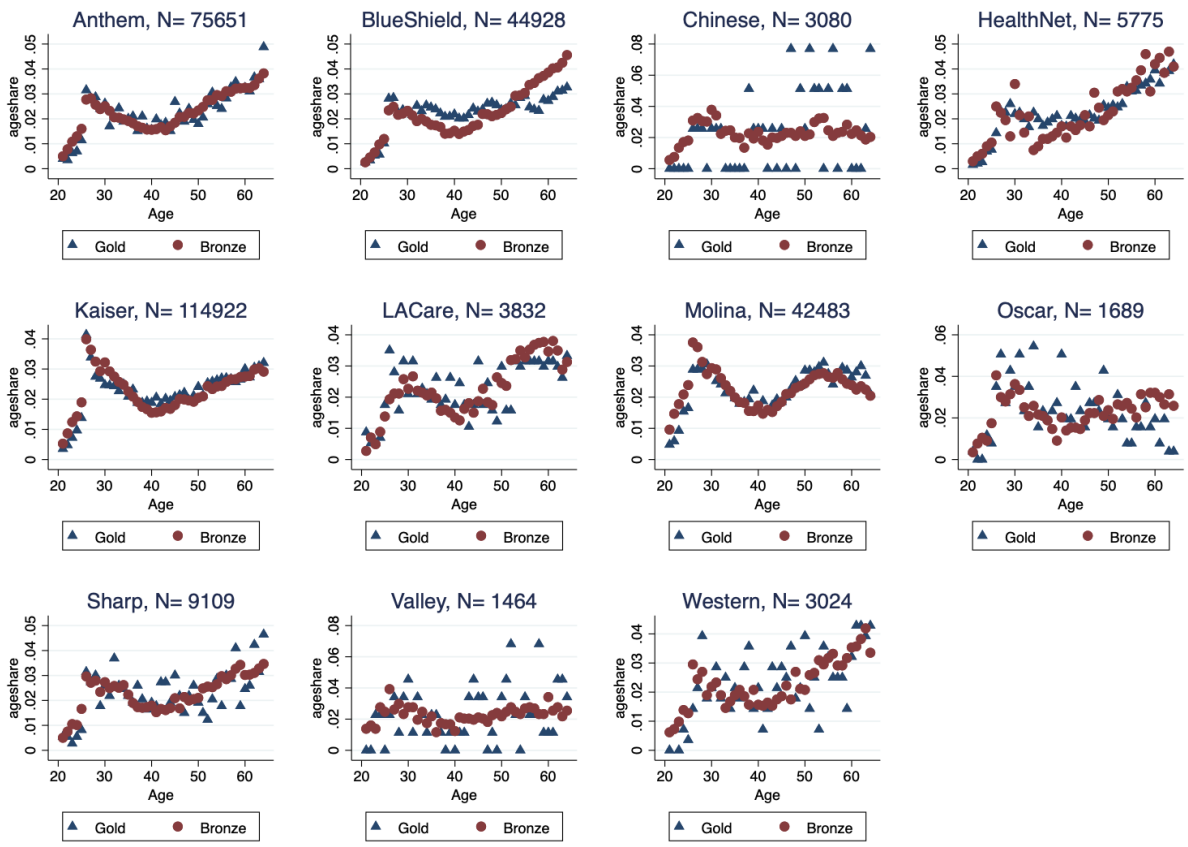
Notes: The figures show the community-rated and fair premium difference between the Gold and Bronze plans using alternative total medical loss distributions across ages. The overall across-age average is the same as the baseline. However, in panel (a), we increase the age differentials of the mean total uncovered medical expenditure, while in panel (b), we decrease the differentials.

Figure C7: Distribution of Gold-Bronze Premium Ratios by Insurers in 2017



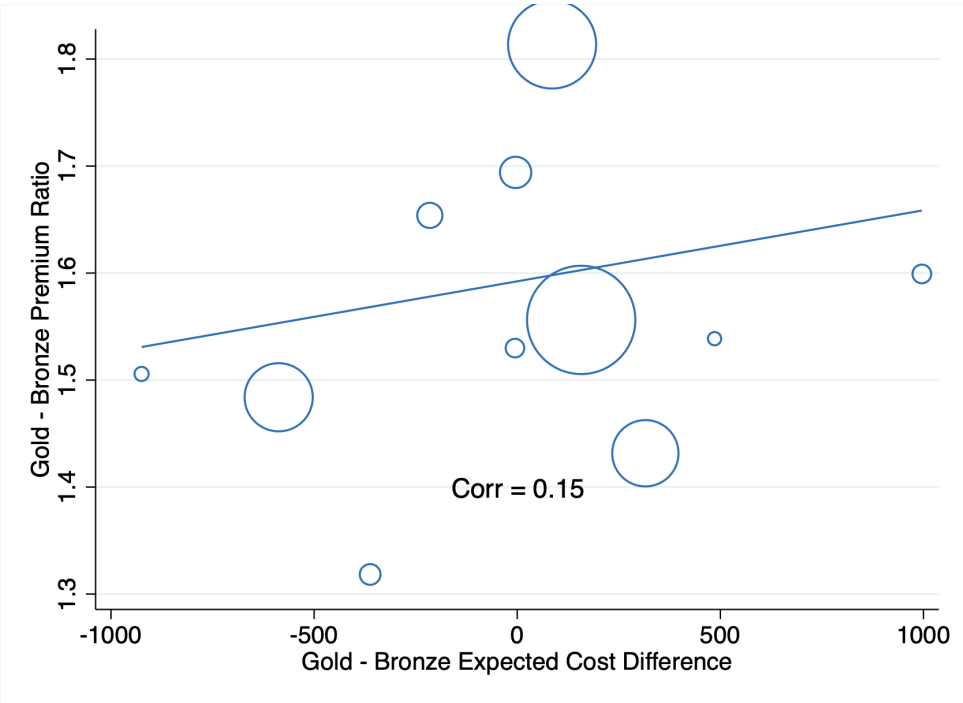
Notes: We calculate the Gold-Bronze premium ratios for all comparable pairs (having the same network type, offered by the same insurer, and in the same rating area), and plot the distribution by insurers in 2017.

Figure C8: Gold and Bronze Enrollees' Age Shares By Insurers



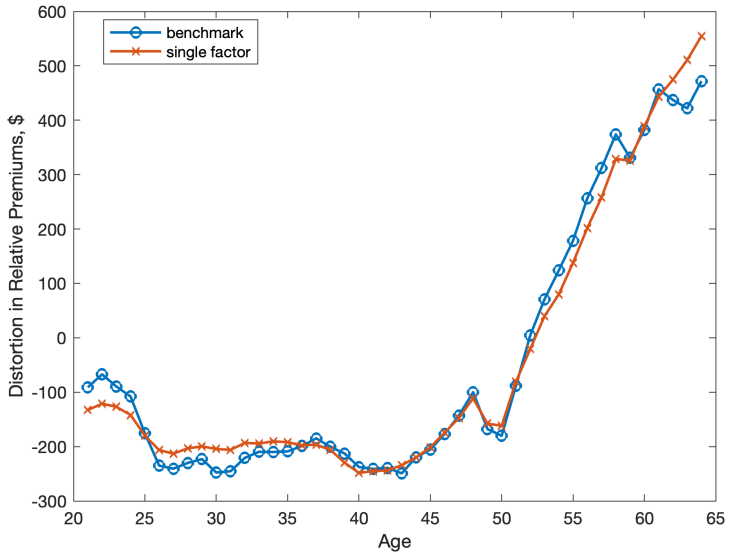
Notes: We calculate, for each 2017 insurer, the share of enrollees of a particular age if they choose that insurer's Bronze plan (red circle) or Gold plan (blue triangle).

Figure C9: Correlation between Gold-Bronze Premium Ratio and Expected Cost Difference



Notes: We plot, for each 2017 insurer, the Gold-Bronze premium ratio (mean if the insurers have multiple ratios for their comparable pairs) against the Gold-Bronze expected cost difference. The expected cost for each insurer and tier is calculated as the inner product of the expected total medical expenditure of each age (imputed from the baseline distribution) and a vector of the age shares of enrollees choosing that insurer and metal tier. The size of the bubble is proportional to the number of enrollees in that insurer. The line fit and the correlation number is weighted by the number of enrollees of that insurer.

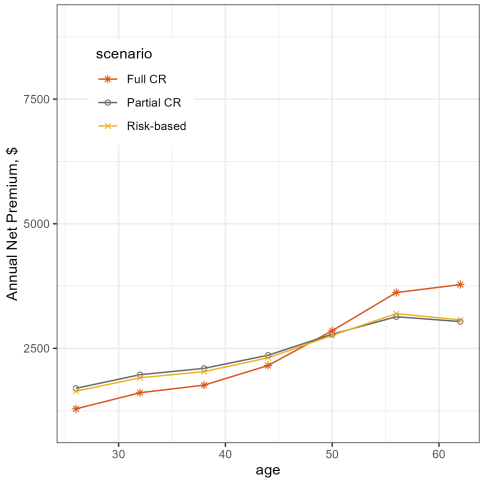
Figure C10: Robustness Check with Single Inflation Factor: Premium Distortions



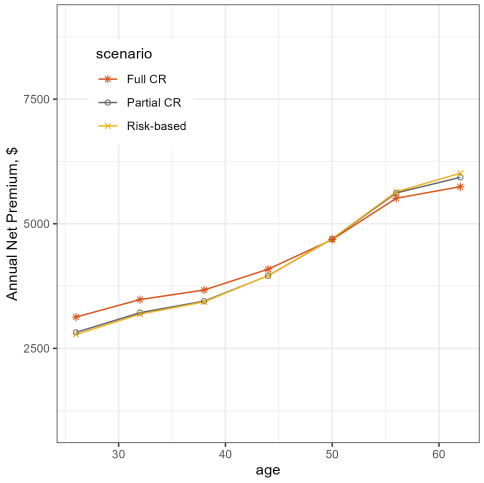
Notes: The y-axis shows the annual premium distortion, calculated as the difference between the Gold-brozne premium gap under the current partial community rating and the risk-based pricing. We plot the value separately for the baseline assumption ("benchmark") and under the assumption that there is a single inflation factor ("single factor").

Figure C11: Equilibrium Net Premiums under Three Pricing Schemes

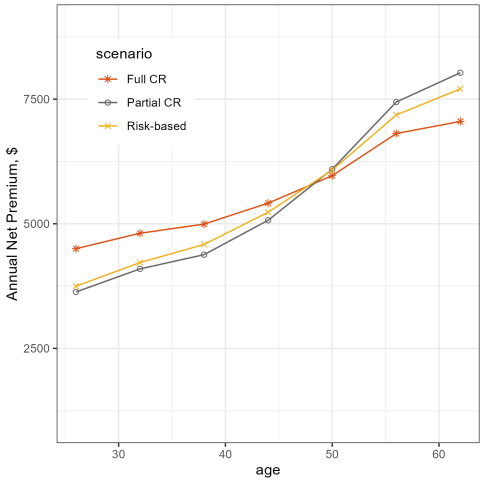
(a) Panel A. Bronze



(b) Panel B. Silver

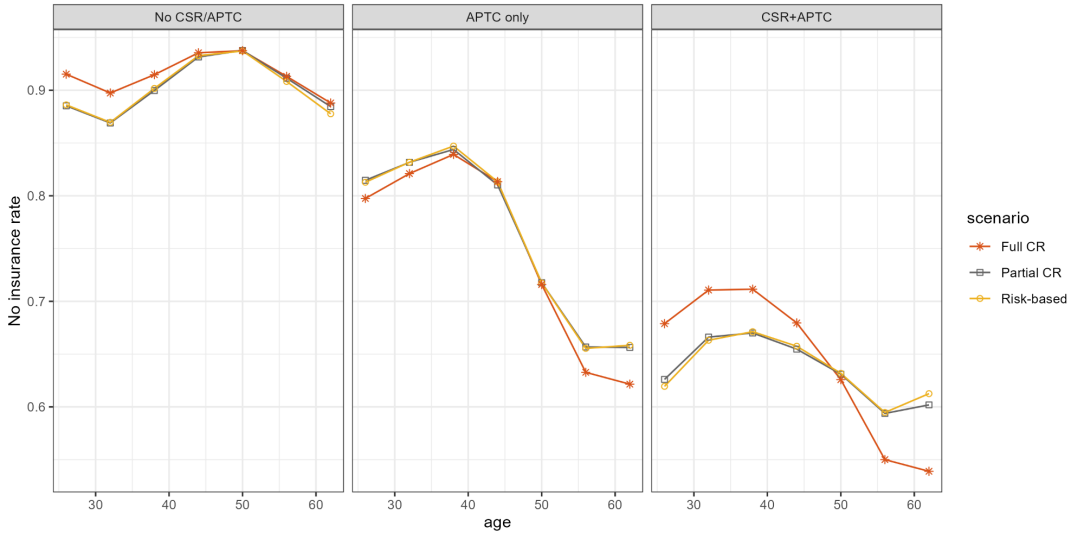


(c) Panel C. Gold



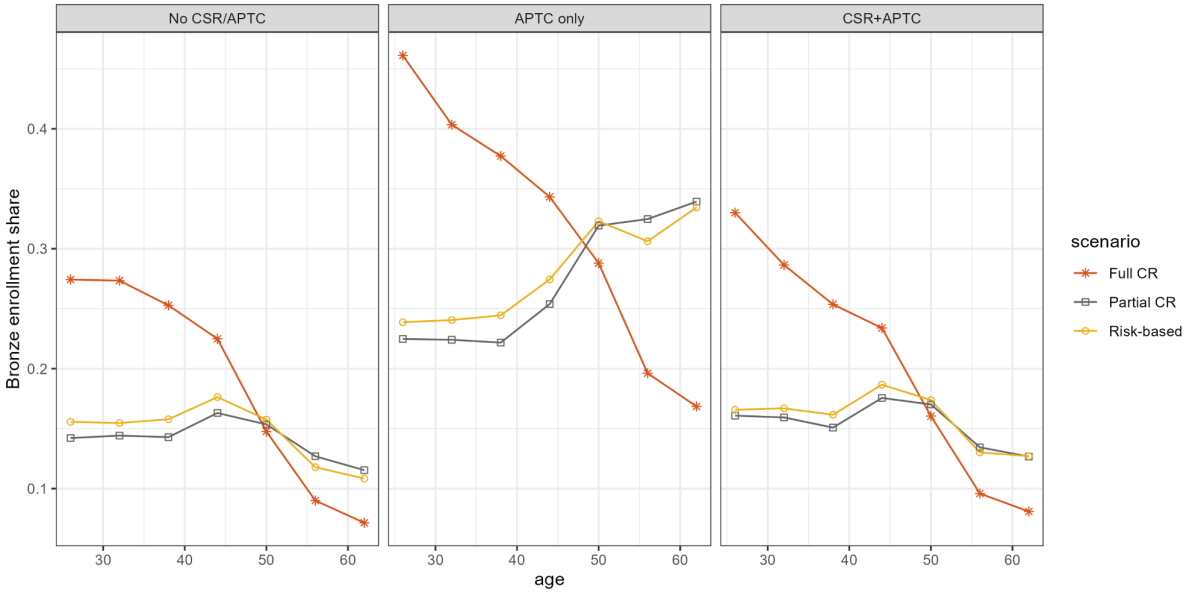
Note: The y-axis indicates the annual net premiums (total premiums minus the APTC subsidies) for individual coverage, averaged across the sample. "Full CR" refers to full community rating with age-specific side transfers. "Partial CR" refers to partial community rating by age (3:1 ratio). "Risk-based" refers to risk-based pricing with age-specific side transfers.

Figure C12: No Insurance Rates By Income



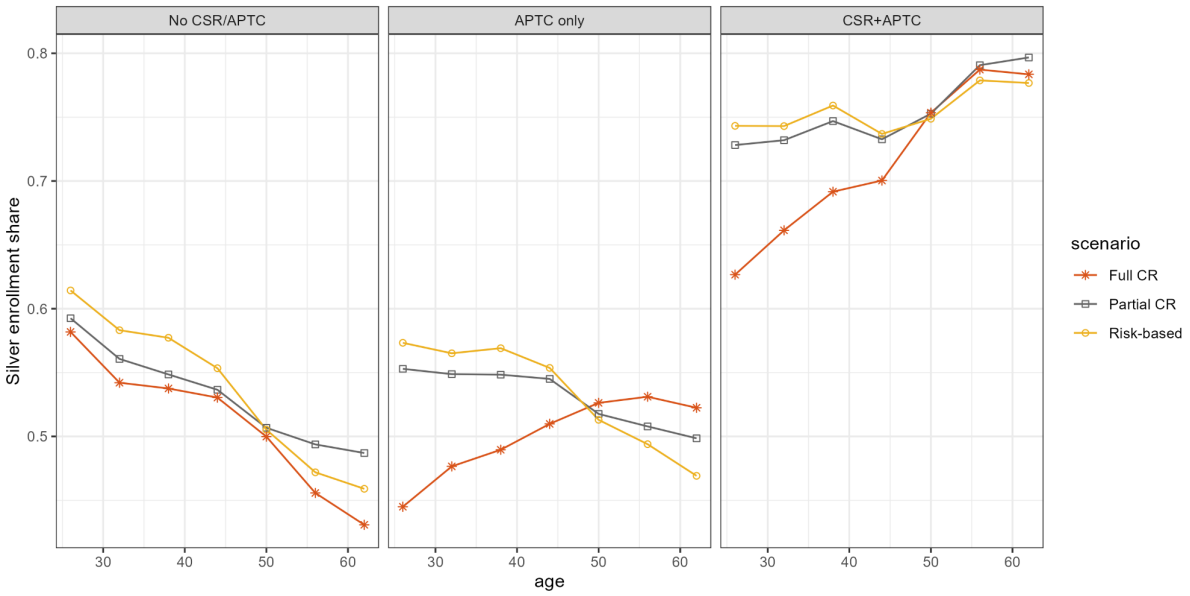
Notes: The three panels represent three types of enrollees by income. "No CSR/APTC" refers to enrollees whose income is above 400 FPL, so they are not eligible for CSR plans or APTC. "APTC" refers to middle-income enrollees (between 250 and 400 FPL) who are only eligible for APTC, and the "CSR+APTC" group refers to enrollees whose income is between 100 to 250 FPL and eligible for both.

Figure C13: Bronze Enrollment Shares Among Enrollees Choosing A Plan By Income



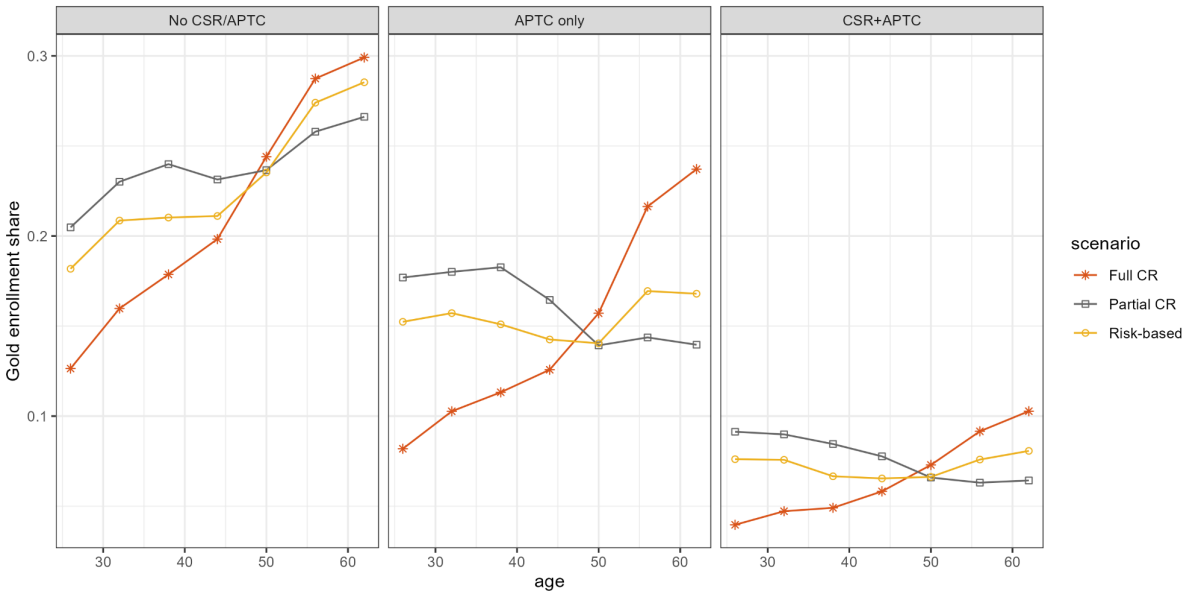
Notes: The three panels represent three types of enrollees by income. "No CSR/APTC" refers to enrollees whose income is above 400 FPL, so they are not eligible for CSR plans or APTC. "APTC" refers to middle-income enrollees (between 250 and 400 FPL) who are only eligible for APTC, and the "CSR+APTC" group refers to enrollees whose income is between 100 to 250 FPL and eligible for both.

Figure C14: Silver Enrollment Shares Among Enrollees Choosing A Plan By Income



Notes: The three panels represent three types of enrollees by income. "No CSR/APTC" refers to enrollees whose income is above 400 FPL, so they are not eligible for CSR plans or APTC. "APTC" refers to middle-income enrollees (between 250 and 400 FPL) who are only eligible for APTC, and the "CSR+APTC" group refers to enrollees whose income is between 100 to 250 FPL and eligible for both.

Figure C15: Gold Enrollment Shares Among Enrollees Choosing A Plan By Income



Notes: The three panels represent three types of enrollees by income. "No CSR/APTC" refers to enrollees whose income is above 400 FPL, so they are not eligible for CSR plans or APTC. "APTC" refers to middle-income enrollees (between 250 and 400 FPL) who are only eligible for APTC, and the "CSR+APTC" group refers to enrollees whose income is between 100 to 250 FPL and eligible for both.