Beyond Vertical Choice: Cost-Sharing Designs in Insurance Menus

Chenyuan Liu*

February 2026

Abstract

Prior work on sorting and the value of offering vertical choices typically assumes plans are ordered by financial generosity. I develop a precise ordering framework for cost-sharing designs under expected utility. I distinguish WTP-ordering (everyone prefers one contract when both are free) from the stronger RP-ordering (the preference survives at actuarially fair prices) and link each to transparent conditions on out-of-pocket schedules. Using the 2015 KFF Employer Health Benefits Survey, I document that most observed menus are not RP-ordered and that a substantial share are not WTP-ordered, with coverage better for smaller losses in one plan and better for larger losses in the other. I then use simulations disciplined by claims data and heterogeneous risk aversion to show that comparative statics and the value of offering choice differ sharply across menu types. Overall, the results caution against extrapolating welfare conclusions from vertically ordered menus to the broader set of designs observed in practice.

*School of Economics and Management, Tsinghua University. Email: liuchy3@sem.tsinghua.edu.cn.

1 Introduction

Health insurance markets routinely offer consumers a menu of plans that differ in cost-sharing attributes such as deductibles, coinsurance, and out-of-pocket maximums. A large literature evaluates whether offering multiple coverage levels improves welfare and how consumers sort across plans. Much of this work treats plan menus as vertically ordered along financial generosity, as in constant coinsurance designs, tiered metal levels, or other monotone coverage differences (Azevedo and Gottlieb, 2017; Ericson and Sydnor, 2017; Marone and Sabety, 2022; Chade et al., 2022; Ho and Lee, 2023).

However, prior work uses different labels for these coverage differences—actuarial value, coverage level, financial generosity, or vertical choice—and it is not always clear whether these terms refer to the same underlying ordering. Moreover, real-world plan menus need not be ordered along financial terms at all. With multiple loss states and multi-dimensional cost-sharing attributes, two plans can be horizontally differentiated along the cost-sharing dimension, so that one covers small losses better while the other covers large losses better. For example, Liu (2025) highlights how adverse selection forces would create a menu featuring such "non-ordered" plans in equilibrium.

In this paper, I develop a precise ordering framework for two-plan menus under the expected-utility (EU) model with risk-averse individuals facing stochastic losses. In much of the IO literature, vertical choice means that everyone prefers one contract to another when both are free; I formalize this as *willingness-to-pay-ordering* (WTP-ordering). A stronger notion is *risk premium-ordering* (RP-ordering), which requires that everyone also prefers one contract when both plans are priced at actuarially fair premiums. Each concept maps to transparent sufficient conditions on out-of-pocket (OOP) schedules: WTP-ordering holds when one plan's OOP is weakly lower at every loss state (with strict inequality somewhere), while RP-ordering additionally requires the lower OOP plan to be weakly flatter in losses (its slope is everywhere no larger).[1] These definitions yield three empirically relevant menu types: WTP- and RP-ordered menus, WTP-ordered but not RP-ordered menus, and menus that are not WTP-ordered.

Under this menu classification, the prior literature mostly studies the value of extra choices for menus that are both WTP- and RP-ordered (e.g., Azevedo and Gottlieb, 2017; Ericson and Sydnor, 2017; Marone and Sabety, 2022; Chade et al., 2022; Ho and Lee, 2023 is an exception). This raises two questions. First, are most real-world plan menus in this category, so that the other cases are empirically negligible? Second, if not, do the main welfare insights from the vertical-choice literature generalize to the other menu types?

I answer these questions in three steps. First, I document the prevalence of each menu type using the Kaiser Family Foundation Employer Health Benefits Survey (KFF EHBS) for 2015. I focus on firms offering exactly two plans that differ in cost-sharing attributes for single coverage.

[1] See Marone and Sabety (2022) for related results under CARA preferences and a moral-hazard formulation. I focus on more general utility functions and abstract from moral hazard.

I summarize each complex plan into a stylized three-arm design with a deductible segment, a coinsurance segment, and a maximum out-of-pocket (MOOP) segment (Liu and Sydnor, 2022). Specifically, I compute each plan's actuarial value (AV) using the calculator developed by the Centers for Medicare and Medicaid Services (CMS), then solve for the single coinsurance rate that reproduces the same AV while holding the plan's deductible and MOOP fixed. This produces a simplified OOP schedule that is a piecewise linear function of total annual spending, allowing a clean classification of menus by OOP ordering. I find that the majority of menus are not RP-ordered: only 14.6% are both WTP- and RP-ordered, 48.0% are WTP-ordered but not RP-ordered, and 37.4% are not WTP-ordered. Public health insurance markets also span different menu types: Swiss public health insurance offers contracts that satisfy both WTP and RP ordering, while Covered California provides clear examples of non-ordered menus.

Second, I use numeric simulations to study how menu type affects willingness-to-pay (WTP), marginal costs, and social surplus. I assume CARA utility with heterogeneous risk aversion and construct loss distributions from Truven MarketScan claims data using a shifted lognormal specification. For comparative statics, I generate a sequence of distributions that satisfy the monotone likelihood ratio property (MLRP) by varying the mean while holding other parameters fixed. I compute WTP for the difference between the two plans in each menu. Separately, I compute the marginal cost (MC) of the incremental coverage using a proportional loading factor that is held fixed at 1.2 throughout the simulations. I find that for RP-ordered menus, the standard intuition holds: WTP for additional coverage increases with risk type and with risk aversion, consistent with common wisdom. In contrast, for menus that are WTP-ordered but not RP-ordered, these monotonicities can fail: WTP can be non-monotone in risk type and risk aversion coefficient (e.g., hump-shaped or decreasing over some range). Finally, for non-WTP-ordered menus WTP is positive for some individuals and negative for other individuals.

Third, I evaluate the value of offering choice under full community rating. The key insight from Marone and Sabety (2022) is that heterogeneity in efficient coverage levels is not, by itself, a sufficient justification for offering multiple coverage levels when premiums cannot reflect individual costs. Under full community rating, consumers sort into each plan based on willingness to pay, while efficiency depends on the comparison of willingness to pay and marginal cost for the coverage difference. A single price can implement the efficient allocation only if individuals with higher WTP also tend to have higher social surplus from the extra coverage, so that price-based sorting aligns with efficiency. I therefore examine the correlation between WTP and $SS = WTP - MC$ across menu types: a positive correlation implies that a single price can be efficient, while a negative correlation implies that no single price can implement the efficient allocation.

In the numeric simulation, I draw 5,000 types that differ in both loss distributions (derived from Truven MarketScan data) and risk aversion (the risk aversion distribution and its correlation with risk types follow Handel (2013) and Marone and Sabety (2022)). I consider two

scenarios: risk types and risk aversion are independent, and risk aversion is positively correlated with risk. I then compute the WTP–SS correlation for each observed menu in the KFF data and compare its distribution across the three menu types (type 1: WTP- and RP-ordered; type 2: WTP-ordered but not RP-ordered; type 3: not WTP-ordered). I find that all KFF menus yield $SS > 0$ for some types and $SS < 0$ for other types, so the value of offering choice remains potentially relevant across the observed designs. Moreover, I find that type 3 menus typically generate more positive correlations: when risk and risk aversion are independent, the correlation distributions are mostly negative for menu types 1 and 2, while type 3 shows substantially more mass near zero and on the positive side. When risk aversion is positively correlated with risk, the distributions shift right for all menu types: type 1 centers near zero, type 2 remains mostly negative, and type 3 becomes clearly positive on average. These results suggest that the common conclusion in the vertical-choice literature about the limited efficiency gains from offering multiple plans may not extend to horizontally differentiated menus, where the WTP–SS relationship can look fundamentally different.

This paper contributes to three strands of literature. First, it relates to the classic theoretical literature on optimal cost-sharing and risk protection of insruance contracts. Theoretical work shows that full insurance is optimal under actuarially fair pricing (Mossin, 1968; Rothschild and Stiglitz, 1970), and that straight-deductible contracts are optimal under proportional loading (Arrow, 1963; Gollier and Schlesinger, 1996). I extend this line by characterizing the risk-protection value of partial insurance menus and the ordering conditions under which it is preserved.

Second, this paper contributes to the broader literature on allocation and efficiency in health insurance markets. Much of this work studies how pricing institutions and policy interventions shape sorting and efficiency without explicitly specifying the design of different contracts (Einav, Finkelstein and Cullen, 2010; Handel, Hendel and Whinston, 2015; Weyl and Veiga, 2016). Relatedly, a growing literature studies the interaction of adverse selection and cost-sharing designs (Liu, 2025; Liu and Sydnor, 2022). I add to the literature on the value of offering multiple contracts (Azevedo and Gottlieb, 2017; Ericson and Sydnor, 2017; Marone and Sabety, 2022; Chade et al., 2022; Ho and Lee, 2023) by allowing menus that are not ordered in cost-sharing designs.

The paper also connects to the literature on product differentiation of health insurance contracts. In health insurance markets, horizontal differentiation is typically associated with non-financial attributes such as hospital networks(Tilipman, 2022; Olivella and Vera-Hernández, 2007; Shepard, 2022; Wagner, 2023). I show that, even along purely financial attributes, menus can be horizontally differentiated when OOP schedules cross.

The rest of the paper is organized as follows. Section 2 develops the conceptual framework and the ordering results. Section 3 presents empirical evidence on the prevalence of menu types in private and public markets. Section 4 presents the simulation exercises on comparative statics

and the value of offering choice. Section 5 concludes.

2 Conceptual Framework

In this section, I formalize the ordering concepts and show how they map to simple properties of OOP schedules. The framework isolates the conditions under which "vertical choice" holds and clarifies when a stronger risk-protection ordering applies. All proofs are collected in Appendix A.

Setup. An individual faces a stochastic loss $x \in [0, \bar{x}]$ with distribution F. Let w denote initial wealth (or income). A contract is a mapping $h(\cdot)$ from losses to out-of-pocket (OOP) spending, and utility is defined over consumption. Preferences are expected utility with $u' > 0$ and $u'' < 0$, and there is no moral hazard. A type θ summarizes both the loss distribution and risk preferences, allowing heterogeneity in expected losses and risk aversion.

Contract space. I focus on contracts that satisfy the following conditions. First, $0 \leq h(x) \leq x$ for all x, so OOP spending is nonnegative and never exceeds the loss. Second, $h(\cdot)$ is continuous with a derivative except at finitely many kink points, which matches the piecewise-linear structure of common plans. Third, $h'(x) \geq 0$, so OOP spending is nondecreasing in losses. Fourth, $h'(x) \leq 1$, so the incremental OOP payment is no larger than the incremental loss. This last condition reflects that health insurance contracts are typically defined over annual cumulative spending, so the enrollee's dollar payment should not rise faster than total spending as losses accumulate over the year. These restrictions capture standard health insurance designs with deductibles, coinsurance, and maximum out-of-pocket (MOOP) limits.

Contract Ordering. With the contract space in place, I first define the standard vertical-choice notion in terms of willingness-to-pay and then introduce the stronger, risk-protection notion.

Definition 1 (WTP Ordering)**.** *Given two contracts h_1 and h_2, willingness-to-pay (WTP) for contract 1 relative to contract 2 is the amount WTP that solves*

$$\mathbb{E}[u(w - h_1(x) - WTP)] = \mathbb{E}[u(w - h_2(x))]\,.$$

The menu is WTP-ordered *if $WTP > 0$ for all types.*

The WTP-ordering is the typical definition of vertical differentiation: everyone prefers one contract when both are free. The primary purpose of insurance, however, is risk protection, so it is natural to ask whether a preference ordering survives when contracts are priced fairly.

Definition 2 (RP Ordering)**.** *Define the actuarially fair premium for contract i as $\pi_i = \mathbb{E}[x - h_i(x)]$. The risk premium (RP) for contract 1 relative to contract 2 solves*

$$\mathbb{E}[u(w - h_1(x) - \pi_1 - RP)] = \mathbb{E}[u(w - h_2(x) - \pi_2)]\,.$$

The menu is RP-ordered *if $WTP > 0$ and $RP > 0$ for all types.*

The RP-ordering is a stronger ordering that requires the preferred contract to deliver strictly more risk protection at actuarially fair prices. These definitions raise a natural question: under what conditions on the OOP schedules such that the contracts are WTP-ordered and, more stringently, RP-ordered? The next two propositions provide simple sufficient conditions that connect the ordering notions to the shape of the OOP schedules.

Proposition 1 (WTP Ordering)**.** *Given two contracts with $h_1(x) \leq h_2(x)$ for all x and strict inequality for some x, then $WTP > 0$ for all types.*

Proposition 1 formalizes the usual IO notion of vertical ordering: if one contract's OOP schedule lies weakly below another's at every loss state, then everyone prefers it when both are free. From the proposition, it follows that all insurance contracts have positive WTP relative to no insurance. Note that WTP-ordering is a partial ordering: many contracts may not be ordered in the WTP sense. In Figure 1, the horizontal axis is loss x, the vertical axis is OOP spending $h(x)$, and each line is a contract's OOP schedule. Panel (a) and Panel (b) show the cases where one line is everywhere below the other, so $WTP > 0$ for all types. In contrast, Panel (c) shows a case where the two contracts violate WTP ordering and are not vertically differentiated.

The WTP ordering does not imply RP ordering. In fact, the contract with lower OOP may deliver less risk protection after premiums adjust to expected payouts. To guarantee that the preferred contract remains preferred at actuarially fair prices, a stronger condition is needed:

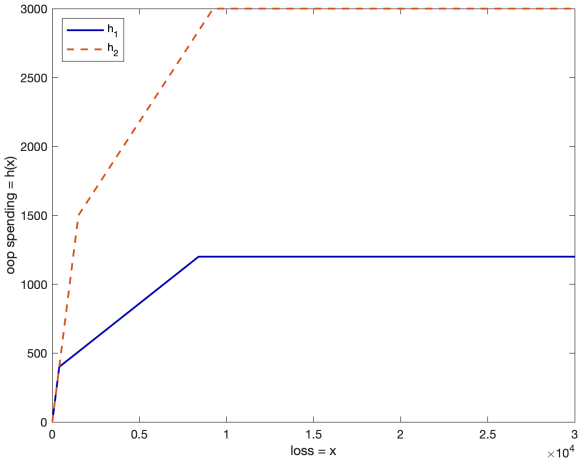
Proposition 2 (RP Ordering)**.** *Given two contracts with $h_1(x) \leq h_2(x)$ for all x and $h_1'(x) \leq h_2'(x)$ almost everywhere, then $WTP > 0$ and $RP > 0$ for all types.*

Proposition 2 adds a slope restriction: to obtain the RP-ordering, the lower-OOP contract must also be weakly flatter in losses for the preference ordering to hold at actuarially fair prices. The condition is intuitive: to provide more risk protection, the "better" contract must reduce OOP not only in levels but also in how fast OOP rises with losses.

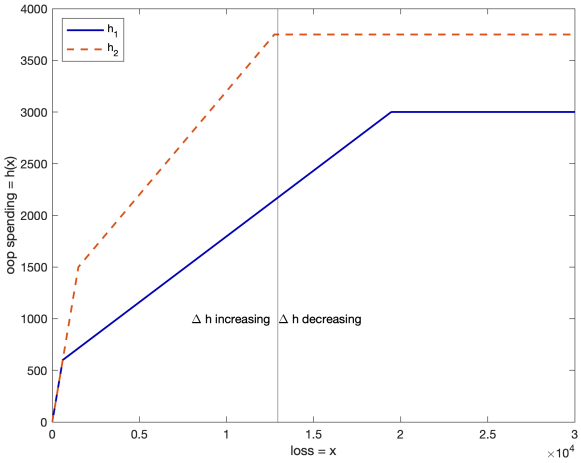
This condition is satisfied for any insurance contract relative to no insurance because $h'(x) \leq 1, \forall x, \forall h$, and it is satisfied for full insurance relative to any partial insurance because $h'(x) \geq 0, \forall x, \forall h$. However, for two partial insurance contracts, the ordering may not hold. Figure 1 (a) illustrates a menu that meets both requirements: one OOP schedule lies below the other and becomes flatter when loss increases. Figure 1 (b) contrasts a case where the lower-OOP contract is not uniformly flatter: the more generous contract hits the MOOP at a larger loss state than

the less generous contract. As a result, the OOP difference increases up to the vertical line, and then starts to decrease. Such a menu violates the $h_1'(x) \leq h_2'(x)$ condition, so RP-ordering fails even though WTP-ordering holds. In other words, for menus as in Figure 1 (b), there might exist risk averse individuals who prefer the less generous coverage to the more generous coverage if both are priced actuarially fairly.

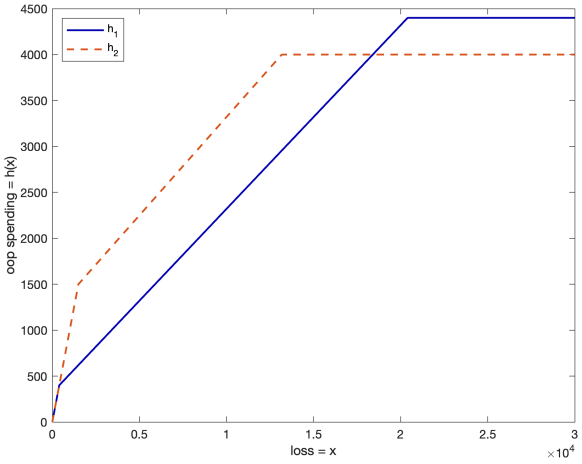
Menu classification. The two notions yield three empirically relevant menu types. Case (i) is WTP- and RP-ordered, illustrated in Figure 1 (a). Case (ii) is WTP-ordered but not RP-ordered, illustrated in Figure 1 (b). Case (iii) is not WTP-ordered, illustrated in Figure 1 (c). The first two are vertically differentiated; the third is horizontally differentiated in the cost-sharing dimension.



(a) WTP- and RP-ordered



(b) WTP-ordered, not RP-ordered



(c) Not WTP-ordered

Figure 1: Stylized OOP schedules for the three menu types.

Note: Each panel plots out-of-pocket spending $h(x)$ against loss x for two contracts. Panel (a) is WTP- and RP-ordered, panel (b) is WTP-ordered but not RP-ordered, and panel (c) violates WTP ordering.

3 Empirical Analysis

The conceptual framework highlights three menu types, but most of the existing literature focuses on menus that are both WTP- and RP-ordered. An open question is whether this focus is simply a matter of empirical relevance: perhaps non-RP-ordered or non-WTP-ordered menus are rare in practice. This section addresses that question by examining the prevalence of each menu type in observed plan designs. I begin with private employer-sponsored markets in the US, where firms and insurers have substantial latitude over cost-sharing design, and then present evidence from regulated public markets where plan rules are largely set by policy.

3.1 Private Health Insurance Markets

I first examine the issue using employer-sponsored insurance, where firms and insurers can freely choose deductibles, coinsurance, and maximum out-of-pocket limits. This setting provides a natural benchmark for the menu types that arise when plan design is driven by market forces rather than regulation.

Data and sample. I use the Kaiser Family Foundation Employer Health Benefits Survey (KFF EHBS) for 2015, a nationally representative survey of employer-sponsored plans and premiums in the US (Liu and Sydnor, 2022). The EHBS covers private and public employers with three or more workers across industries and firm-size categories and reports plan attributes such as plan network type (e.g., HMO vs. PPO), cost-sharing features, premiums, and enrollment. I restrict attention to the single-coverage (which has complete plan details) and to firms that report a high-deductible plan (deductible above $1,000 and paired with an HSA/HRA) and exactly one lower-deductible plan, yielding 331 firms. The two plans may differ in network type, and the KFF data do not contain detailed network information; the analysis here focuses on their cost-sharing designs. The sample construction mirrors Liu and Sydnor (2022).

To compare plans on a common basis, I translate each plan's complex cost-sharing rules into a simplified three-arm design, following the procedure introduced by Ericson et al. (2021) and used in Liu and Sydnor (2022). The procedure first computes each plan's actuarial value (AV) given the full cost-sharing designs for all benefits using the CMS AV calculator, then solves for a single coinsurance rate that reproduces the same AV while holding the deductible and MOOP fixed. This yields a simplified plan with only three basic cost-sharing features—deductible, coinsurance, and MOOP—so that OOP is a simple function of total annual spending. The coinsurance rate in this simplified plan approximates the average fraction of medical bills paid out of pocket above the deductible. I then compare the two plans' simplified schedules to assign each menu to one of the three types.

Main results. The distribution of menu types is far from the vertical-ordered benchmark. Table 1 reports the fractions by menu type. Only 14.6% of menus are both WTP- and RP-

ordered (case 1), 48.0% are WTP-ordered but not RP-ordered (case 2), and 37.4% are not WTP-ordered (case 3). The conversion to a simplified three-arm design may affect the classification of cases 1 and 2 because the estimated coinsurance rates affect the loss level at which a contract reaches the maximum out-of-pocket limit (MOOP). The classification of case 3 menus is unaffected by the simplification because it depends only on the deductible and MOOP values of the two contracts.

Table 1: Fraction of Menus by Type in the KFF EHBS Sample

Menu type	(1) WTP & RP ordered	(2) WTP ordered, not RP ordered	(3) Not WTP ordered
Fraction of menus	14.6%	48.0%	37.4%

Note: Fractions are computed from the KFF EHBS 2015 sample of 331 firms with two cost-sharing contracts for single coverage. Menu types are classified using the simplified three-arm design described above.

3.2 Public-market examples

Public markets provide a complementary setting in which plan designs are heavily shaped by regulation rather than by firms. The Swiss and Covered California examples illustrate how different regulatory environments can generate different menu types.

Swiss market. In Switzerland, deductible options and cost-sharing rules are tightly regulated. Table 2 shows a representative deductible menu where coinsurance and the cap are fixed by policy while deductibles vary across options. The implied OOP schedules are nested: higher deductibles shift the OOP schedule upward in a way that preserves ordering at every loss level. As a result, the menu is WTP- and RP-ordered. Figure 2 illustrates the ordered OOP schedules for this menu.

Table 2: Deductible choice menu

Deductible	Co-insurance rate	Cap co-insurance
CHF 300	10%	CHF 700
CHF 500	10%	CHF 700
CHF 1,000	10%	CHF 700
CHF 1,500	10%	CHF 700
CHF 2,000	10%	CHF 700
CHF 2,500	10%	CHF 700

Note: This table reports the regulated Swiss deductible menu.

Covered California. Covered California is California's ACA marketplace, where standardized plan designs are set by the state. Table 3 reports the 2017 standardized designs, and I use the same AV-based simplification to map them into three-arm schedules. The Bronze and Gold

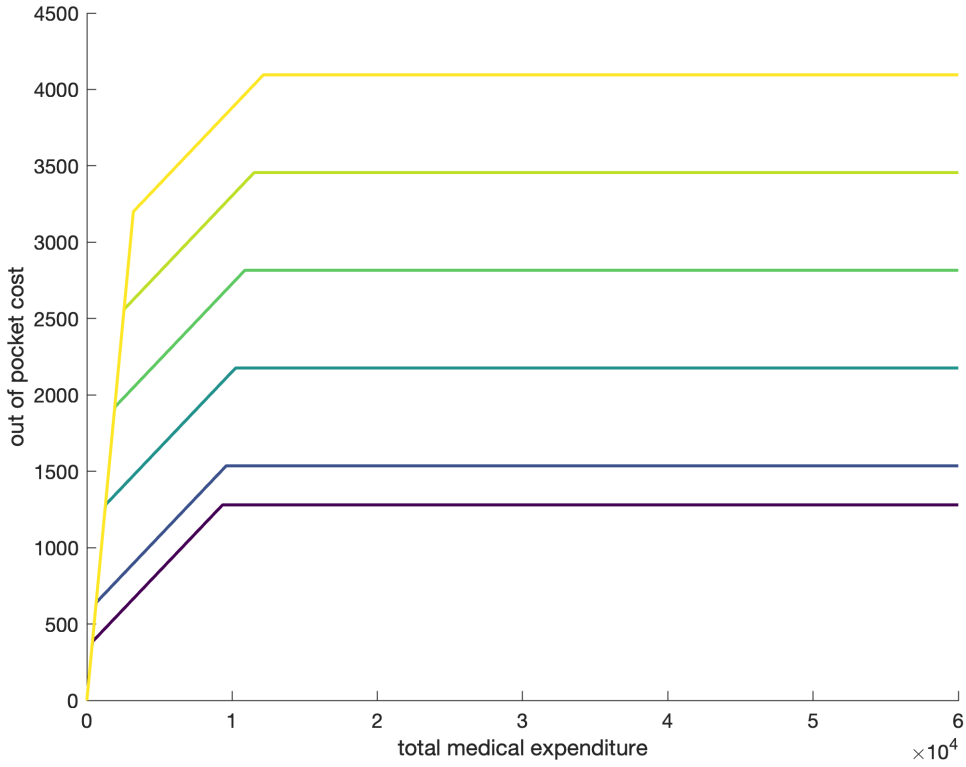


Figure 2: Swiss market example: ordered OOP schedules.

Note: The figure plots OOP schedules implied by the deductible menu in Table 2. The horizontal axis is total loss and the vertical axis is OOP spending.

plans are vertically differentiated: Bronze has a higher deductible, higher MOOP, and much higher coinsurance, yielding an OOP schedule that lies above Gold for all loss levels. Moreover, because the Bronze plan hits the MOOP at a smaller loss state, the two plans are WTP-ordered but are not RP-ordered.

By contrast, the Gold and Silver plans are not vertically ordered. Gold has a lower deductible but a higher coinsurance rate in the simplified design; Silver has a higher deductible but lower coinsurance. These features lead to crossing OOP schedules, so the Gold vs. Silver comparison fails WTP ordering. Figure 3 shows the resulting non-ordered schedules.

Table 3: 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

Note: This table reports the deductibles and MOOPs in the standardized plan designs for California in 2017. The simplified coinsurance rate in the final column is calculated to match AV using the CMS calculator. All designs are for individual, first-tier in-network coverage and are used to construct the OOP schedules in Figure 3.

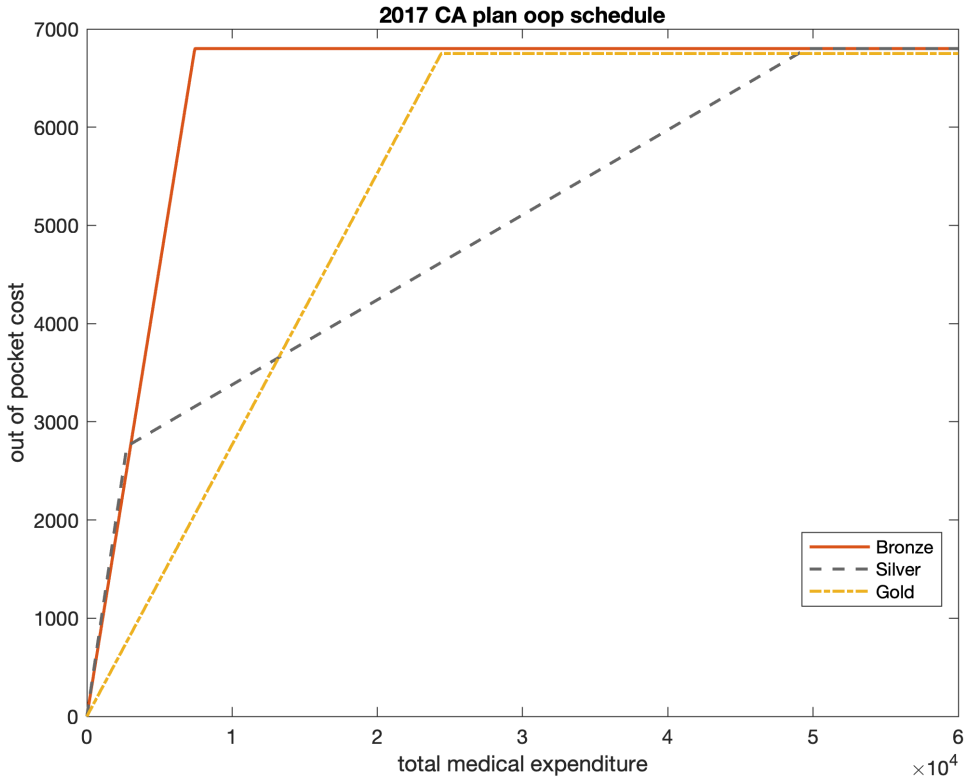


Figure 3: Covered California example: non-ordered OOP schedules.

Note: The OOP schedules correspond to the standardized plan designs in Table 3 after applying the same AV-based three-arm simplification.

4 Numeric Simulation

The empirical evidence shows that non-RP-ordered and non-WTP-ordered menus are common. This section uses simulation to evaluate whether insights from the vertical-choice literature (which focuses on case 1 menus) extend to the broader set of menus observed in practice. The simulations provide a clean environment in which preferences and loss distributions are disciplined by data and the literature but choice frictions are absent, allowing us to isolate the mechanisms linking menu design to welfare.

4.1 Simulation Setup

Across all exercises, I assume that individuals have CARA utility $u(x) = -\exp(-\gamma x)$, where γ is the absolute risk aversion coefficient. I use CARA preferences both for comparability with the existing literature and because it allows transparent comparative statics in risk aversion. Social marginal cost (MC) of additional coverage is modeled using proportional loading: for a menu with contracts h_1 and h_2, $MC = \lambda\,\mathbb{E}[h_1(x) - h_2(x)]$ with $\lambda = 1.2$. This approach differs from models where social costs arise through moral hazard; here the goal is to isolate how menu design and risk protection interact in a setting without behaviorally complex responses.

4.2 Comparative Statics

The first set of exercises studies how WTP and MC differences between the two contracts respond to risk type and risk aversion. To obtain a range of loss distributions that are well ordered, I use the 2013 Truven MarketScan claims data, and group individuals into 44 types based on their ages (21–64). For each age, I then estimate a three-parameter shifted lognormal distribution of x:

$$\log(x + \kappa) \sim N(\mu, \sigma^2). \tag{1}$$

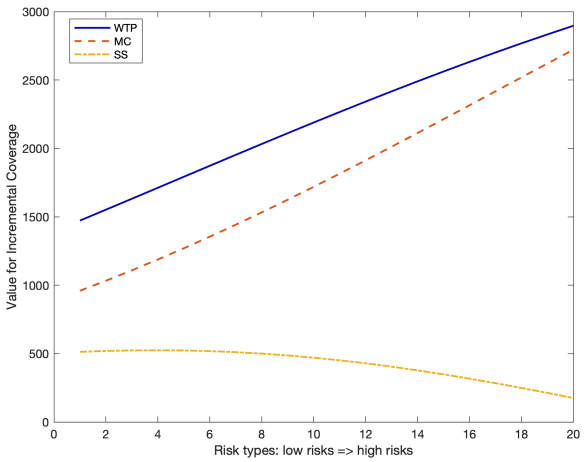
I then fix the average σ and κ across the 44 distributions and vary μ across the full range to generate a sequence of distributions that satisfy the monotone likelihood ratio property (MLRP). The method creates a sequence of distributions that are empirically relevant.

To study how WTP and MC change with risk types, I hold risk aversion at $\gamma = 0.0004$ (the mean and median estimate from Handel, 2013) and vary the loss distributions. To study how WTP change with risk aversion, I hold the loss distribution at age 40 and vary γ.

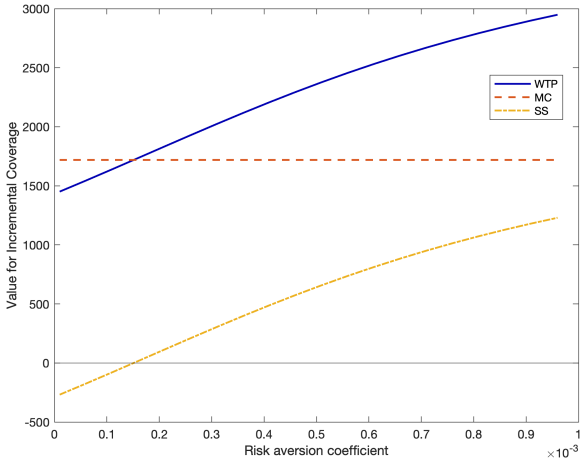
I consider three stylized menus corresponding to the three menu types. Case 1 (WTP and RP ordered) uses two straight-deductible plans with \$1,000 and \$5,000 deductibles. Case 2 (WTP ordered but not RP ordered) uses two plans with zero deductibles and a \$4,000 MOOP; one has coinsurance 0.5 and the other has coinsurance 1 (so the less-coverage plan is a straight-deductible plan with a \$4,000 deductible). Case 3 (not WTP ordered) uses two plans with coinsurance 0.1 in the interior but different deductibles and MOOPs: one has deductible 0 and MOOP \$6,000, the other has deductible \$1,000 and MOOP \$2,000. For each menu I compute WTP and MC for the coverage difference between the two plans.

The comparative statics deliver three key patterns. First, for the WTP and RP-ordered menu, the standard intuition holds: both WTP and MC increase with risk type, and WTP increases with risk aversion (Figures 4a–4b). In these panels, the horizontal axis is risk type (left) or the risk aversion coefficient γ (right). The solid blue line shows WTP and the dashed red line shows MC for the two-contract difference; both move upward as risk increases, and WTP shifts upward with higher risk aversion. This is the typical scenario people have in mind when modeling the value of vertical insurance contracts.

However, for menus that are WTP ordered but not RP ordered, these monotonicities break down: WTP need not be monotone in risk type or risk aversion, and MC can be non-monotone in risk type. In Figures 5a–5b, the horizontal axis again reports risk type (left) or γ (right), and the vertical axis reports WTP (solid blue line), MC (dashed red line) and SS (WTP - MC, the dotted yellow line) for the coverage difference. The plotted lines show that WTP can flatten or decline over ranges of risk or risk aversion, reflecting that the contract providing lower OOP in levels is not more preferred to more risky or risk averse individuals. Also note that MC is not monotone with regard to risk types. Intuitively, given two contracts with the same MOOP, higher risks are likely to hit the MOOP for both plans, resulting in the difference in coverage



(a) RP-ordered menu: WTP vs. risk type

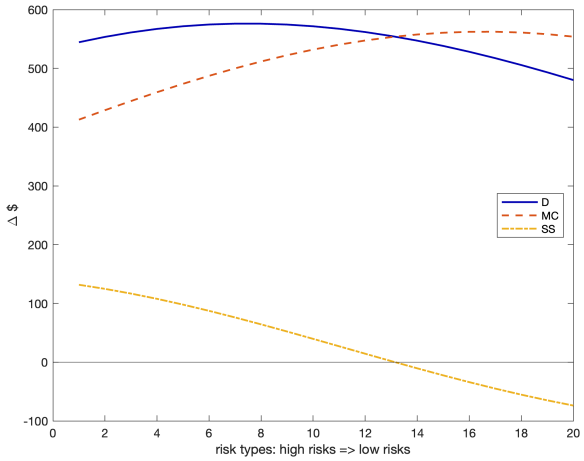


(b) RP-ordered menu: WTP vs. risk aversion

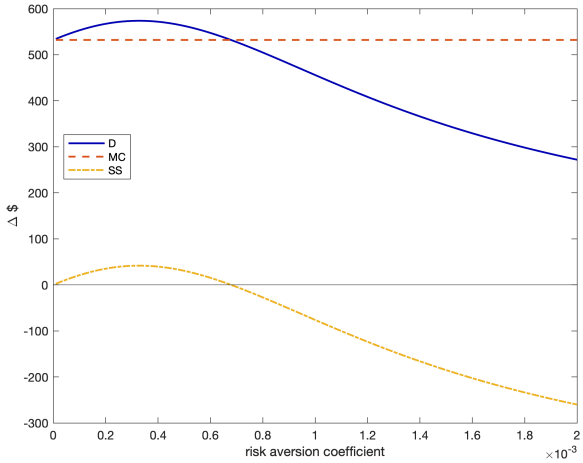
Figure 4: Comparative statics for RP-ordered menus.

Note: Solid blue line is WTP and dashed red line is MC for the coverage difference between two contracts in a RP-ordered menu. Left panel varies risk type; right panel varies risk aversion. Simulated under CARA preferences as described in the text.

between the two plans for smaller losses not very different for them. In contrast, such different coverage may be very different for individuals more likely to experience smaller losses. The non-monotonicity of MC between two contracts that are WTP ordered but not RP ordered is also highlighted by Li, Liu and Sydnor (2025).



(a) WTP-ordered, not RP-ordered: WTP vs. risk type



(b) WTP-ordered, not RP-ordered: WTP vs. risk aversion

Figure 5: Comparative statics for WTP-ordered but not RP-ordered menus.

Note: Solid blue line is WTP, dashed red line is MC, and dotted yellow line is $SS = WTP - MC$ for the coverage difference between two contracts. Left panel varies risk type; right panel varies risk aversion.

Finally, for non-WTP-ordered menus, WTP can be negative for some types (Figure 6). The horizontal axis is risk type and the vertical axis is WTP; the curve crosses zero, indicating that some types would prefer one contract and others prefer the other contract when both are free.

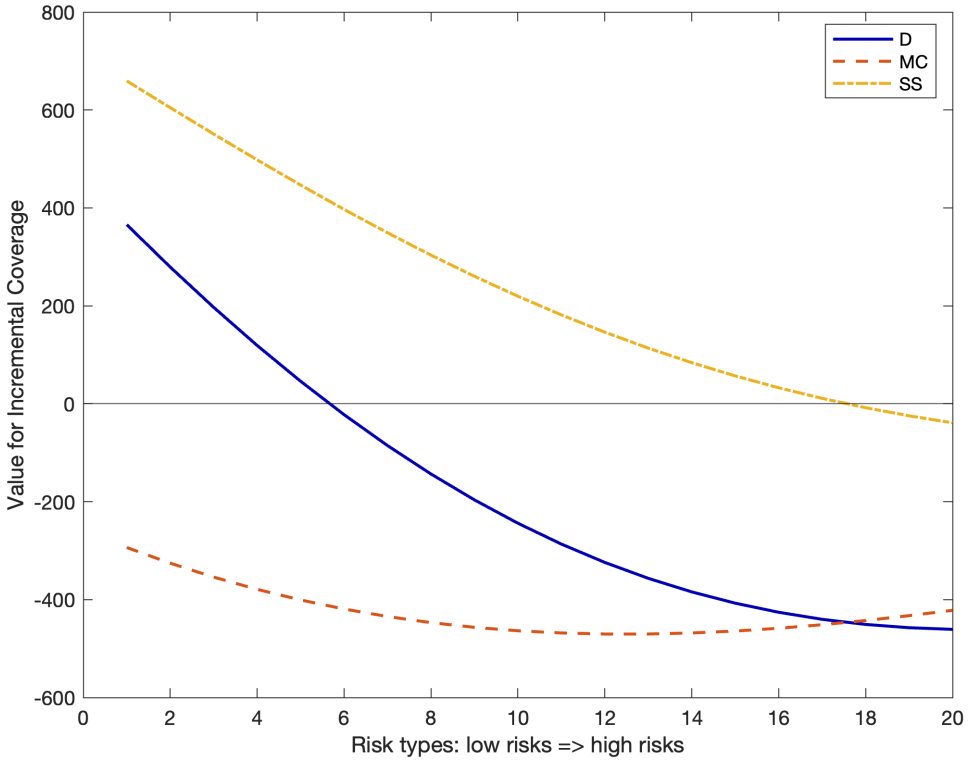


Figure 6: Comparative statics, non-WTP-ordered menu: WTP can be negative.

Note: The figure plots WTP for the coverage difference between two contracts in a non-WTP-ordered menu as risk type varies. WTP crosses zero, implying that some types prefer the lower-coverage contract even when both are free.

4.3 The Value of Offering Choice

The second exercise evaluates whether offering two plans improves efficiency relative to a single plan. I follow the framework of Marone and Sabety (2022). A social planner can set any prices but must charge the same price to everyone for a given contract (full community rating). Efficient allocation assigns individuals to contract 1 (instead of contract 2) when WTP exceeds MC for the difference between the two contracts and to contract 2 when WTP is below MC. Under full community rating, a single price partitions individuals by their WTP. A single price can achieve efficient allocation only if higher-WTP individuals also tend to have higher social surplus from extra coverage. Figure 7 illustrates the logic: if WTP crosses MC from above (Panel A), a single price can implement the efficient assignment; if WTP crosses from below (Panel B), no single price can do so. Define SS = WTP - MC. The correlation between WTP and SS summarizes whether a single price can implement the efficient allocation: positive correlations are favorable, while negative correlations indicate that no single price can sort individuals efficiently.

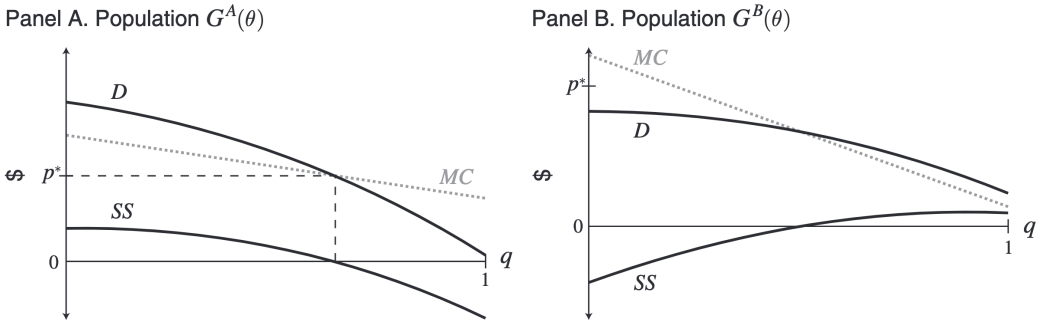


FIGURE 1. EXAMPLES IN WHICH VERTICAL CHOICE (A) *IS* AND (B) *IS NOT* EFFICIENT

Figure 7: Efficiency with full community rating (relabeling D as WTP), based on Marone and Sabety (2022).

Note: WTP denotes willingness to pay for the higher-coverage contract and MC denotes social marginal cost. If WTP crosses MC from above (Panel A), a single price can implement the efficient allocation; if it crosses from below (Panel B), it cannot.

Following this framework, I simulate WTP, MC and SS for each menu in the KFF dataset used in section 3.1 under a pre-specified population. Each individual in the population is characterized by their loss distributions and risk aversion coefficient. I assume that all the medical expenditure follow a shifted log normal distribution with parameters (μ, σ, κ), and the three parameters are jointly normally distributed:

$$\begin{pmatrix} \mu \\ \sigma \\ \kappa \end{pmatrix} \sim N(\bar{\mu}, \Sigma) \tag{2}$$

To estimate the distribution, I apply the k-means clustering method to Truven MarketScan

data using 2012 characteristics (age, gender, total medical expenditure, pre-existing conditions, and employment status) and forming 100 clusters (Liu, 2025). For each cluster, I fit a shifted lognormal distribution to the 2013 spending. Finally, I estimate $\bar{\mu}$ and Σ using maximum likelihood estimation based on the 100 fitted distributions. Table 4 reports the estimated parameters.

Table 4: Estimated Parameters for the Shifted Lognormal Distribution

	μ	σ	κ
$\bar{\mu}$	6.9782	1.4617	155.6740

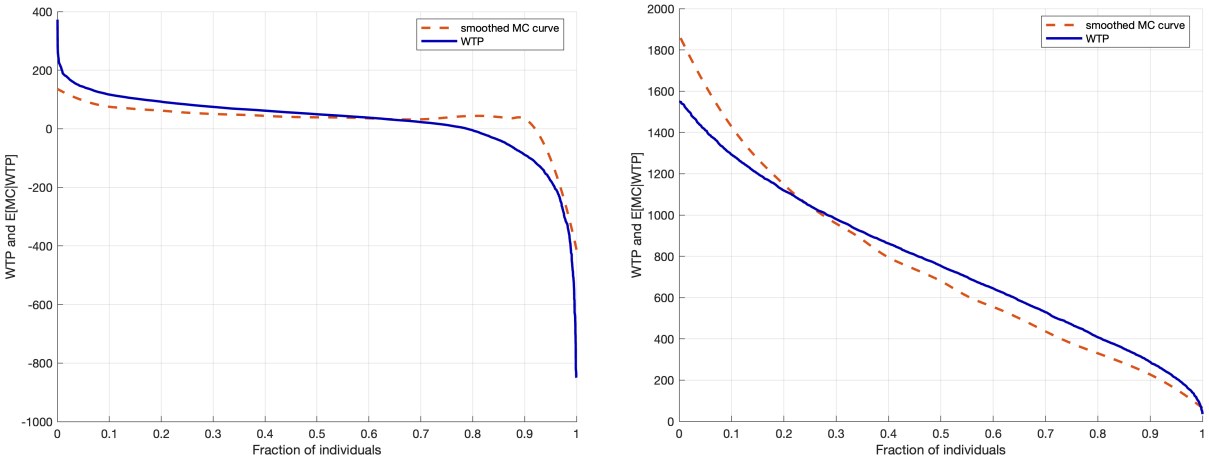
	μ	σ	κ
	0.0014	-0.0002	0.0082
Σ ($\times 10^3$)	-0.0002	0.0001	-0.0020
	0.0082	-0.0020	1.1950

Note: The top panel reports the mean vector $\bar{\mu}$. The lower panel reports the covariance matrix Σ (shown in units of 10^3). Estimates are based on 100 clusters from 2012 MarketScan characteristics and fitted 2013 spending distributions.

For risk aversion, I consider two scenarios. In the first case, I assume risk aversion is independent of risk types. I model $\log(\gamma)$ as drawn from a normal distribution with mean and standard deviation from Handel (2013). In the second case, I assume risk aversion is correlated with risk types. Specifically, I model $(\mu, \sigma, \kappa, \log \gamma)$ as jointly normal and impose a correlation between μ and γ based on Marone and Sabety (2022). (I assume there is no correlation between σ and $\log(\gamma)$ and κ and γ.)

After obtaining the distributions of the parameters, I draw 5000 individuals and compute WTP, MC, and SS = WTP - MC for the coverage difference of each menu. Figure 8 shows two menus from the simulations when risk types and risk aversion are independent: the left panel is the case where a single price can be efficient (Panel (a) in Figure 7) and the right panel is a backward-sorting case (Panel (b) in Figure 7).

The question then is whether the left-panel pattern or the right-panel pattern is more likely for each menu type. I summarize sorting outcomes by tabulating the correlation between WTP and SS for each menu and plotting the distribution by menu type. The exercise delivers several insights. First, in no menu do I find SS positive for all individuals or SS negative for all individuals, so a single plan is never uniformly optimal. Second, the correlation between WTP and SS varies substantially across menu types and across the two scenarios (independent versus positively correlated risk and risk aversion). Figure 9a shows the correlation distribution when risk aversion is independent of risk types. The blue, red, and yellow lines indicate the distributions for type 1–3 menus, respectively. Figure 9b shows the same distributions when risk aversion and risk types are positively correlated. In both cases, type 3 menus seem more likely to yield a positive correlation than the other two types. Interestingly, WTP-ordered but not RP-ordered menus appear most likely to yield a negative correlation, suggesting that such



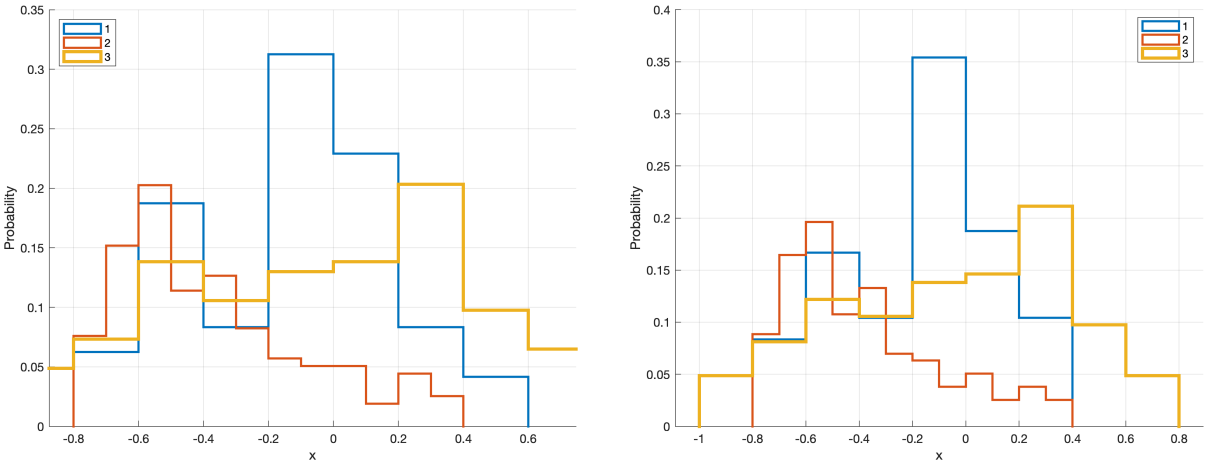
(a) A single price can be efficient

(b) A single price can never be efficient

Figure 8: Examples of menus in simulations.

Note: Each panel plots WTP and MC against the fraction of individuals, ordered by WTP, for a simulated KFF menu when risk types and risk aversion are independent. The left panel corresponds to the case where WTP crosses MC from above, while the right panel corresponds to backward sorting (WTP crosses MC from below).

menus are less likely to deliver an efficient allocation under a single price. This pattern reinforces the main lesson of the paper: conclusions about the value of offering choice depend critically on menu type and do not carry over automatically from vertically ordered menus to the broader set of designs observed in practice.



(a) Risk aversion independent of risk type

(b) Risk aversion positively correlated with risk type

Figure 9: Correlation between WTP and SS by menu type under two scenarios.

Note: Each panel reports the distribution of correlations between WTP and SS = WTP - MC across KFF menus under the indicated preference scenario. Positive correlations imply that a single price can implement the efficient allocation; negative correlations imply that it cannot.

5 Conclusion

This paper distinguishes between WTP-ordering and the stronger RP-ordering and links each to simple conditions on OOP schedules. The empirical evidence shows that non-RP-ordered and non-WTP-ordered menus are common, both in employer-sponsored markets and in regulated public markets. These findings suggest that the narrow focus on vertically ordered menus in the prior literature misses a large share of observed plan designs.

The simulation results reinforce this point. Comparative statics and the value of offering choice can differ sharply across menu types, and conclusions drawn from vertically ordered menus do not automatically extend to horizontally differentiated designs. The prevalence of non-vertical choices in health insurance markets suggests room for more theoretical work on why such menus are offered and whether they are socially beneficial. A related direction is to revisit competition in health insurance markets when cost-sharing designs introduce meaningful horizontal differentiation, alongside the more commonly studied non-financial dimensions.

References

Arrow, Kenneth J. 1963. "Uncertainty and the Welfare Economics of Medical Care." *American Economic Review*, 53(5): 941–973.

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

Chade, Hector, Victoria R. Marone, Amanda Starc, and Jeroen Swinkels. 2022. "Multidimensional Screening and Menu Design in Health Insurance Markets." National Bureau of Economic Research Working Paper 30542. Revised May 2024.

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 Marzilli, 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.

Ericson, Keith Marzilli, Philipp Kircher, Johannes Spinnewijn, and Amanda Starc. 2021. "Inferring Risk Perceptions and Preferences Using Choice from Insurance Menus: Theory and Evidence." *Economic Journal*, 131(634): 713–744.

Gollier, Christian, and Harris Schlesinger. 1996. "Arrow's Theorem on the Optimality of Deductibles: A Stochastic Dominance Approach." *Economic Theory*, 7(2): 359–363.

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

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

Ho, Kate, and Robin S. Lee. 2023. "Health Insurance Menu Design for Large Employers." *The RAND Journal of Economics*, 54(4): 598–637.

Li, Honglin, Chenyuan Liu, and Justin R. Sydnor. 2025. "Community Rating and Vertical Price Distortions in Insurance Menus." National Bureau of Economic Research Working Paper 34367.

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

Liu, Chenyuan, and Justin Sydnor. 2022. "Dominated Options in Health Insurance Plans." *American Economic Journal: Economic Policy*, 14(1): 1–25.

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

Mossin, Jan. 1968. "Aspects of Rational Insurance Purchasing." *Journal of Political Economy*, 76(4): 553–568.

Olivella, Pau, and Marcos Vera-Hernández. 2007. "Competition among Differentiated Health Plans under Adverse Selection." *Journal of Health Economics*, 26(2): 233–250.

Rothschild, Michael, and Joseph E. Stiglitz. 1970. "Increasing Risk: I. A Definition." *Journal of Economic Theory*, 2(3): 225–243.

Shepard, Mark. 2022. "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange." *American Economic Review*, 112(2): 578–615.

Tilipman, Nicholas. 2022. "Employer Incentives and Distortions in Health Insurance Design: Implications for Welfare and Costs." *American Economic Review*, 112(3): 998–1037.

Wagner, Myles. 2023. "How Can Menu Design Address Adverse Selection? Evidence from a Health Insurance Exchange." Working paper, September 28, 2023.

Weyl, E. Glen, and André Veiga. 2016. "Pricing Institutions and the Welfare Cost of Adverse Selection." SSRN Working Paper, 2016.

A Proofs

Proof of Proposition 1. Fix any type with loss distribution F and utility u strictly increasing. The assumption $h_1(x) \leq h_2(x)$ for all x with strict inequality on a set of positive probability implies $w - h_1(x) \geq w - h_2(x)$ almost surely and strictly on a set of positive probability. By strict monotonicity, $u(w - h_1(x)) \geq u(w - h_2(x))$ almost surely and strictly on a set of positive probability, so $\mathbb{E}[u(w - h_1(x))] > \mathbb{E}[u(w - h_2(x))]$. The compensating variation WTP that makes the individual indifferent therefore satisfies $WTP > 0$. □

Proof of Proposition 2. Under the stated conditions, $h_1(x) \leq h_2(x)$ implies $WTP > 0$ by Proposition 1. Let π_i be the actuarially fair premium for contract i. Define $\Delta(x) = h_2(x) - h_1(x) \geq 0$. The slope condition $h_1'(x) \leq h_2'(x)$ implies $\Delta(x)$ is weakly increasing in x, so contract 1 provides relatively more indemnity in higher-loss states. Write post-premium wealth as

$$W_i(x) = w - h_i(x) - \pi_i, \quad i = 1, 2.$$

Since $\pi_1 - \pi_2 = \mathbb{E}[\Delta(x)]$, we have

$$W_1(x) = W_2(x) + \Delta(x) - \mathbb{E}[\Delta(x)].$$

Let $g(x) = \Delta(x) - \mathbb{E}[\Delta(x)]$, so $\mathbb{E}[g(x)] = 0$ and g is weakly increasing. Because h_2 is increasing in x, $W_2(x)$ is weakly decreasing in x. To show second-order stochastic dominance, use the integrated CDF criterion. Let F_i denote the CDF of W_i. For any threshold y, define $x_i(y)$ as the (weak) inverse of W_i so that $W_i(x_i(y)) = y$. Since $W_1(x) \geq W_2(x)$ for all x with strict inequality on a set of positive probability, we have $x_1(y) \geq x_2(y)$ for all y. Then

$$\int_{-\infty}^{y} \big(F_1(t) - F_2(t)\big)\, dt = \int_{x_2(y)}^{x_1(y)} \big(W_2'(x) - W_1'(x)\big)\, dF_X,$$

and $W_2'(x) - W_1'(x) = h_1'(x) - h_2'(x) \leq 0$ by the slope condition. Hence the integral is nonpositive for all y, with strict inequality for some y when the slope condition is strict on a set of positive probability. Therefore W_1 second-order stochastically dominates W_2, which implies $\mathbb{E}[u(W_1)] \geq \mathbb{E}[u(W_2)]$ for all concave u, with strict inequality for strictly concave u. Hence $RP > 0$ for all types. □