Community Rating and Distortions of Relative Prices in Insurance Menus

Honglin Li (University of Wisconsin, Madison)[1]

Chenyuan Liu (Tsinghua University)

Justin Sydnor (University of Wisconsin, Madison & NBER)

This paper investigates the distortionary effects of community rating on relative prices in insurance markets with multiple products. We show that common community-rating approaches may generate counterintuitive distortions in the price *differentials* between plan options that run in the opposite direction as the distortions in overall price levels. We document that partial age-based community rating leads to this pattern in private health insurance exchanges in the U.S. The distortions are large enough that older individuals often face and sometimes choose dominated options. We discuss the potential effects of alternative menu-based approaches to community rating that would reduce relative price distortions.

[1] Contact information: Honglin Li (hli434@wisc.edu), Chenyuan Liu (liuchy3@sem.tsinghua.edu.cn), Justin Sydnor (justin.sydnor@wisc.edu). We thank workshop participants at the University of Wisconsin, Madison and the Munich Risk and Insurance Center at LMU for helpful comments.

Community rating is a common feature in many insurance markets. Motivated potentially 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 an insurance contract for those with higher expected risk and raises the price for those with lower expected risk.

How does community rating affect insurance markets with multiple differentiated products? Prior literature on this topic has focused on the potential for inefficiencies arising from adverse selection (see e.g., Einav, Finkelstein and Cullen, 2010; Geruso and Layton, 2017). If individuals with different risk levels sort into different insurance products and the price of those products reflects the average cost of enrollees, this can lead to (partial) unraveling of the market and inefficient levels of insurance. Risk adjustment schemes and related regulations that help to decouple insurance prices from selection patterns are a primary policy remedy for handling this consequence of community rating.

We examine the distortions in the relative prices of product options that arise with common approaches to community rating even when effective risk-adjustment schemes are in place. We develop a simple conceptual framework to analyze the prices of insurance policies with multiplicative community rating factors. The framework allows for both full community rating, in which each insurance product has the same price for all risk types, and partial community rating, in which regulation allows for a prescribed amount of variation in prices across risk groups. Community rating naturally leads to prices for insurance plans that do not reflect individual risk. As a consequence, the relative prices of different plan options will also typically not align with individual-level differences in the expected value of insurance options.

While this basic point that community rating leads to distortions in the relative prices of insurance options is already well known, our work highlights a novel insight: distortions in price *differentials* between product options can run in the opposite direction as the distortions in the overall *level* of prices. Community rating will tend to make insurance cheaper for those with high risk. One would expect then that in markets with differentiated products that community rating would especially lower the price of the highest value insurance products. This is indeed the typical logic in studies of adverse selection in vertically differentiated insurance contracts, in which community rating (initially) leads the marginal price for higher coverage options to be below the marginal cost of providing that additional insurance to the higher-risk types (Cutler & Reber, 1998; Einav, Finkelstein, and Cullen, 2010; Handel, 2013). In these cases, community rating lowers the overall level of prices for the higher risk types and also lowers their relative cost for higher-value insurance options. We show, though, that this will not always be the case. Even though higher risk types are receiving an overall transfer from community rating, the price differentials they face for more valuable insurance options can actually be higher than what they would face under risk-based pricing. This situation

arises with non-linear insurance contracts when the expected value of different insurance options varies proportionally more for lower-risk types than higher risk types. While it can occur at any level of community rating, it will be more likely with partial than full community rating.

We demonstrate the empirical relevance of this insight by analyzing the prices of private health insurance plans in the United States. In the aftermath of the Affordable Care Act of 2010, private health insurance markets in the U.S. involve a combination of strong pricing regulations and multiple product options, which makes them similar to health insurance markets in some other countries like Switzerland and the Netherlands. The Affordable Care Act established regulations that require community rating for plans but allows for prices to vary by geography, age, and, in some states, by smoking status. We focus attention on the pricing by age. Federal regulations allow for prices to rise with age using regulated age-rating factors, capped at a 3:1 age-rating ratio for 64 year olds relative to 21 year olds. At the time 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 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 covering approximately 60 percent of expected medical expenses for an average population established by the Center for Medicare and Medicaid Services, while Silver, Gold, and Platinum plans target 70, 80, and 90 percent actuarial values 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. We combine these data with national-level medical-utilization data through the NBER's license with Truven Marketscan that allows us to estimate the expected value of plan options at different ages.

We show that while under the community-rating rules older enrollees face price levels that are below their expected cost, 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.[2] In reality the average observed annual price differential for 64-year

[2] 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.

olds for gold vs. bronze plans was $4,300. In contrast, the additional price for gold vs. bronze-tier plans for 30-year olds 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 we estimate 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 the majority of enrollees 55 years and older in California had 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. Perhaps as a consequence of 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, the majority of them selected a plan that was dominated by a lower-tier option offered by the same insurer.

The reason for these surprising patterns lies in an interaction between the non-linearity of health insurance contracts and the community-rating rules that are applied across all plans regardless of the plan design. While regulations require that plans in each tier are designed to achieve a target level of actuarial value for a pre-determined average population, the actual actuarial value at different ages differs. For younger individuals, the actuarial value differences between the more generous gold-tier design and the bronze-tier designs are larger than the 20 percent actuarial-value difference targeted by regulation. For example, we estimate that for 25-year olds the gold tier plan covers about 80% of expected medical expenses, but the high-deductible-health-plan bronze-tier plan covers only 54% of expected medical expenses. 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 out-of-pocket maximum limits with both plan designs and we estimate much more similar actuarial value 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 and compressed for younger individuals relative to the expected cost differentials of the different tiers for each age group.

While Covered California offers an attractive market for our primary analysis due to availability of pricing and enrollment data at the individual level and standardized plan options, the basic findings are

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 any of our analysis, it does make comparisons of cross-tier prices a little more complicated to display.

likely to hold in many other private health insurance exchanges. We provide evidence that the patterns in CA on dominated options appear to hold broadly outside of California as well.

The distortions in the relative prices across plan options in multi-product markets occur because community rating is typically implemented at the level of insurance plans.[3] We examine how different prices might be under an alternative "menu-based" approach where community rating in multi-product markets aimed to create an overall transfer in the level of prices but allowed the differentials in prices to reflect expected cost for different risk types.[4] We illustrate the impact of such a change in the California system. We begin by calculating the implied average transfer under community rating by age, estimating that on average that the oldest enrollees receive a transfer of more than $1,000 per year, with the youngest enrollees providing a transfer of about the same magnitude. We then recalculate prices holding fixed current enrollment patterns and this level of average transfer but requiring that the marginal price between coverage tiers reflect the actuarial-fair differential for each age. We estimate that in such a system, the oldest enrollees would face prices that are about $400 less per year for gold-tier plans and see an increase of around $100 per year in bronze-tier plans. Younger enrollees would see a moderate increase in the price of gold-tier plans of around $100 to $200 per year, and a modest decrease in bronze-tier prices on the order of around $50 per year.

Our work contributes to a few streams of prior literature. First, these results contribute to a literature that has raised questions about the value of providing multiple insurance options, particularly vertically differentiated options of different coverage levels, with community rating. Many studies have noted the potential issues that can arise with adverse selection distortions across plan options under community rating (e.g., Cutler & Reber, 1998; Einav, Finkelstein, and Cullen, 2010; Handel, 2013, Geruso and Layton, 2017). Geruso et al., (2023) highlight that policies needed to combat intensive-margin selection issues, such as risk adjustment, may conflict with policies to support extensive margin selection. Our results suggest that common community-rating approaches, and the partial age-rating policies in ACA private health exchanges in particular, have this tension. Our analysis of the menu-based alternative to community rating shows that this alternative would bring down the cost of the lowest-coverage plan options for younger enrollees, going in the direction of better supporting extensive-margin selection into the market. In other related work, Marone and Sabety (2022) highlight that even in markets with strong pricing regulations, the benefits of vertical choice in insurance markets can be limited if prices cannot reflect individuals' costs (see also Ericson and Sydnor, 2017). In their framework, private information is the primary reason that price

[3] This is true generally with community rating and not limited to the partial community rating in ACA markets.
[4] This could be achieved, in theory, via regulations that mandate differential rating factors by plan design or by allowing for risk-based pricing of all plans but a series of taxes and subsidies that vary by risk-level.

differentials could not equate to costs. Our analysis focuses specifically on the implications for community rating as it is typically implemented on price differentials in a way that prior literature has not. We highlight that even without private information, community rating can result in non-intuitive distortions in relative pricing across plan options.

Second our analysis also identifies a new way in which pricing distortions, and dominated options in particular, can arise in insurance menus. Prior research on dominated options has primarily focused on how they 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). Rasmussen and Anderson (2021) provide evidence of the emergence of dominant options on private health insurance exchanges as a consequence of changes in pricing practices in response to the federal government ending payments for required cost-sharing subsidies. Our work shows that dominant options can arise for older individuals in the ACA as a natural consequence of otherwise-well-intentioned core pricing regulations. These findings emphasize the need for regulators to understand the potentially nuanced impact of pricing regulations, especially in environments with non-linear insurance contracts.

Finally, our results provide additional evidence of the potential value of finding ways of providing better information for both consumers and policymakers that makes it easier to 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 & Sydnor, 2020; Abaluck and Gruber, 2020; Gruber, et al., 2021; Heiss, McFadden, Winter, Wuppermann, and Zhou, 2021; Handel, Kolstad, Minten, and Spinnewijn, forthcoming). 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 people who fall back on simple heuristics when making choices, since older individuals might naturally assume that their higher risk means that the most generous coverage options are right for them and vice versa for younger individuals. Our work highlights the importance of both designing market structures and providing decision support in a way that does not financially disadvantage less financially 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 focus on a simple setting where there are two types of consumers in the market, denoted H and L. Let the share of L types in the market be denoted γ. The types differ in their level of total expected loss prior to insurance. Let θ^H and θ^L denote the expected losses. Without loss of generality, we assume that $\theta^H > \theta^L$. In the context of health insurance, we can think of H as denoting older and less healthy individuals and θ^H and θ^L as denoting the expected total medical expenditures of unhealthy and healthy types respectively.

We consider the simplest case for multiproduct markets, where there are two plans, j, and k, which are exogenously determined. Let m denote the plan index with $m = j, k$. The expected cost of providing plan m to type i is λ_m^i. We assume that $\lambda_j^i > \lambda_k^i, \forall i$, which implies that plan j has the higher expected cost for all types. In the context of health insurance with plans vertically differentiated by coverage level, this assumption would imply that the expected covered expenditures are higher under plan j for all types, or equivalently that plan j has higher "generosity" for all types.

The actuarially fair premium of plan m for type i would charge the type-specific expected value so that:

$$p_{im}^{fair} = \lambda_m^i. \tag{1}$$

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. We denote the allowed multiplicative risk-rating factor as a and define it as the allowed ratio of premium charged to H versus type L. Full community rating, in which both types pay the same price, occurs when $a = 1$. Cases in which $1 < a < \frac{\theta_H}{\theta_L}$ would typically be called "partial" community rating, reflecting the fact that some pricing variation is allowed but not as much as the ratio of full expected losses. We assume that a mixture of competition and regulations result in final premiums that are set as follows:

$$p_{Hm}^{cr} = \frac{a}{\gamma + (1-\gamma)a}(\gamma\lambda_m^L + (1-\gamma)\lambda_m^H), \tag{2}$$

$$p_{Lm}^{cr} = \frac{1}{\gamma + (1-\gamma)a}(\gamma\lambda_m^L + (1-\gamma)\lambda_m^H)\,. \tag{3}$$

The term in parentheses $(\gamma\lambda_m^L + (1-\gamma)\lambda_m^H)$ is simply the population-level expected cost if the full population were enrolled in plan m. The denominator in the ratios $\frac{a}{\gamma+(1-\gamma)a}$ and $\frac{1}{\gamma+(1-\gamma)a}$ in equations 2 and 3, respectively, is a normalization for a break-even condition that ensures that if the entire population enrolls in plan m that the collected premium will equal the expected cost. While other normalizations are possible and one could instead imagine break-even conditions based on observed (rather than population-level) enrollment, this normalization comes close to matching the regulatory environment under the Affordable Care Act described in Section 3.1.

2.2 Transfers and Relative Price Distortion under Community Rating

We begin by considering the level of transfers inherent in the premiums for each plan. We define the transfer received by type i under plan m as the difference between the actuarially-fair and community-rated premiums for that plan for that type. Using the community-rating rule for premiums defined in Equations 2 and 3 above, this gives:

$$T_{Hm} \equiv p_{Hm}^{fair} - p_{Hm}^{cr} = \frac{\gamma}{\gamma + (1-\gamma)a}(\lambda_m^H - a\lambda_m^L), \tag{4}$$

$$T_{Lm} \equiv p_{Lm}^{fair} - p_{Lm}^{cr} = \frac{(1-\gamma)}{\gamma + (1-\gamma)a}(a\lambda_m^L - \lambda_m^H) = -\frac{(1-\gamma)}{\gamma}T_{Hm}. \tag{5}$$

The high-cost type H receives a positive transfer in the pricing of plan m if and only if $a < \frac{\lambda_m^H}{\lambda_m^L}$. That is, there is a positive transfer to H if the allowed pricing ratio between the types is less than the ratio of the fair premiums (i.e., expected costs) for that plan. The transfer to/from the L types is the opposite of the transfer for the H types scaled by the relative population weight of H types over L types.

Now consider the price *differentials* between the two plan options that each type observes under community rating. Recalling the assumption that plan j has the higher expected cost for all types, we can define the actuarially fair relative premiums (a positive value) for each type as:

$$\Delta_i^{fair} \equiv p_{ij}^{fair} - p_{ik}^{fair} = \lambda_j^i - \lambda_k^i. \tag{6}$$

Then we can calculate an expression for the difference in premiums between plans under community rating for each type, which gives:

$$\Delta_H^{cr} = p_{Hj}^{cr} - p_{Hk}^{cr} = \frac{a}{\gamma + (1-\gamma)a}\left(\gamma\Delta_L^{fair} + (1-\gamma)\Delta_H^{fair}\right), \tag{7}$$

$$\Delta_L^{cr} = p_{Lj}^{cr} - p_{Lk}^{cr} = \frac{1}{\gamma + (1-\gamma)a}\left(\gamma\Delta_L^{fair} + (1-\gamma)\Delta_H^{fair}\right). \quad (8)$$

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 with the community-rated premium differential for the high-risk types be less than, equal to or greater than the fair premium differential for the high risk type? To see this we note that the question of interest is:

$$\Delta_H^{cr} = \frac{a}{\gamma + (1-\gamma)a}\left(\gamma\Delta_L^{fair} + (1-\gamma)\Delta_H^{fair}\right) \lessgtr \Delta_H^{fair}, \quad (9)$$

and this in turn can be easily shown to be equivalent to the following condition:

$$a \lessgtr \frac{\Delta_H^{fair}}{\Delta_L^{fair}}. \quad (10)$$

So the direction of the distortion in the price differential between plans under community rating for the high-risk type depends on how the allowed community rating factor compares to the ratio of the fair premium differentials between the high-risk and low-risk types. If the rating factor is less than the ratio of fair premium differentials, then the community-rated premium differentials will be compressed for the high type and they will be essentially subsidized on the margin for more generous coverage. However, the reverse will be true if the rating factor exceeds the ratio of fair premium differentials. It is easy to show that for the low-risk types, the relevant condition for the price differentials simply flips, so that the direction of the distortion (if any) runs in the opposite direction for the two types.

In many situations, it will be true that $a < \frac{\Delta_H^{fair}}{\Delta_L^{fair}}$ and the intuitive case in which community rating leads to a reduction in the relative price for higher-valued plans for the high-risk type will hold. For example, it will hold for any community-rating factor $a \in \left[1, \frac{\theta_H}{\theta_L}\right]$ for linear plan designs that cover a constant share σ_m of the loss. For linear plan designs, $p_{im}^{fair} = \sigma_m\theta^i$ and $\Delta_i^{fair} = (\sigma_j - \sigma_k)\theta^i$. In these cases, the ratio of fair premium differences is simply the ratio of the expected losses since each plan's cost is just proportional to expected losses.

In other cases, though, the ratio of the fair premium differentials may be smaller and the condition can reverse. For example, health insurance plans often have non-linear designs, incorporating combinations of deductibles, co-insurance, and limits on out of pocket costs. In these situations, the ratio of the difference in fair premiums across types can be smaller than the ratios of the level of the fair premiums. This occurs if the lower-risk types have proportionally larger differences in coverage between the plan options.

2.3 Numerical example

For a simple numerical example to illustrate this possibility, consider a market in which 75% of individuals are low-risk types (L) and the other 25% are high-risk types (H). Further, suppose that there are two possible underlying loss sizes, $2,500 and $50,000. The types differ in the probability of experiencing each loss size. Assume the low-risk types (L) having an 90% chance of the small loss and 10% chance of the large loss, while the high-risk types (H) have an 10% chance of the small loss and an 90% chance of the large loss. While clearly stylized, this can be thought of as representing a health insurance market that has more young than old and where younger individuals are more likely to incur modest health expenditures, while the older have much higher expected costs.

Now suppose that there are two plan options with non-linear cost-sharing designs in the market. The following table shows the plan designs:

Table 1. Numerical Example Insurance Plan Designs

Plan design	Deductible	Coinsurance (paid by insured)	Out of pocket limit
Generous	$0	10%	$2,500
Less generous	$1,500	10%	$2,500

The generous plan here has no deductible, covers 90% of costs up to an out-of-pocket maximum of $2,500. The less generous plan is very similar but has a $1,500 deductible before coverage kicks in. For small losses ($2,500), the generous plan will cover $2,250 (90% of the cost), while the less generous plan will only cover $900 (90%*(2,500-1,500)). On the other hand, for the large $50,000 losses, the out-of-pocket limit would be hit with both the generous and less-generous plans, so in each case the plan would cover the same $47,500 of losses.

Table 2 analyzes the fair and potential community-rated premiums for this situation. Columns 1 and 2 show the fair premiums for each type in each plan. The expected cost for the high-risk types is much greater for both plan designs. The fair premium ratio between *H* and *L* types is 6.34 for the generous plan and 7.71 for the less generous plan. However, the difference in premiums in this example are actually higher for the low-risk types. The reason is that they are much more likely to suffer the small-size losses for which the two plans differ in coverage.

Table 2. Numerical Example Premiums

	Fair premiums			Full CR premiums (a = 1)		Partial CR premiums (a = 2)	
	High type (1)	Low type (2)	Ratio (H/L) (3)	High type (4)	Low type (5)	High type (6)	Low type (7)
Generous plan (j)	$42,975	$6,775	6.34	$15,825	$15,825	$25,320	$12,660
Less generous plan (k)	$42,840	$5,560	7.71	$14,880	$14,880	$23,808	$11,904
Difference (j - k)	$135	$1,215	0.11	$945	$945	$1,512	$756

This then becomes a case where community rating will naturally lead to an increase in the price differential for the more generous plan for the high-risk types. We can see this in the subsequent columns, which show community rated premiums under full community rating ($a = 1$) and an example of partial community rating ($a = 2$). In the case of full community rating, we see that all types pay the same premium for each plan (reflecting the expected cost for the population), and that this represents a large transfer from the low-risk to high-risk types for both plans. However, the bottom row shows that with full community rating, the price differential for the high-risk types is $945, more than three times the fair price differential for H types. Low-risk types, in contrast, see a reduction in the price differential. For the partial community rating case, we note that the risk-rating factor of 2 is below the ratio of fair premiums for both plans, and as a result the high-risk types again get a transfer and see lower premium levels than with risk-based pricing. But the price differentials are substantially inflated for the high-risk types. In fact, at a price differential of $1,512, the less generous plan now *dominates* the generous plan for the high-risk types.

This example is particularly stark because it is a case where the fair price differentials are actually smaller for the high-risk types so that $\frac{\Delta_H^{fair}}{\Delta_L^{fair}} < 1$. We note, though, that this needs not be the case to get the distortion in price differentials to run in the opposite direction as the distortion to price levels. For that to happen we only need to observe $\frac{\Delta_H^{fair}}{\Delta_L^{fair}} < a$. This, in turn, helps to highlight that for the price distortion to run in this direction will be more common under partial community rating ($a > 1$) than for full community rating ($a = 1$), but can happen in either case. In the following sections we provide evidence that in the California private health insurance exchanges, the fair price differentials between plan options are larger for the older (higher-risk) types, but the key ratio of interest is below the allowed community rating factor.

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 in the Covered California health insurance market context and supplement the analysis with the Federal Affordable Care Act exchange data. 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 the California health insurance marketplace offering ACA-compliant plans to individuals and small businesses. The market features vertical choices, where plans are ranked by actuarial value in four metal tiers: Bronze (60%), Silver (70%), Gold (80%), and Platinum (90%).[5] Unlike the ACA markets in many other states where insurers offer a wide range of plan designs (Liu, 2023), 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 Gold and Silver tiers, the cost-sharing designs are standardized, except for a few benefits for which there is a copay version and coinsurance version. 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 regulates insurers to set premiums based on all enrollees choosing their plan, not the pool selecting into a specific plan. The single risk pool requirement is consistent with the premium setting equations in 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 intend to eliminate insurers' incentive to risk select healthy 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, which specifies the premium ratio for each age group relative to the 21-year-old. Insurers set a base premium for the 21-year-old for a plan, and premiums 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

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

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, representing more than 400,000 enrollees annually. Moreover, the market highlights many features common in 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, maximum out-of-pocket, coinsurance rates, etc.) for plans in the market from the official Covered California website.[6] 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., actuarial value, 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.[7] The datasets contain information on other plan attributes, including carrier names, plans available region, premiums, and plan types (HMO, PPO, or EPO.)

Covered California enrollment data. The 2015 – 2020 Covered California enrollment data were acquired via a public records request following previous literature (e.g. Saltzman, 2019; Tebaldi, 2022). 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 chooses by matching the plan information—less than 1% of enrollees 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.

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

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' actuarial values. 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 4.1 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 of Medicare and Medicaid Services.[8] We collect each plan's attributes from these files, including the cost-sharing design, insurer, plan type, network ID, formulary ID, launched county, and premiums.

3.3 Adjustments to Medical Expenditure and Plan Data

Our goal in the analysis in the next section is to compare, for each age and each 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. We extend the analysis in Section 2.1 into multiple types and have

$$p_{im}^{fair} = \lambda_m^i. \tag{11}$$

$$p_{ij}^{cr} = a_i \frac{\sum_{s=21}^{64} \gamma_s \lambda_j^s}{\sum_{s=21}^{64} \gamma_s a_s} + \tau, \forall i, m, \tag{12}$$

where γ_i is the population weights of age i, a_i is the age-rating factor, and τ is a level shift adjustment to all premiums to make sure the system breaks even.[9] The key component of the calculation is λ_m^i, which we

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

[9] To calculate τ, we first set $\tau = 0$, solve p_{ij}^{cr} and p_{ij}^{fair}. We then calculate the system-level budget using the observed 2017 California age by metal tier enrollment share, ω_{ij}. $\tau = \sum_i \sum_j \omega_{ij}(p_{ij}^{fair} - p_{ij}^{cr})$.

estimate using the Truven MarketScan data and the Covered California Plan design. The population weights are estimated using the California enrollment data.

To obtain the estimated distribution of medical costs, we estimate each age's total medical expenditure distribution using the 2013 Truven MarketScan data. We 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 second term is a common additive adjustment term to all plans that guarantees that the system of premiums would be expected to break even accounting for the observed enrollment shares by age and metal tier in the California health insurance exchange. 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 A. The key assumption here is that while the level of costs may differ in California, that the distribution of costs by age is proportional in California to what is observed in the national data.[10]

The second step in our calculations is to create simplified plan designs to represent the cost-sharing of plans in California. We consider individual coverage for the in-network services of Gold, standard Silver, and Bronze plans. In reality, the cost-sharing rules of a plan might be complicated: different services may have different coinsurance rates or copayments applied. Following Liu and Sydnor (2021), 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 leg is a coinsurance range, and the third leg is the MOOP range, under which consumers out-of-pocket spending is capped. We convert using the CMS actuarial value calculator for the ACA market. The calculator estimates the actuarial value 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

[10] One small 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 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 both to the national estimates available at that time and to our estimates. At the very least it suggests that the sort of age differentials we use in this paper are consistent with the type of analysis that Covered California was getting from professional consultants.

between the deductible and MOOP such that the plan's actuarial value is the same as the reported actuarial value.

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 drugs for three age groups conditional on the total expenditure level. Appendix Figure 1 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. The Bronze with a health savings account (HSA) has a 40% coinsurance rate for all services in the second leg, so the design is already a simplified three-arm design, and 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 similar λ_m^i. Indeed, Appendix Figure 2 shows that the two λ_m^i lines are extremely close to each other.

Table 3 shows the simplified plan design for the year 2017. For all the following analysis, 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. Using the simplified plan designs and applying the adjusted Truven distributions described above, we calculate the actuarially fair and community-rated premiums using Equations (4) and (5).

Table 3. California 2017 Simplified Plan Designs

Plan	Deductible	MOOP	Simplified Co-insurance rate (Paid by consumers)
Gold	0	6750	0.28
Standard Silver	2750	6800	0.09
Bronze	6800*	6800	0.92
Bronze - HDHP	4800	6550	0.40

Note: This table shows the deductible and maximum out of pocket limits 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 actuarial value 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.

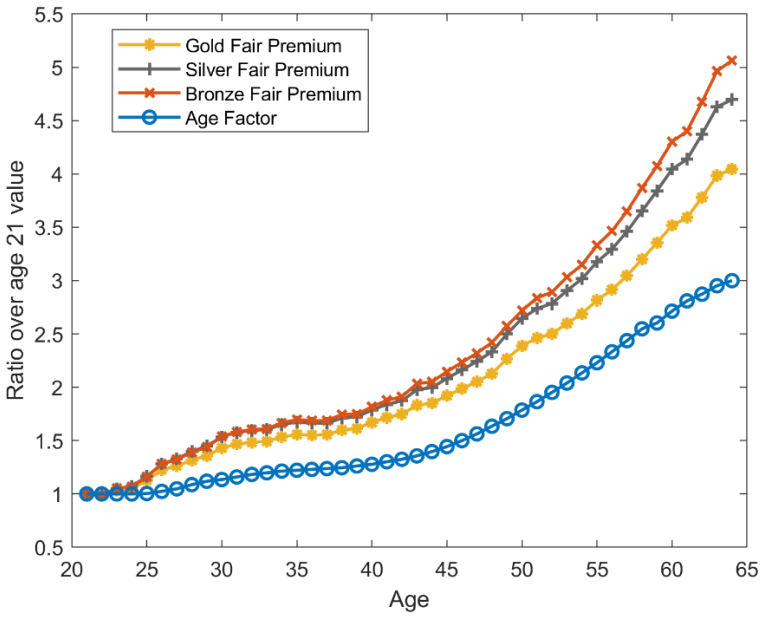
All designs are for individual, first-tier in-network coverage. *We use a slightly different procedure and have a slightly different interpretation of the Simplified Co-insurance rate for the Bronze plan. The Bronze plan has a deductible of 6800, the same as the MOOP, while some services are covered 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 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 maximum out of pocket limit of $6,800.

4. Empirical results

4.1 Price Distortion with CA Plans

We first illustrate how the current age-based community rules distort the plan premiums relative to risk-based pricing. Figure 1 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-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.

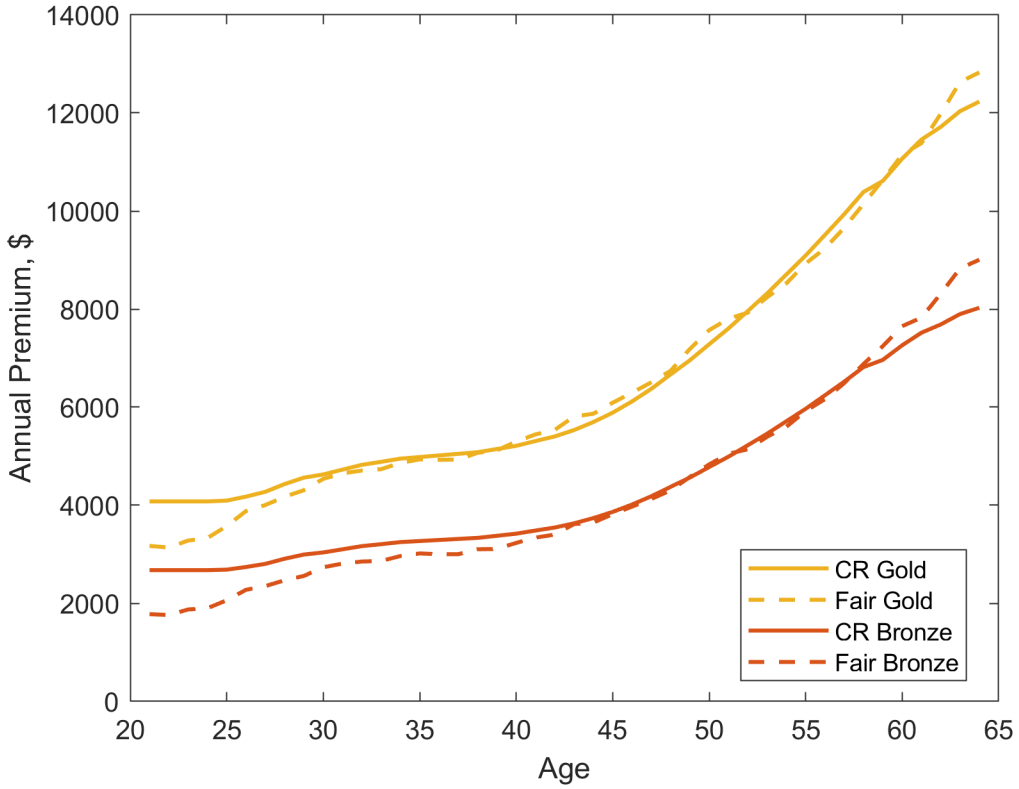
Figure 1. ACA Age Factors and Fair 2017 California Premiums



Note: We use the 2017 California plan design to calculate the figure. First, we construct the medical expenditure distributions for each age using the 2013 Truven MarketScan data. Next we inflate the distributions to match the average premium levels in California (see text and Appendix A for details). Third, we transform the California standardized plan design into a simplified three-arm design using the 2017 actuarial value calculator published by CMS. Fourth, we apply the simplified design on the medical expenditure distribution to simulate the fair premiums for each age and metal tier. 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 2 illustrates the distortions in premium levels for Gold and Bronze plans by age. For both plans, the community rating 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 1).

Figure 2. Premium Distortion under Community Rating: California 2017

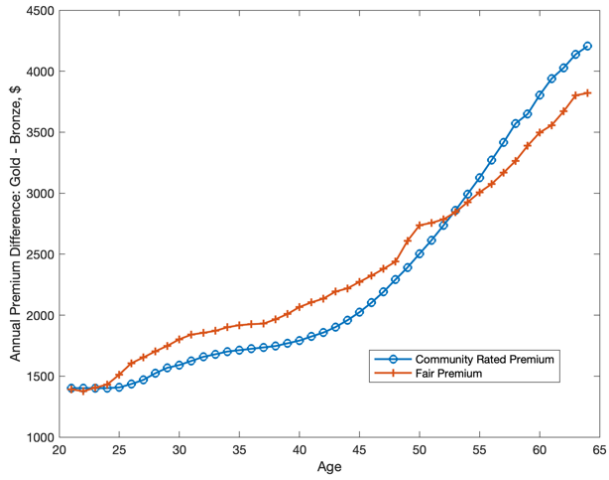


Note: We use the 2017 California plan design to calculate the figure. "CR" indicates community rated premiums, while "Fair" indicates premium level under risk-based pricing. The y-axis indicates the annual premiums for individual coverage.

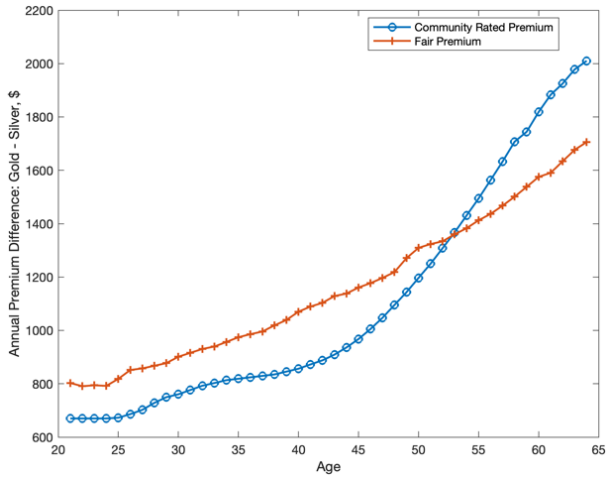
However, the relative premium distortion is different by age. Figure 3 Panel A shows the premium difference between Gold and Bronze plans for each age. The red line with crosses represents the risk-based actuarially fair premium differences, while the blue circle represents the community-rated premium differences. We find that the age-rating rules make the premium gap larger than the fair level for the older population and smaller for the younger population. We find a similar pattern between Gold and Silver plans (Panel B) and little relative premium distortion between Silver and Bronze plans.

Figure 3. Premium Differences Across Metal Tiers: California 2017

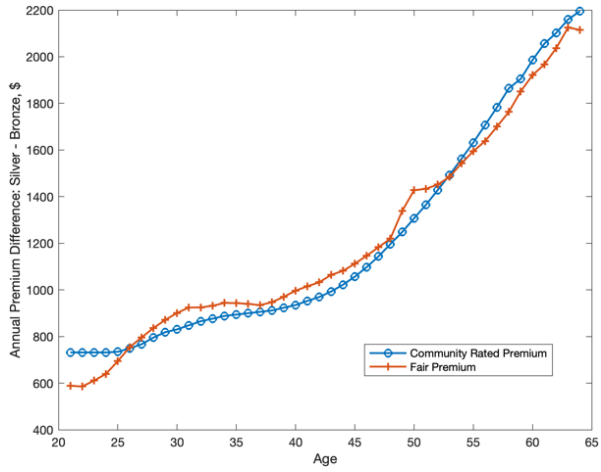
Panel A. Gold and Bronze



Panel B. Gold and Silver



Panel C. Silver and 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 summary, we find that the age-based community rating rules increase the premium gap between Gold and Bronze plans for the older population in the Covered California market.

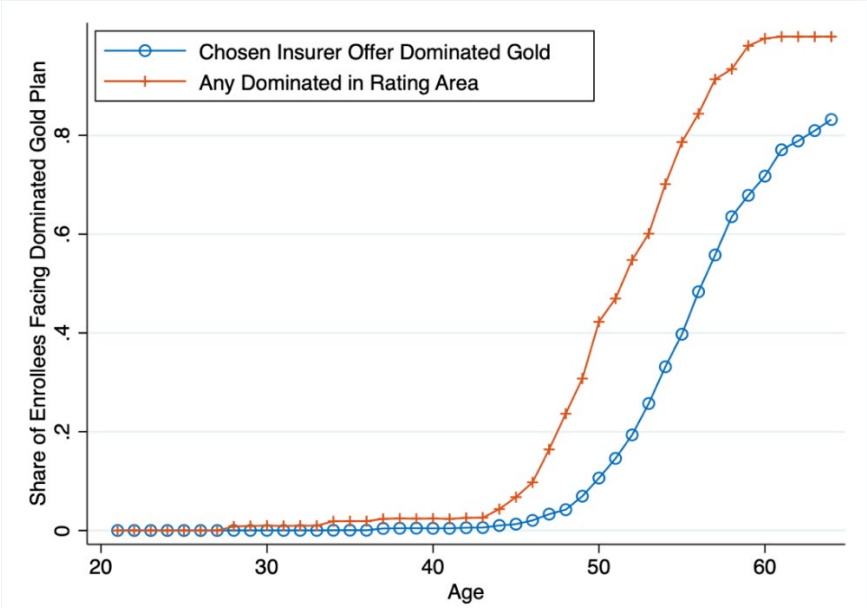
4.2 Dominated Gold Plans in Covered California

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. [11] 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 otherwise similar Gold-Bronze plan pairs. 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 might 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 extra 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 in our analysis.

Figure 4. Share Facing Dominated Gold Plan by Age: California 2015 – 2020



Note: Data from enrollment information through Covered California 2015-2020. Sample include individuals not eligible for CSR-variation plans. We construct Gold-Bronze pairs by requiring they are offered by the same insurer, in the same rating area, and with the same plan type (HMO, PPO, EPO, POS). 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.

Second, for each pair, we calculate the 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. The value gives the minimum premium gap such that Bronze

[11] 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.

would dominate Gold. 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.[12] Figure 4 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 shows that the premium gap distortion resulting from the community rating rule is large enough to make dominated Gold plans prevalent for 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.

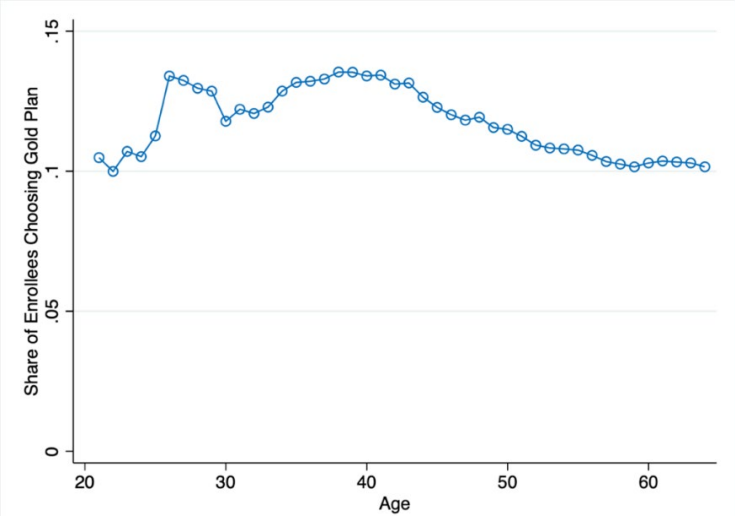
4.3 Gold Enrollment Share by Age in Covered California

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 5 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 a number of reasons for these age-related enrollment patterns, we note that they are at least consistent with the distortions in relative premiums we document above that make Gold-tier plans relatively expensive for the oldest enrollees.

Second, we find that the existence of the dominated Gold plans creates room for mistakes in choosing plans. Among enrollees not eligible for CSR, about 2% of them choose a dominated Gold plan. Figure 6 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 choose a dominated Gold plan, and more than 46% of enrollees aged above 54 who choose a Gold plan choose a dominated Gold plan.

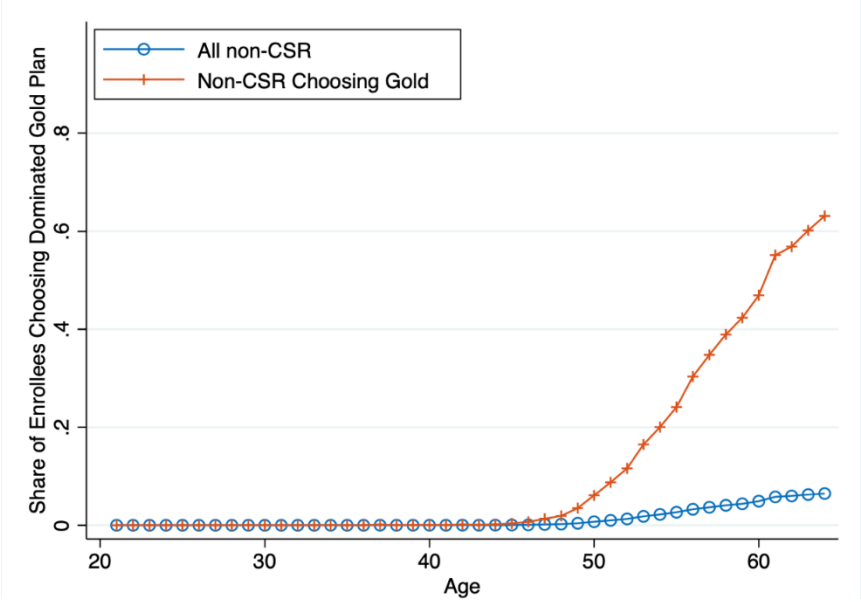
[12] We exclude the CSR-eligible population, because these individuals can choose from more generous plans with the Silver premium, making Gold plans possibly dominated by these CSR-variation Silver plans. Such dominance has been documented in the literature (DeLeire et al., 2017). However, it is not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

Figure 5. Share Choosing Gold Plan by Age: California 2015-2020



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

Figure 6. Share Choosing Dominated Gold Plan by Age: California 2015 – 2020

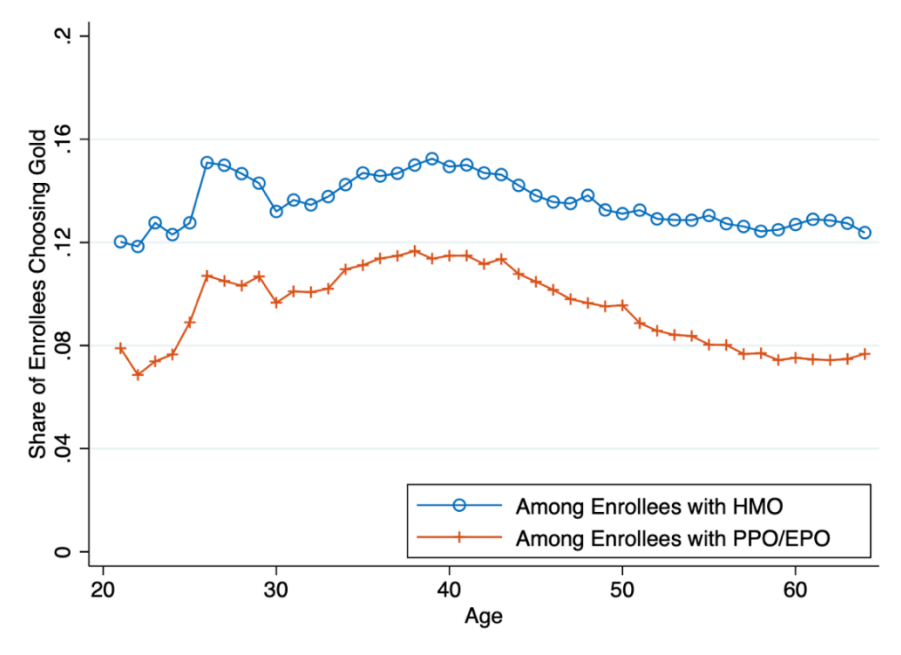


Note: Data from enrollment information through Covered California 2015-2020. Sample include individuals not eligible for CSR-variation plans. We calculate the share of enrollees choosing a dominated Gold plan among all non-CSR eligible enrollees and non-CSR eligible enrollees who select a Gold plan at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

Finally, we find that the Gold enrollment share decreases more among enrollees choosing a PPO plan. Conventionally, we would expect that older (and higher expenditure) enrollees would prefer more coverage (thus a Gold plan) and more broad 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.

Figure 7. Share Choosing Gold by Plan Type and Age: California 2015-2020



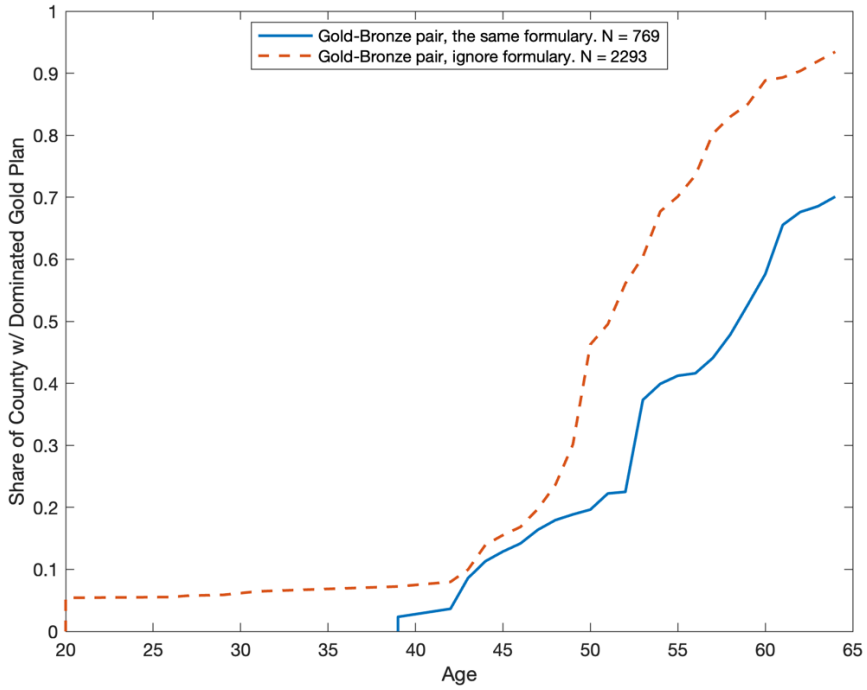
Note: Data from enrollment information through Covered California 2015-2020. Sample includes individuals not eligible for CSR-variation plans. We calculate the share of enrollees choosing Gold plan separately for two types of enrollees: those choosing an HMO; those choosing a PPO.

4.4 Dominated Gold Plans in the Federal Exchange

The above findings in Covered California are likely prevalent in many other markets. 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 of the two plans having 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 2.4 to calculate, for each county, the minimum age at which there is at least one dominated Gold plan.

Figure 8 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 other plan attributes. If we allow Gold and Bronze plans to have possibly different drug formulary, then about 50% of counties have at least one dominated Gold plan for enrollees older than age 50.

Figure 8. Share Facing Dominated Gold Plan by Age: Federal Exchange 2017



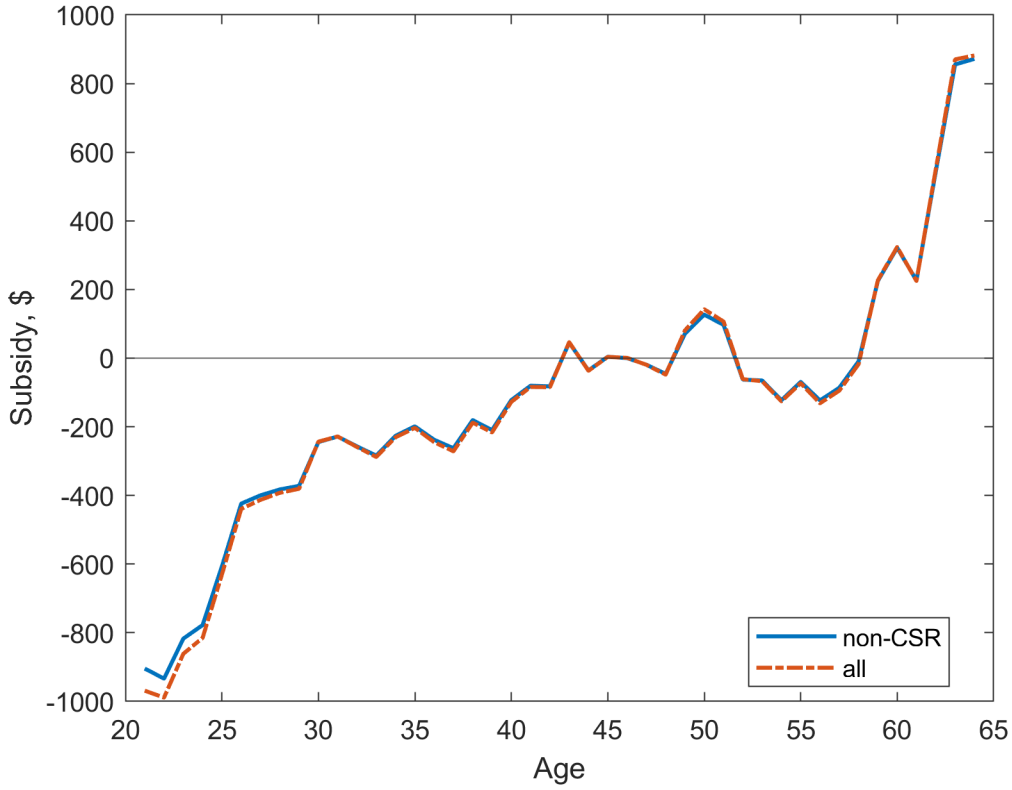
Note: All data from states using Healthcare.gov in 2017. Plan designs are for individual, first-tier in-network coverage. We transform all plan designs into simplified three-arm design using the 2017 AV Calculator. We construct two types of Gold-Bronze pairs: (1) with the same insurer, plan type (HMO, PPO, EPO or POS), network ID, national network status, launched in the same county; (2) on top of (1), also require both plans have the same drug formulary (indicated by the same drug formulary ID). After constructing the Gold-Bronze pairs, we then calculate the age at which Gold would be dominated by Bronze. Finally, we calculate, in each county, whether there exists at least one Gold plan dominated by their respective Bronze pair for a particular age. The solid blue line indicates the pairs constructed using (1), while the dashed red line indicates the pairs constructed using (2). We exclude counties without any comparable pairs in each case, and the number of counties used in the calculation is labeled in the legend.

4.5 Counterfactuals with Menu-based Community Rating

The preceding sections have documented the ways in which community rating regulations lead to distortions in the price differentials between metal tiers in the Covered California market. These distortions are ultimately a consequence of the fact that community rating is implemented across all plans in the same way. It is conceptually possible, however, to design a 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. This 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).

Figure 9. Implied Subsidy by Age: California 2017

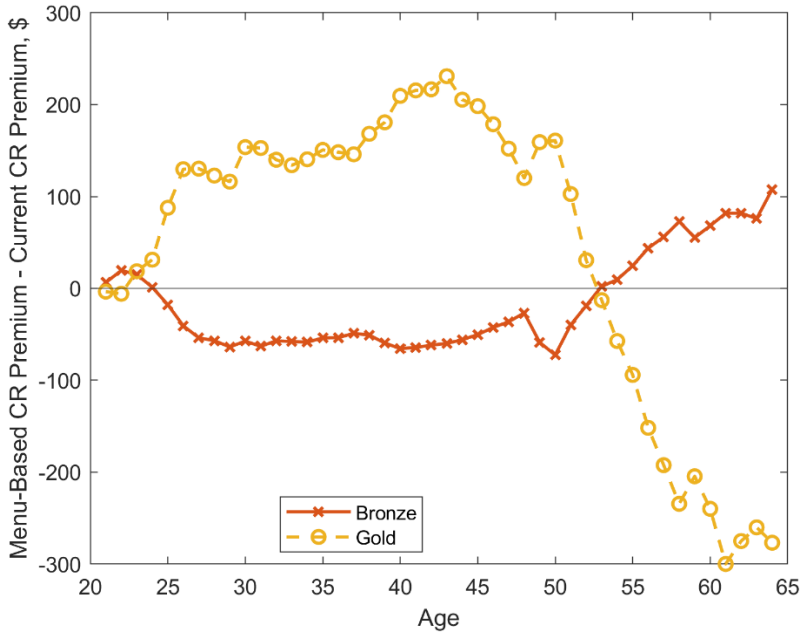


Note: We use the 2017 California plan design to calculate the figure. We first calculate the subsidy level for each age and metal tier by subtracting the annual fair premium from the community rated premiums. We then calculate the average subsidy level for a particular age by weighting the subsidy for 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 either the non-CSR eligible enrollees (the blue solid line) or all enrollees (the red dashed line).

In this section we generate simple counterfactuals to show how different plan premiums would be under this type of alternative system. To do this, we first calculate the amount of the average transfer received or given at each age. We calculate these values by taking the difference between the observed premium for each individual's selected plan tier and the estimated fair premium level they would face under risk-based pricing. We average these values for each age. Figure 9 shows the results. Those above 55 years old receive positive subsidies, maxing out at transfers almost $1,000 for the oldest enrollees. Those under

30, in contrast, are effectively paying average surcharges, peaking at around $1,000 for the youngest enrollees.

Figure 10. Premium Change Under Menu-Based Community Rating



Note: We use the 2017 California plan design to calculate the figure. The y-axis indicates the annual premium difference (for individual coverage) between two pricing rules: the current community rating premiums and the fair premium with age-specific side transfers. Positive number indicates the counterfactual premiums are higher than the current premiums.

We can then calculate how the net premiums for each plan tier would change if premiums were risk-rated and people faced the calculated average surcharge/subsidy for their age. Figure 10 graphs these values by age for both the gold-tier and bronze-tier plans. The biggest change in premiums would occur for gold-tier plans for older individuals. Gold-tier prices would fall by as much as $300 per year for the oldest enrollees. The biggest increases in gold-tier premiums would occur for those around age 45, who would see increases of around $200 per year. In contrast, bronze-tier prices would rise for those above age 55 and fall modestly for those from age 25 to 55.

5. 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 high-risk vs. low-risk actuarially fair premium differentials between plan options.

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.

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 at times 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 have impacts on the efficiency of the system 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.

The counterfactual we introduce in Section 4.5 highlights that it may be feasible to consider different approaches to community rating in multiproduct markets. Our approach involves generating a set of community-rating transfers that affect the overall level of prices that different risk types face but allowing the differentials between plan options to reflect the type-specific expected marginal cost between plans. Removing relative pricing distortions in this way could potentially help restore the underlying logic of offering multiple products in selection markets. This type of system, though, would be useful in situations where plans can be effectively priced for expected cost based on observable risk type but would not address situations that fundamentally involve private information. In the context of health insurance, many aspects of risk are potentially observable and could be priced if community-rating regulations allowed the full use of observable data. Yet there may still be elements of private information, such as the extent of moral hazard response to additional cost sharing (Marone and Sabety, 2022), that are impossible to fully capture in pricing. A useful direction for future research may be to explore more fully the extent to which available data could predict individual-level variation in expected costs across insurance products and the degree to which people have private information relative to the best-available risk-based pricing approach in these settings.

References:

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, 2020. "Less Is More: Improving Choices by Limiting Choices in Health Insurance Markets." Working paper.

Bhargava, Saurabh, George Loewenstein, and Justin Sydnor, 2017. "Choose to Lose: Health Plan Choices from a Menu with Dominated Option." *The Quarterly Journal of Economics* 132(3): 1319–72.

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.

DeLeire, Thomas, Andre Chappel, Kenneth Finegold, and Emily Gee. 2017. "Do Individuals Respond to Cost-Sharing Subsidies in Their Selections of Marketplace Health Insurance Plans?" *Journal of Health Economics* 56 (December): 71–86.

Einav, Liran, Amy Finkelstein, and Mark R. Cullen, 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices." *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.

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." *The Review of Economics and Statistics*, 105(2): 237-257.

Handel, Benjamin R., 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review*, 103(7): 2643-82.

Handel, Benjamin, Jonathan Kolstad, Thomas Minten, and Johannes Spinnewijn. Forthcoming. "The Socio-Economic Distribution of Choice Quality: Evidence from Health Insurance in the Netherlands." *American Economic Review: Insights*.

Heiss, Florian, Daniel McFadden, and Joachim Winter. 2010. “Mind the Gap! Consumer Perceptions and choices of Medicare Part D Prescription Drug Plans.” *Research Findings in the Economics of Aging*.

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): 2737–81.

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, no. 537.

Liu, Chenyuan, 2023. “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 Anderson, David, 2021. “When All That Glitters Is Gold: Dominated Plan Choice on Covered California for the 2018 Plan Year.” *The Milbank Quarterly*, 99(4), 1059–1087.

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

Samek, Anya and Justin Sydnor, 2020. “Impact of Consequence Information on Insurance Choice.” NBER Working Paper 28003.

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

Tebaldi, Pietro. Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA. No. w29869. National Bureau of Economic Research, 2022.

Appendix A. Adjusting Truven Distribution to Match to California Premiums

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}.

We then mean-shifted the medical expenditure distributions to match the observed California premium levels using the following steps.

First, we calculate the 2017 California average premiums for Bronze, Silver, and Gold plans. Among all CA enrollees, we calculate the fraction of individuals in each carrier-rating area and year. We use this as weights to calculate the weighted average observed premiums of plans in 2017 separately for each metal tier for the 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_{ij}^{o} denote plan j's premium for age i, where j is Bronze, Silver and Gold.

Next, we assume that all ages have the same inflation factor k_j for plan j. Let x_i denote the uncovered expenditure level for age i. We multiple x_i by k_j, and then calculate the community rated premiums, p_{ij}^{s} in 2017 as a function of k_j. Specifically, let λ_{ij} denote the expected covered costs for age i under plan j. We calculate λ_{ij} as

$$\lambda_{ij}(k_j) = \sum_{s} oop_j(k_j x_{is})\mu_{is},$$

where $oop_j(\cdot)$ is the cost-sharing design of plan j. Given $\lambda_{ij}(k_j)$, we then calculate the community rated premiums as follows:

$$p_{ij}^{cr}(k_j) = a_i \frac{\sum_{s=21}^{64} \gamma_s \lambda_j^s}{\sum_{s=21}^{64} \gamma_s a_s}, \forall i, m,$$

where γ_i is the population weights of age i, and a_i is the age-rating factor. The resulting p_{ij}^{cr} may results in a non-neutral budget for the system. Let m_{ij} denote the fraction of individuals with age i choosing plan j (as a fraction of the total population.) Then we calculate the budget of the system under p_{ij}^{cr} as

$$b(k_j) = \sum_i \sum_j m_{ij}\lambda_{ij}(k_j) - \sum_i \sum_j m_{ij} p_{ij}^{cr}(k_j).$$

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

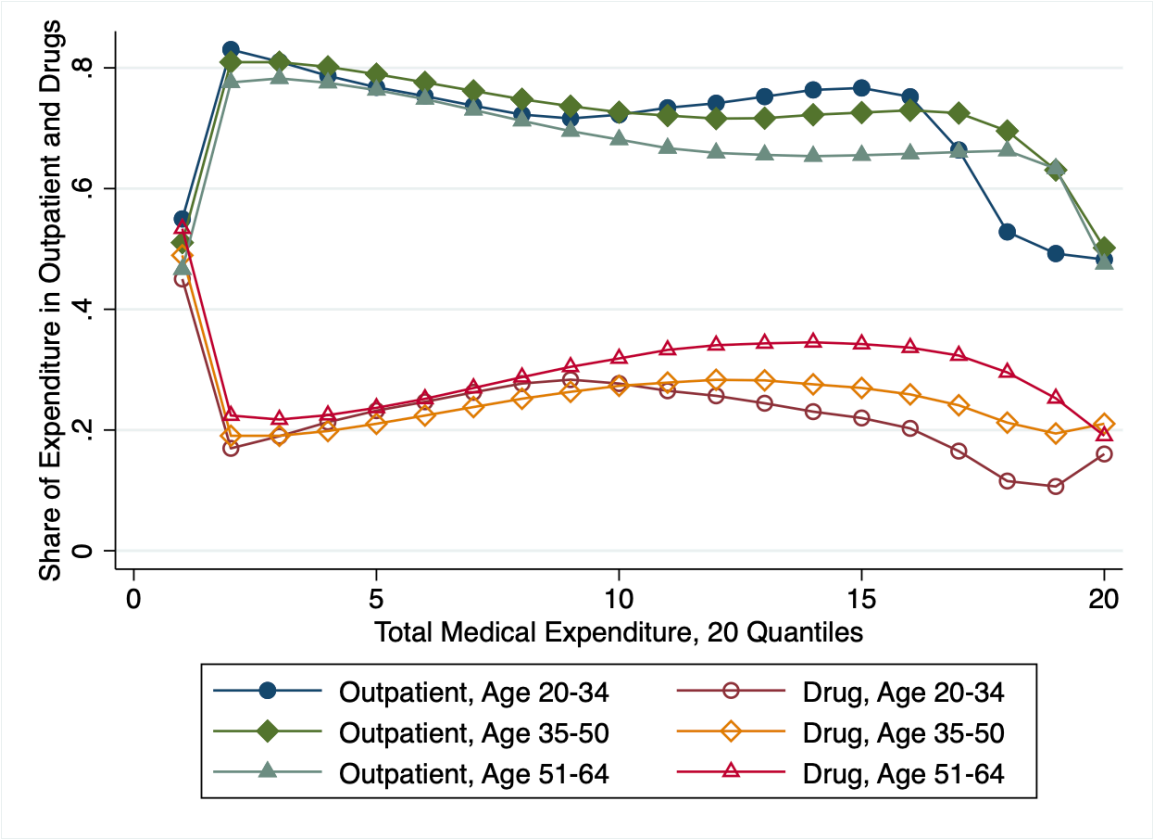
$$p_{ij}^{s}(k_j) = p_{ij}^{cr}(k_j) + b(k_j).$$

Finally, we solve for k_j by minimizing $\sum_i (p_{ij}^{s} - p_{ij}^{o})^2$.[13]

[13] In our calculation, even without the adjustment $b(k_j)$, the budget is almost neutral.

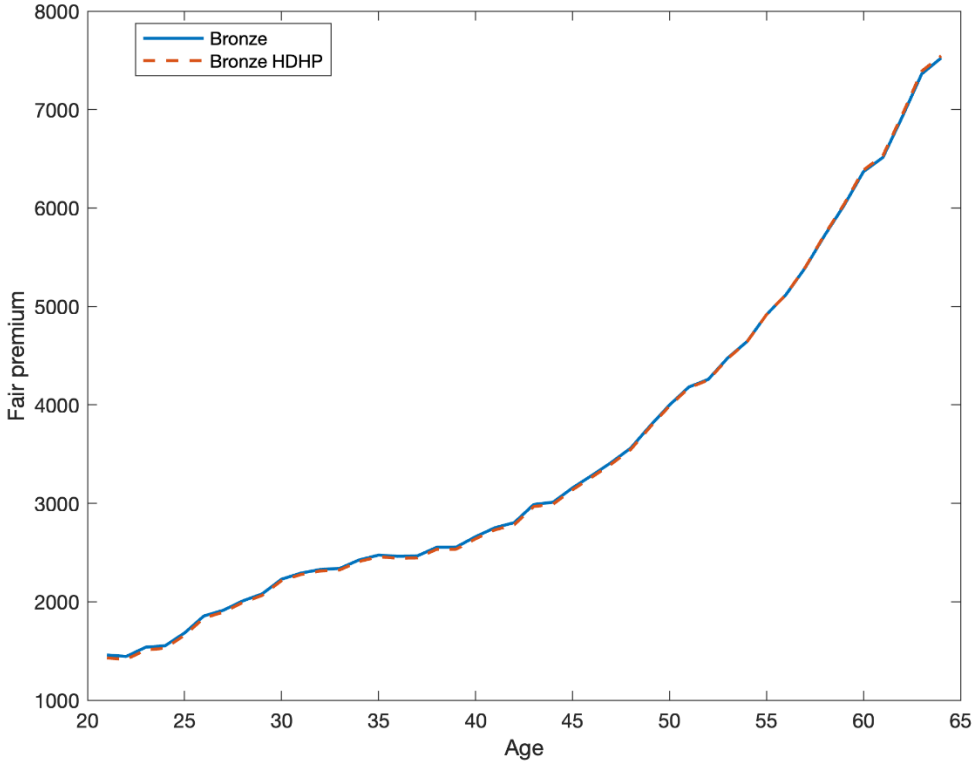
Appendix B. More Figures for Empirical Analysis.

Appendix Figure 1. Service Share and Total Medical Expenditure Level by Age



Note: 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.

Appendix Figure 2. Simplifying Plan Designs: Robustness Checks



Note: In California, and in some years, there might be two similar designs within a metal tier. For example, in 2017, there are two Bronze plans: one normal Bronze plan and another HDHP Bronze plan. The HDHP Bronze plan has a simplified three-arm cost-sharing rules by design. While the normal Bronze plan has more complicated designs and are 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.