ONLINE APPENDIX

When Should There Be Vertical Choice in Health Insurance Markets?

Victoria R. Marone Adrienne Sabety

September 2021

Appendix A Derivations and Proofs

A.1 Derivation of Willingness to Pay

The expected utility of a type-θ consumer with initial income $\hat{y}$ for contract x at premium p is given by $U(x, p, \theta)$, as defined in Equation 1 and repeated here:

$$U(x, p, \theta) = \mathbb{E}_l[\ u_\psi(\hat{y} - p - c_x^*(l, \omega, x) + b^*(l, \omega, x))].$$

The corresponding certainty equivalent $CE(x, p, \theta)$ solves $u(CE(x, p, \theta)) = U(x, p, \theta)$. It can be expressed as:

$$\begin{aligned} CE(x, p, \theta) &= u_\psi^{-1}(U(x, p, \theta)) \\ &= EV(x, \theta) + \hat{y} - p + u_\psi^{-1}(U(x, p, \theta)) - EV(x, \theta) - \hat{y} + p \\ &= EV(x, \theta) + \hat{y} - p - RP(x, p, \theta), \end{aligned}$$

where $EV(x, \theta) + \hat{y} - p$ is the expected payoff and $RP(x, p, \theta)$ is the risk premium associated with the lottery. In particular,

$$\begin{aligned} EV(x, \theta) &= \mathbb{E}_l[\ b^*(l, \omega, x) - c_x^*(l, \omega, x)\ ] \\ &= \mathbb{E}_l[\ b^*(l, \omega, x_0) - c_x^*(l, \omega, x_0) + v(l, \omega, x)\ ], \text{ and} \\ RP(x, p, \theta) &= EV(x, \theta) + \hat{y} - p - u_\psi^{-1}(U(x, p, \theta)), \end{aligned} \tag{A.1}$$

where as before $v(l, \omega, x) = b^*(l, \omega, x) - b^*(l, \omega, x_0) - \big(c_x^*(l, \omega, x) - c_x^*(l, \omega, x_0)\big)$. A consumer's willingness to pay for contract x relative to the null contract x_0 is equal to $\tilde{p}$ that solves:

$$\begin{aligned} CE(x, \tilde{p}, \theta) &= CE(x_0, 0, \theta) \\ EV(x, \theta) + \hat{y} - \tilde{p} - RP(x, \tilde{p}, \theta) &= EV(x_0, \theta) + \hat{y} - RP(x_0, 0, \theta) \\ \tilde{p} &= EV(x, \theta) - EV(x_0, \theta) + RP(x_0, 0, \theta) - RP(x, \tilde{p}, \theta). \end{aligned}$$

To obtain a closed-form expression for willingness to pay, we assume constant absolute risk aversion, and thus that the risk premium RP does not depend on residual income.[1] In this case, marginal willingness to pay for contract x relative to the null contract is given by:

$$\begin{aligned} WTP(x,\theta) &= EV(x,\theta) - EV(x_0,\theta) + RP(x_0,\theta) - RP(x,\theta) \\ &= \mathbb{E}_l[\ c^*_{x_0}(l,\omega,x_0) - c^*_x(l,\omega,x_0) + v(l,\omega,x)\ ] + \Psi(x,\theta), \end{aligned}$$

where $\Psi(x,\theta) = RP(x_0,\theta) - RP(x,\theta)$. If the null contract provides a riskier distribution of payoffs than contract x, $\Psi(x,\theta)$ will be positive.

A.2 Definitions and Proofs

Assumptions. Consider the model in Section II.A. Suppose contracts $x \in X$ are characterized by increasing, continuous, and concave out-of-pocket cost functions $c_x : \mathbb{R}_+ \to \mathbb{R}_+$, where $c_x(m) \leq m\ \forall\ m$ and which are differentiable almost everywhere with $c'_x \in [0,1]$, where c'_x denotes the derivative wherever it exists. Suppose consumers have type $\theta = (F,\omega,\psi) \in \Delta^c(\mathbb{R}) \times \mathbb{R}_{++} \times \mathbb{R}_{++} =: \Theta$.[2] Given health state realization $l \in \mathbb{R}$, contract premium p, and initial income $\hat{y}$, suppose consumers value healthcare spending $m \in \mathbb{R}_+$ according to $u_\psi(\hat{y} - p + b(m;l,\omega) - c_x(m))$, where $b(m;l,\omega) = (m-l) - \frac{1}{2\omega}(m-l)^2$ and where $u_\psi(x) = -\exp(-\psi x)$.

Under these assumptions, social surplus is given by $SS(x,\theta) = \Psi(x,\theta) - \textsc{scmh}(x,\theta)$, where

$$\begin{aligned} \Psi(x,\theta) &= RP(x_0,\theta) - RP(x,\theta) \\ &\qquad \text{where}\ \ RP(x,\theta) = \psi^{-1}\log\Big(\mathop{\mathbb{E}}_{l\sim F}\big[\exp\big(-\psi(z_x(l,\theta) - \bar{z}_x(\theta))\big)\big]\Big), \\ \textsc{scmh}(x,\theta) &= \mathop{\mathbb{E}}_{l\sim F}\Big[\tfrac{\omega}{2}\big(1 - c'_x(m^*(l,\omega,x))\big)^2\Big], \end{aligned}$$

and where $z_x(l,\theta) = \hat{y} - p + b(m^*(l,\omega,x);l,\omega) - c_x(m^*(l,\omega,x))$ and $\bar{z}_x(\theta) = \mathbb{E}_{l\sim F}[z_x(l,\theta)]$. Appendix C.2 solves for privately optimal spending $m^*(l,\omega,x) = \text{argmax}_m\,(b(m;l,\omega) - c_x(m))$ when contracts are piecewise linear with a deductible, coinsurance rate, and out-of-pocket maximum. As m^* never falls on a kink, $c'_x(m^*)$ always exists. The indirect benefit from privately optimal spending is given by $b(m^*(l,\omega,x);l,\omega) = \frac{\omega}{2}(1 - c'_x(m^*(l,\omega,x))^2)$. Willingness to pay is given by $WTP(x,\theta) = \bar{z}_x(\theta) - \bar{z}_{x_0}(\theta) + \Psi(x,\theta)$.

[1] In Equation A.1, $\hat{y} - p$ cancels out completely. This assumption is most reasonable when marginal premiums between relevant plans are small relative to initial income.

[2] $\Delta^c(\mathbb{R})$ denotes the the set of continuous probability measures on the Borel σ-algebra of $\mathbb{R}$.

Definitions. We say that a given contract is "higher coverage" than another if it provides both a higher certainty equivalent payoff $WTP(x, \theta)$ as well as greater risk protection $\Psi(x, \theta)$. This notion of coverage level is slightly stronger that what is implied by vertical differentiation alone. We use it because it allows our model to have the following desirable properties:

(i) the value of risk protection is increasing in coverage level;

(ii) the social cost of moral hazard is increasing in coverage level;

(iii) efficient coverage level is increasing in risk aversion;

(iv) efficient coverage level is decreasing in the moral hazard parameter.

Definitions 1 and 2 formalize the distinction between vertical differentiation and coverage level ordering. Propositions 1 and 2 provide the conditions on contracts that yield each ordering. Briefly, vertical differentiation requires only a relation on contracts' *level* of out-of-pocket costs, while coverage level ordering (as defined) also requires a relation on contracts' *marginal* out-of-pocket costs. A higher-coverage contract must have an out-of-pocket cost function that is everywhere below and everywhere flatter than a lower-coverage contract.

Implications. The most important reason we use this notion of coverage level is that it allows extrapolation of social surplus across coverage levels. Namely, it implies that social surplus is single peaked in coverage level. Proposition 3 states this formally. Single-peakedness allows one to infer, for example, that if a given contract is less-than-socially-optimal coverage for all households, the same would be true of any lower level of coverage.

Proofs are provided below. Of the four stated properties of the model, property (i) is true by definition, property (ii) is established in the proof of Proposition 3, and properties (iii) and (iv) are proved in Lemmas 2 and 3, respectively.

Definition 1. Contracts $x', x \in X$ are *vertically differentiated* (with x' preferred) if and only if $WTP(x', \theta) \geq WTP(x, \theta) \ \forall \ \theta \in \Theta$.

Definition 2. Given $x', x \in X$, contract x' is *higher coverage* than contract x if and only if $WTP(x', \theta) \geq WTP(x, \theta) \ \forall \ \theta \in \Theta$ and $\Psi(x', \theta) \geq \Psi(x, \theta) \ \forall \ \theta \in \Theta$. We denote this relationship by writing $x' \geq x$.

Proposition 1. *Contracts $x', x \in X$ are vertically differentiated (with x' preferred) if and only if $c_{x'}(m) \leq c_x(m) \ \forall \ m$.*

Proposition 2. *Given $x', x \in X$, contract x' is higher coverage than contract x if and only if $c_{x'}(m) \leq c_x(m) \ \forall \ m$ and $c'_{x'}(m) \leq c'_x(m)$ almost everywhere.*

Proposition 3. *Social surplus is single peaked in coverage level. That is, fixing $\theta \in \Theta$ and $x, x', x'' \in X$ where $x \leq x' \leq x''$: if $SS(x'', \theta) \geq SS(x', \theta)$, then $SS(x', \theta) \geq SS(x, \theta)$.*

Proof of Proposition 1: Contracts $x', x \in X$ are vertically differentiated (with x' preferred) if and only if $c_{x'}(m) \leq c_x(m)\ \forall\ m$.

Fix $\theta \in \Theta$. Let $Z_x =: z_x(l, \theta)$ be the random payoff induced by health state distribution F. At any health state l, lower out-of-pocket costs deliver higher payoffs:

$$\begin{aligned} Z_x = z_x(l, \theta) &= \hat{y} - p + b(m^*(l, \omega, x); l, \omega) - c_x(m^*(l, \omega, x)) \\ &\leq \hat{y} - p + b(m^*(l, \omega, x); l, \omega) - c_{x'}(m^*(l, \omega, x)) \\ &\leq \hat{y} - p + b(m^*(l, \omega, x'); l, \omega) - c_{x'}(m^*(l, \omega, x')) = z_{x'}(l, \theta) = Z_{x'}\ , \end{aligned}$$

where the second inequality holds by the optimality of $m^*(l, \omega, x')$. [$\Leftarrow$] $Z_{x'}$ therefore first-order stochastically dominates Z_x, and the result follows because u_ψ is increasing. [$\Rightarrow$] If $c_{x'}(\tilde{m}) > c_x(\tilde{m})$ for some $\tilde{m}$, the first inequality fails to hold for consumer type $\tilde{\omega}$ at health state realization $\tilde{l}$ at which $m^*(\tilde{l}, \tilde{\omega}, x) = \tilde{m}$. Such a consumer type exists for any $\tilde{m}$ we might choose because as ω approaches zero, privately optimal utilization approaches the health state, meaning any m can be approached arbitrarily closely. As c_x is continuous, if $c_{x'}(\tilde{m}) > c_x(\tilde{m})$, the same will be true in a neighborhood of $\tilde{m}$. A consumer with health state distribution $\tilde{F}$ degenerate on $\tilde{l}$ would strictly prefer contract x. By continuity, a consumer with a health state distribution that is sufficiently concentrated at $\tilde{l}$ would also prefer contract x. □

Proof of Proposition 2: Contract x' is higher coverage than contract x if and only if $c_{x'}(m) \leq c_x(m)\ \forall\ m$ and $c'_{x'}(m) \leq c'_x(m)$ almost everywhere.

By Proposition 1, $c_{x'}(m) \leq c_x(m)\ \forall\ m$ is necessary and sufficient for the contracts to be vertically differentiated. It remains to show that $c'_{x'}(m) \leq c'_x(m)$ almost everywhere is necessary and sufficient for $\Psi(x', \theta) \geq \Psi(x, \theta)$. Fix $\theta \in \Theta$. Let $\dot{Z}_x =: z_x(l, \theta) - \bar{z}_x(\theta)$ be the mean-zero random payoff induced by health state distribution F. Differentiating $\dot{Z}_x$ with respect to the health state realization l:

$$\begin{aligned} \frac{d\dot{Z}_x}{dl} &= \frac{\partial b}{\partial l}(m^*(l, \omega, x); l, \omega) \\ &\leq \frac{\partial b}{\partial l}(m^*(l, \omega, x'); l, \omega) = \frac{d\dot{Z}_{x'}}{dl} \\ &\leq 0. \end{aligned}$$

That is, the payoff is weakly decreasing in the health state, and is doing so faster for contract x than for contract x'. The first equality holds by the envelope theorem. Because $\frac{\partial^2 b}{\partial l \partial m} = \omega^{-1} \geq 0$, the first inequality holds as long as $m^*(l, \omega, x)$ is increasing in x. The second inequality holds because $\frac{\partial b}{\partial l} = \omega^{-1}(m^*(l, \omega, x) - l) - 1 \leq 0$, or in other words, because moral hazard spending never exceeds ω.[3] [$\Leftarrow$] Lemma 1 shows that $m^*(l, \omega, x)$ is increasing in x as long as $c'_{x'}(m) \leq c'_x(m)$. $\dot{Z}_x$ is therefore a mean preserving spread of $\dot{Z}_{x'}$, and would be preferred by any risk-averse expected utility maximizer: $\mathbb{E}_{l \sim F}[u_\psi(\dot{Z}_{x'})] \geq \mathbb{E}_{l \sim F}[u_\psi(\dot{Z}_x)]$. The result follows because $-\psi^{-1} \log(-x)$ is increasing. [$\Rightarrow$] If $c_{x'}(\tilde{m}) > c_x(\tilde{m})$ for some $\tilde{m}$, the first inequality fails to hold for consumer type $\tilde{\omega}$ at health state realization $\tilde{l}$ at which $m^*(\tilde{l}, \tilde{\omega}, x) = \tilde{m}$. Such a consumer type exists for any $\tilde{m}$ we might choose because as ω approaches zero, privately optimal utilization approaches the health state, meaning any m can be approached arbitrarily closely. As c_x is continuous, if $c_{x'}(\tilde{m}) > c_x(\tilde{m})$, the same will be true in a neighborhood of $\tilde{m}$. At $\tilde{l}$, the payoff would therefore be decreasing faster in the health state under contract x' than under contract x, and x would provide strictly more risk protection to a consumer with health state distribution $\tilde{F}$ sufficiently concentrated around $\tilde{l}$. □

Proof of Proposition 3. Social surplus is single peaked in coverage level. That is, fixing $\theta \in \Theta$ and $x, x', x'' \in X$ where $x \leq x' \leq x''$: if $SS(x'', \theta) \geq SS(x', \theta)$, then $SS(x', \theta) \geq SS(x, \theta)$.

Let $\tilde{c}_x(l) = c_x(m^*(l, \omega, x))$ be the indirect out-of-pocket cost function for consumer type ω under contract x.[4] As θ is fixed throughout the proof, we omit ω as an argument in $\tilde{c}_x(l)$. Similarly, let $\tilde{c}'_x(l) = c'_x(m^*(l, \omega, x))$ be the indirect marginal out-of-pocket cost function. Note that because $m^*(l, \omega, x)$ is increasing in x (see Lemma 1) and contracts are concave, $\tilde{c}'_{x''}(l) \leq \tilde{c}'_{x'}(l) \leq \tilde{c}'_x(l)$ wherever these derivatives exist.

Next, for each contract $x \in \{x, x', x''\}$, calculate the cutoff values of the health state l that determine which segment of the piecewise linear out-of-pocket cost function the consumer of type θ will choose. Appendix C.2 describes this procedure and provides formulas for the cutoffs. As the contracts we consider have at most three segments, each contract has at most three cutoffs: one at which positive healthcare utilization begins and two separating the segments of the out-of-pocket cost function.[5] Considering the three cutoff values of our three candidate

[3] Note that this statement would not be true under the "multiplicative" specification of preferences proposed by Einav et al. (2013) and used in Ho and Lee (2021). In that case, $\frac{\partial b}{\partial l}$ becomes positive at a certain health state level, and the payoff $z_x(l, \theta)$ begins increasing in the health state. The conditions given in Proposition 2 would therefore not be sufficient to guarantee coverage level ordering in that context.

[4] The line labelled c^* in Figure A.2 represents the function $\tilde{c}_x(l)$ in that example.

[5] The proof extends trivially to piece-wise linear out-of-pocket functions with a different number of segments.

contracts simultaneously, the space of health states (the real line) is divided into at most 10 regions. Denote these regions by $\{R_r\}_{r=1}^{10}$, where $R_r = (l_r^{lb}, l_r^{ub})$ and $l_r^{ub} = l_{r+1}^{lb}$.[6] The lower bound of the first region is $-\infty$ and the upper bound of the final region is ∞. For each contract x in each region R_r, out-of-pocket costs are linear in the health state, and so can be written $\tilde{c}_{x,r}(l) = \gamma_{x,r} + l\,\tilde{c}'_{x,r}$, with intercept $\gamma_{x,r}$ and slope $\tilde{c}'_{x,r}$. As before, higher coverage contracts are flatter: $c'_{x'',r} \leq c'_{x',r} \leq c'_{x,r}\ \forall\ r$.

Extend this notation to the consumer's payoff $z_x(l,\theta)$. Omitting θ, the payoff in region r under contract x can now be written:

$$\begin{aligned} z_x(l) &= \hat{y} - p_x + \tfrac{\omega}{2}(1 - \tilde{c}'_x(l)^2) - \tilde{c}_x(l) \\ &= \hat{y} - p_x + \tfrac{\omega}{2}(1 - \tilde{c}'^{\,2}_{x,r}) - \gamma_{x,r} - \tilde{c}'_{x,r}\, l, \quad l \in R_r. \end{aligned}$$

The payoff is linear in the health state with slope and intercept determined by the relevant segment of the indirect out-of-pocket cost function. To isolate the effects of level from the effects of slope, it is useful to express the payoff in terms of differences from its mean in a given region. To this end, write:

$$z_x(l) = \bar{z}_{x,r} - \tilde{c}'_{x,r}(l - \bar{l}_r), \quad l \in R_r$$

where $\bar{l}_r = \mathbb{E}_{l|R_r}[l]$ is the conditional expectation of the health state in region r with respect to the consumer's health state distribution F, and $\bar{z}_{x,r} = z_x(\bar{l}_r)$ is the conditional expectation of the payoff. Note that because higher coverage contracts deliver everywhere higher payoffs (see proof of Proposition 1): $\bar{z}_{x'',r} \geq \bar{z}_{x',r} \geq \bar{z}_{x,r}\ \forall\ r$. Each contract is now fully characterized by the payoff function it generates, which in turn is fully described by its mean and slope in each region: $\{\bar{z}_{x,r}, \tilde{c}'_{x,r}\}_{r=1}^{10}$. Higher coverage contracts generate both higher and flatter payoffs in every region. Expressing the payoff function in this way allows us think about changing a contract's slope while holding its expected payoff fixed, and vice versa.

We now proceed in two steps. We first show that the social cost of moral hazard $\textsc{scmh}(x,\theta)$ is increasing and "convex" in coverage level. As coverage level itself has no cardinal interpretation, the idea of convexity is applicable with respect to the slope of contracts' indirect out-of-pocket cost functions $\tilde{c}'_{x,r}$. We then show that the value of risk protection $\Psi(x,\theta)$ is increasing and "concave" in coverage level, where the idea of concavity is again applicable with respect to $\tilde{c}'_{x,r}$. Note that the tradeoff between risk protection and moral hazard operates entirely through the slope of the out-of-pocket cost function. The level of out-of-pocket costs impacts only the value of risk protection, and does so monotonically. As $SS(x,\theta) = \Psi(x,\theta) - \textsc{scmh}(x,\theta)$, these

[6] As we have assumed F is continuously distributed, there is zero mass on region boundaries.

two steps imply $SS(x,\theta)$ is itself concave in the slope of the out-of-pocket function. Single-peakedness in coverage level follows from the fact that this slope is monotonic in coverage level.

1. $\textsc{scmh}(x,\theta)$ *is increasing and "convex" in coverage level.*

First, split the expectation between the defined regions, omitting θ as an argument:

$$\begin{aligned}\textsc{scmh}(x) &= \mathop{\mathbb{E}}_{l\sim F}\left[\tfrac{\omega}{2}(1-\tilde{c}'_x(l))^2\right] \\ &= \textstyle\sum_{r=1}^{10}\pi_r\left[\tfrac{\omega}{2}(1-\tilde{c}'_{x,r})^2\right],\end{aligned}$$

where $\pi_r = Pr(l \in R_r | l \sim F)$ is the probability of realizing a health state in region R_r. Taking the derivative with respect to the slope of the indirect out-of-pocket cost function in a given region:

$$\frac{d\textsc{scmh}(x)}{d\tilde{c}'_{x,r}} = -\pi_r\omega(1-\tilde{c}'_{x,r}) \leq 0.$$

As $\textsc{scmh}(x)$ is decreasing in $\tilde{c}'_{x,r}$ in any region, it is increasing in coverage level. Taking the second derivative:

$$\frac{d^2\textsc{scmh}(x)}{d\tilde{c}'^{\,2}_{x,r}} = \pi_r\omega \geq 0.$$

The social cost of moral hazard is therefore increasing in the slope of the indirect out-of-pocket cost function $\tilde{c}'_{x,r}$ at an increasing rate. It is unaffected by changes in $\bar{z}_{x,r}$.

2. $\Psi(x,\theta)$ *is increasing and "concave" in coverage level.*

First, split the expectation between the defined regions, omitting θ as an argument:

$$\begin{aligned}\Psi(x) &= RP(x_0) - \psi^{-1}\log\left(\mathbb{E}_l\left[\exp\left(-\psi(z_x(l)-\bar{z}_x)\right)\right]\right) \\ &= RP(x_0) - \psi^{-1}\log\left(\textstyle\sum_{r=1}^{10}\pi_r\mathbb{E}_{l|R_r}\left[\exp\left(-\psi(z_x(l)-\bar{z}_x)\right)\right]\right),\end{aligned}$$

where $\pi_r = Pr(l \in R_r | l \sim F)$ is the probability of realizing a health state in region R_r. Taking the derivative with respect to the slope of the indirect out-of-pocket cost function in a given region:

$$\frac{d\Psi(x)}{d\tilde{c}'_{x,r}} = \left(\mathbb{E}_l[\exp(-\psi z_x(l))]\right)^{-1}\pi_r\mathbb{E}_{l|R_r}[\exp(-\psi z_x(l))(\bar{l}_r - l)] \leq 0.$$

Because the function $\exp(-\psi x)$ is convex and the payoffs $z_x(l)$ are decreasing in the health state, worse-than-average health states ($l \geq \bar{l}_r$) receive more weight than better-than-average health states ($l \leq \bar{l}_r$), and the expression is nonpositive. Taking the second derivative:

$$\frac{d^2\Psi(x)}{d\tilde{c}'_{x,r}{}^2} = \psi \left[\left(\frac{\pi_r \, \mathbb{E}_{l|R_r}[\exp(-\psi z_x(l))(\bar{l}_r - l)]}{\mathbb{E}_l[\exp(-\psi z_x(l))]} \right)^2 - \left(\frac{\pi_r \, \mathbb{E}_{l|R_r}[\exp(-\psi z_x(l))(\bar{l}_r - l)^2]}{\mathbb{E}_l[\exp(-\psi z_x(l))]} \right) \right] \leq 0.$$

The first term is the squared conditional expectation of $(\bar{l}_r - l)$. The second term is the conditional expectation of $(\bar{l}_r - l)^2$. Because x^2 is convex, the expression is nonpositive by Jensen's inequality.

□

Lemma 1. *Healthcare utilization is increasing in coverage level.*

Proof. Fix $l \in \mathbb{R}$, $\omega \in \mathbb{R}_{++}$, and $x, x' \in X$ where $x \leq x'$. Optimal utilization $m^*(l, \omega, x) = \operatorname{argmax}_m (b(m; l, \omega) - c_x(m))$. Consider $m, m' \in \mathbb{R}_+$ where $m \leq m'$:

$$\begin{aligned} b(m'; l, \omega) - c_{x'}(m') - [b(m'; l, \omega) - c_x(m')] &= c_x(m') - c_{x'}(m') \\ &\geq c_x(m) - c_{x'}(m) \\ &= b(m; l, \omega) - c_{x'}(m) - [b(m; l, \omega) - c_x(m)], \end{aligned}$$

where the inequality holds because $c_{x'}(m) \leq c_x(m)$ and $c'_{x'}(m) \leq c'_x(m)$ guarantees c is submodular in m and x. The objective $b(m; l, \omega) - c_x(m)$ is therefore supermodular and standard monotone comparative statics imply $m^*(l, \omega, x)$ is increasing in x.

□

Lemma 2. *Efficient coverage level is increasing in risk aversion.*

Proof. Fix $\theta \in \Theta$. Efficient coverage level $x^{eff} = \operatorname{argmax}_x (RP(x_0, F, \omega, \psi) - RP(x, F, \omega, \psi) - \textsc{scmh}(x, F, \omega))$. As the insurer is risk-neutral, the social cost of moral hazard is unaffected by ψ. Differentiating $RP(x, F, \omega, \psi)$ with respect to ψ:

$$\frac{dRP(x, \theta)}{d\psi} = -\psi^{-1} \left[RP(x, \theta) + (\mathop{\mathbb{E}}_{l \sim F}[\exp(-\psi \dot{Z}_x)])^{-1} \mathop{\mathbb{E}}_{l \sim F}[\exp(-\psi \dot{Z}_x) \dot{Z}_x] \right],$$

where $\dot{Z}_x =: z_x(l, \theta) - \bar{z}_x(\theta)$. The first term in the brackets, $RP(x, \theta)$, is shown to be decreasing in x in Proposition 2. The second term represents a weighted average of deviations from mean payoffs, where the weights correspond to the utility weight at that realization. As $\dot{Z}_x$ becomes less risky as x increases (see proof of Proposition 2), this term is also decreasing in x. $\frac{dSS(x,\theta)}{d\psi}$ is therefore increasing in x, and standard monotone comparative statics imply x^{eff} is increasing in ψ.

□

Lemma 3. *Efficient coverage level is decreasing in the moral hazard parameter.*

Proof. Fix $\theta \in \Theta$. Efficient coverage level $x^{eff} = \text{argmax}_x (\Psi(x,\theta) - \textsc{scmh}(x,\theta))$, where $\textsc{scmh}(x,\theta) = \mathbb{E}_{l\sim F}[\frac{\omega}{2}(1 - c'_x(m^*(l,\omega,x)))^2]$. Differentiating $\textsc{scmh}(x,\theta)$ with respect to ω:

$$\frac{d\textsc{scmh}(x,\theta)}{d\omega} = \mathop{\mathbb{E}}_{l\sim F}[\tfrac{1}{2}(1 - c'_x(m^*(l,\omega,x)))^2] \leq 0.$$

Note that contracts are piecewise linear and $c'_x \in [0,1]$. Because $m^*(l,\omega,x)$ is increasing in x (see Lemma 1) and contracts are concave, $c'_x(m^*(l,\omega,x))$ is decreasing in x and $\frac{dSCMH(x,\theta)}{d\omega}$ is increasing in x. $\frac{dSS(x,\theta)}{d\omega}$ is therefore decreasing in x, and standard monotone comparative statics imply x^{eff} is decreasing in ω. □

Appendix B Additional Analysis

B.1 Estimation of Plan Cost-sharing Features

A crucial input to our empirical model is the cost-sharing function of each plan. While Table 1 describes plans using the deductible and in-network out-of-pocket maximum, plans are in reality characterized by a much more complex set of payment rules, including copayments, specialist visit coinsurance, out-of-network fees, and fixed charges for emergency room visits. To structurally model moral hazard, we make the huge simplification that healthcare is a homogenous good over which the consumer chooses only the quantity to consume. We then model this decision as being based in part on out-of-pocket cost. To that end, our empirical model requires as an input a univariate function that maps total healthcare spending into out-of-pocket cost.

A natural choice might be to use the deductible, nonspecialist coinsurance rate, and in-network out-of-pocket maximum. However, in our setting, the out-of-pocket cost function described by these features does not correspond well to what we observe in the claims data. In particular, we often observe out-of-pocket spending amounts that exceed plans' in-network out-of-pocket maximum. Because of this, we take a different approach.

We define plan cost-sharing functions by three parameters: a deductible, a coinsurance rate, and an out-of-pocket maximum. Taking the true deductibles as given (since these correspond well to the data), we estimate a coinsurance rate and an out-of-pocket maximum that minimizes the sum of squared residuals between predicted and observed out-of-pocket cost. We observe realized total healthcare spending for each household in the claims data. Predicted out-of-pocket cost is calculated by applying the deductible and supposed coinsurance rate and out-of-pocket maximum. "Observed" out-of-pocket cost is either observed directly in the claims data (if a household chose that plan) or else calculated counterfactually. We carry out this procedure separately for each plan, year, and family status (individual or family).[7] Figure A.1 shows an example of the data and estimates for a particular plan: Moda - 3 for individual households in 2012. Table A.3 presents the estimated cost-sharing features for all plans in all years.

B.2 Variation in Plan Menu Generosity

Measuring Plan Menu Generosity. We want to measure the likelihood that a household would choose generous health insurance coverage when presented with a particular plan menu. We refer to this measure as "plan menu generosity." At a simple level, if plan menus consisted of only a single plan, the assignment to higher coverage would obviously constitute a "more generous menu" than the assignment to lower coverage. But plan menus in our setting are more complex. They contain multiple plans and many possible permutations of plan choice sets, and plans vary by their actuarial value, the identity of their insurer, their associated employee premium, and their potential HSA/HRA and vision/dental contribution. All of these factors likely influence households' plan choices.

In order to construct usable measures of plan menu generosity, we transform these multi-dimensional objects using a conditional logit model that excludes all household observables. This specification allows us to predict the probability that a given household would choose a given plan when presented with a given plan menu *as if* the household had been acting like the average household in the data. Variation in the resulting predicted choice probabilities is driven only by variation in plan menus, and not by variation in (observed or unobserved) household characteristics.

Abstracting from the dimension of time for now, we define $plan_{jk}$ as an indicator for the plan j chosen by household k. We estimate the following conditional logit model:

$$plan_{jk} = \underset{j \in \mathcal{J}_d}{\text{argmax}}(\alpha p_{jd} + \alpha^{VD} p_{jd}^{VD} + \alpha^{HA} p_{jd}^{HA} + \nu_j + \epsilon_{jk}), \tag{B.1}$$

[7] So that the cost-sharing estimates are not affected by large outliers, we drop observations where out-of-pocket spending was above \$20,000 or total healthcare spending was above \$100,000.

where $\mathcal{J}_d$ is the set of plans available in the school district-family type-occupation type combination d (to which household k belongs), p_{jd} is the employee premium, p_{jd}^{VD} is the vision/dental subsidy, and p_{jd}^{HA} is the HSA/HRA contribution. Plan characteristics are captured nonparametrically by plan fixed effects ν_j. All household-specific determinants of plan choice are contained in the error term ϵ_{jk}. Estimated parameters are presented in Table A.4, separately for each year of our data. As expected, households dislike premiums, prefer higher HSA/HRA and vision/dental subsidies, and prefer higher-coverage plans to lower-coverage plans.

We use the choice probabilities implied by Equation B.1 to construct our measures of plan menu generosity. Given plan menu $\mathbf{menu}_d \equiv \{p_{jd}, p_{jd}^{VD}, p_{jd}^{HA}, \nu_j\}_{j \in \mathcal{J}_d}$, we denote the predicted probability that plan j is chosen as ρ_{jd}.[8] Our measures of plan menu generosity are the probability a household would choose a given insurer and the expected actuarial value of a household's plan choice conditional on insurer, respectively given by:

$$\begin{aligned} \rho_{fd} &= \sum_{j \in \mathcal{J}_d^f} \rho_{jd}, \\ \widehat{AV}_{fd} &= \sum_{j \in \mathcal{J}_d^f} (\frac{\rho_{jd}}{\rho_{fd}}) AV_j, \end{aligned} \tag{B.2}$$

where $\mathcal{J}_d^f$ is the set of plans in $\mathbf{menu}_d$ offered by insurer f.

Explaining Plan Menu Generosity. Because the majority of the variation in coverage level lies within Moda, we focus on explaining plan menu generosity using the predicted actuarial value among Moda plans. We first compare plan menu generosity to observed household health (see Table A.5). We can in all years reject the hypothesis that household risk scores are correlated with plan menu generosity, conditional on family structure. We also find that plan menus are consistently most generous for single employee coverage and least generous for employee plus family coverage. This pattern is consistent with our understanding of OEBB's benefit structure, and is common in employer-sponsored health insurance.

We further explore which covariates, in addition to family structure, can explain variation in plan menu generosity. Table A.6 presents three additional regressions of predicted actuarial value on employee-level covariates (part-time versus full-time status, occupation type, and union affiliation), as well as on school district-level covariates (home price index and percent of Republicans).[9] Employees are either part-time or full-time. There are eight mutually exclusive

[8] Formally: $\rho_{jd} = \frac{\exp(U_{jd})}{\sum_{g \in \mathcal{J}_d} \exp(U_{gd})}$, where $U_{jd} = \alpha p_{jd} + \alpha^{VD} p_{jd}^{VD} + \alpha^{HA} p_{jd}^{HA} + \nu^j$.

[9] Possible employee occupation types are licensed administrator, non-licensed administrator, classified, community college non-instructional, community college faculty, confidential, licensed, substitute, and superintendent.

employee occupation types; the regressions omit the type "Licensed Administrator." There are five mutually exclusive union affiliations, and employees may not be affiliated with a union; the regressions omit the non-union category. We calculate the average home price index (HPI) in a school district by taking the average zip-code level home price index across employees' zip-code of residence.[10] *Pct. Republican* measures the percent of households in a school district that are registered as Republicans as of 2016.[11]

We find that plan menus are less generous for part-time employees, are substantially less generous for substitute teachers, and are more generous for employees at community colleges. Certain union affiliations are also predictive of more or less generous plan menus. Across school districts, plan menu generosity is decreasing in both the logged home price index and the percent of registered Republicans.

B.3 Reduced-form Estimates of Moral Hazard

While our primary sample consists of data from 2009–2013, we conduct our reduced-form analysis of moral hazard using only data from 2008.[12] The OEBB marketplace began operating in 2008, so that year all employees chose from this set of plans for the first time. This "active choice" year permits us to look cleanly at how plan choices and healthcare spending depended on plan menus without also having to account for how prior-year plan menus affected current-year plan choices. While our structural model will capture these dynamics, we feel they are better avoided at this stage.

We estimate how plan menus—choice sets and prices—affect plan choices, and in turn how

"Licensed" refers to the possession of a teaching license. Within each type, an employee can be either full-time or part-time. Possible family types are employee only; employee and spouse; employee and child(ren); and employee, spouse, and child(ren).

[10] We use 5-digit zip-code-level home price indices from Bogin, Doerner and Larson (2019). The data and paper are accessible at http://www.fhfa.gov/papers/wp1601.aspx.

[11] Data on percent of registered voters by party is available at the county level; we construct school-district-level measures by taking the average over employees' county of residence. Voter registration data in Oregon can be downloaded at https://data.oregon.gov/api/views/6a4f-ecbi.

[12] The cost-sharing features of 2008 plans are presented in Table A.1; they are very similar to the plans offered in 2009. We apply the same sample construction criteria to our 2008 sample, except that households must be present for one prior year.

plan choices affect total healthcare spending, as described by Equations (B.3) and (B.4):

$$plan_k = f(\mathbf{menu}_d, \mathbf{X}_k, \xi_k), \tag{B.3}$$

$$y_k = g(plan_k, \mathbf{X}_k, \xi_k). \tag{B.4}$$

Here, $plan_k$ represents the plan chosen by household k, $\mathbf{menu}_d$ represents the plan menu available to the school district-family type-occupation type combination d (to which household k belongs), $\mathbf{X}_k$ are observable household characteristics, ξ_k are unobservable household characteristics, and y_k is total healthcare spending. Because household characteristics appear in both equations, the standard challenge in estimating the effect of $plan_k$ on y_k is that a household's chosen plan is correlated with its unobservable characteristics ξ_k. Our identifying assumption is that plan menus are independent of household unobservables ξ_k conditional on household observables $\mathbf{X}_k$.

We parameterize $plan_k$ to be an indicator variable for the identity of the insurer and a continuous variable for the plan actuarial value. We then parameterize Equation B.4 according to

$$\log(y_k) = \delta_f \mathbf{1}_{f(k)=f} + \gamma \log(1 - AV_{j(k)}) \mathbf{1}_{f(k)=Moda} + \beta \mathbf{X}_k + \xi_k, \tag{B.5}$$

where $\mathbf{1}_{f(k)=f}$ is an indicator for the insurer chosen by household k and $AV_{j(k)}$ is the actuarial value of the plan chosen by household k. The parameter δ_f represents insurer-specific treatment effects on total spending.[13] Our parameter of interest is γ, which represents the responsiveness of total spending to plan generosity, holding the insurer fixed (at Moda).[14] We follow the literature in formulating the model so that γ represents the elasticity of total healthcare spending with respect to the average out-of-pocket price per dollar of total spending.[15]

We estimate Equation B.5 using two-stage least squares, instrumenting for the chosen insurer ($\mathbf{1}_{f(k)=f}$) and actuarial value ($AV_{j(k)}$) using $\mathbf{menu}_d$. As instruments, we use the measures of plan menu generosity constructed in Appendix B.2. Namely, we instrument for $\mathbf{1}_{f(k)=f}$ using using ρ_{fd} and for $\log(1 - AV_{j(k)}) \mathbf{1}_{f(k)=Moda}$ using $\log(1 - \widehat{AV}_{d,Moda}) \rho_{d,Moda}$. Table A.7 reports the estimates. We report only the coefficient of interest (γ), but all specifications also contain insurer fixed effects, as well as controls for household risk score and family structure. The

[13] These may arise due to "supply side" effects arising from differences in provider prices, provider networks, or care management practices, or due to "demand side" effects from differences in average plan generosity.

[14] We do not try to estimate a moral hazard elasticity among the plans offered by Kaiser and Providence because there is so little variation in coverage level.

[15] To accommodate the fact that 2 percent of households have zero spending, we add 1 to total spending.

first column presents the parameters estimated without instruments, and the second column presents the instrumental variables estimates. Comparing the coefficients in columns 1 and 2, we find that moral hazard explains 46 percent of the observed relationship between plan generosity and total healthcare spending. Our overall estimate of the elasticity of demand for healthcare spending in the population is -0.27. The standard benchmark estimate from the RAND health insurance experiment is -0.2 (Manning et al., 1987; Newhouse, 1993).

Heterogeneity. Columns 3 and 4 of Table A.7 introduce heterogeneity in γ by household health. For each household type (individual or family), we classify households into quartiles based on household risk score, where Q_n denotes the quartile of risk (Q_4 is highest risk). We construct separate instruments for each of the eight household types by estimating the logit model in Equation B.1 for only that subsample of households. We find noisy but large differences in γ across household risk quartiles and between individual and family households.

Variation in γ could reflect either heterogeneity in the intensity of treatment (extent of exposure to varying marginal prices of healthcare across plans), or heterogeneity in treatment effect (different responsiveness to varying marginal prices of healthcare across plans), or both. While this analysis cannot distinguish between these two effects, we find suggestive evidence that the heterogeneity at least in part reflects differential treatment intensity. The remainder of this section presents an analysis that compares the realized spending outcomes of households in different risk quartiles with the variation in plan cost-sharing features that gives rise to different end-of-year marginal out-of-pocket prices. We find that the household types for which we estimate higher γ are also more likely to be exposed to varying marginal out-of-pocket costs. Distinguishing variation in treatment intensity from variation in treatment effect is an important advantage of our structural model.

Appendix C Estimation Details

C.1 Fenton-Wilkinson Approximation

Because there is no closed-form solution for the distribution of the sum of lognormal random variables, the Fenton-Wilkinson approximation is widely used in practice.[16] Under this approximation, the distribution of the sum of draws from independent lognormal distributions can be represented by a lognormal distribution. The parameters of the approximating distribution

[16] See Fenton (1960), and for a summary, Cobb, Rumí and Salmerón (2012).

are chosen such that its first and second moments match the corresponding moments of the true distribution of the sum of lognormals. In our application, the sum of lognormals is the household's health state distribution, and the lognormals being summed are the individuals' health state distributions. An individual's health state $\tilde{l}^i$ is assumed have a shifted lognormal distribution:

$$\log(\tilde{l}^i + \kappa_i) \sim N(\mu_i, \sigma_i^2).$$

All parameters may vary over time (since individual demographics vary over time), but t subscripts are omitted here for simplicity. The moment-matching conditions for the distribution of the household-level health state $\tilde{l}$ are:

$$E(\tilde{l} + \kappa_k) = \sum_{i \in \mathcal{I}_k} E(\tilde{l}^i + \kappa_i), \tag{C.1}$$

$$Var(\tilde{l} + \kappa_k) = \sum_{i \in \mathcal{I}_k} Var(\tilde{l}^i + \kappa_i), \tag{C.2}$$

$$E(\tilde{l}) = \sum_{i \in \mathcal{I}_k} E(\tilde{l}^i), \tag{C.3}$$

where $\mathcal{I}_k$ is the set of individuals in household k. Equation C.1 sets the mean of the household's distribution equal to the sum of the means of each individual's distribution. Equation C.2 matches the variance. Because we have a third parameter to estimate (the shift, κ_k), we use a third moment-matching condition to match the first moment of the unshifted distribution, shown in Equation C.3.

Under the approximating assumption that $\tilde{l} + \kappa_k$ is distributed lognormally, and substituting the analytical expressions for the mean and variable of a lognormal distribution, these equations become:

$$\exp(\mu_k + \frac{\sigma_k^2}{2}) = \sum_{i \in \mathcal{I}_k} \exp(\mu_i + \frac{\sigma_i^2}{2})$$

$$\big(\exp(\sigma_k^2) - 1\big)\exp(2\mu_k + \sigma_k^2) = \sum_{i \in \mathcal{I}_k} \big(\exp(\sigma_i^2) - 1\big)\exp(2\mu_i + \sigma^2)$$

$$\exp(\mu_k + \frac{\sigma_k^2}{2}) - \kappa_k = \sum_{i \in \mathcal{I}_k} \exp(\mu_i + \frac{\sigma_i^2}{2}) - \kappa_i$$

Given a guess of the parameters to be estimated (the individual-level parameters), this leaves three equations in three unknowns, and we can solve for the household-level parameters. The solutions for μ_k, σ_k^2, and κ_k are:

$$\sigma_k^2 = \log[1 + \left[\sum_{i \in \mathcal{I}_k} \exp(\mu_i + \frac{\sigma_i^2}{2})\right]^{-2} \sum_{i \in \mathcal{I}_k} (\exp(\sigma_i^2) - 1)\exp(2\mu_i + \sigma_i^2)]$$
$$\mu_k = -\frac{\sigma_k^2}{2} + \log[\sum_{i \in \mathcal{I}_k} \exp(\mu_i + \frac{\sigma_i^2}{2})]$$
$$\kappa_k = \sum_{i \in \mathcal{I}_k} \kappa_i$$

Given these algebraic solutions for the parameters of a household's health state distribution, we can work backward to estimate which individual-level parameters best explain the observed data on individual-level demographics and household-level healthcare spending. A key advantage of using this approximation instead of simply simulating the true distribution of the sum of lognormals is that we can use quadrature to integrate the distributions of health states, thereby limiting the number of support points needed for numerical integration.

C.2 Estimation Algorithm

We estimate the model using a maximum likelihood approach similar to that described by Revelt and Train (1998) and Train (2009), with the appropriate extension to a discrete/continuous choice model in the style of Dubin and McFadden (1984). The maximum likelihood estimator selects the parameter values that maximize the conditional probability density of households' observed total healthcare spending, given their plan choices.

The model contains four dimensions of unobservable heterogeneity: risk aversion, household health, the moral hazard parameter, and the T1-EV idiosyncratic shock. The last we can integrate analytically, but the first three we must integrate numerically; we denote these as $\beta_{kt} = \{\psi_k, \mu_{kt}, \omega_k\}$. We denote the full set of parameters to be estimated as θ, which, among other things, contains the parameters of the distribution of β_{kt}. Given a guess of θ, we simulate the distribution of β_{kt} using Gaussian quadrature with 27 support points, yielding simulated points $\beta_{kts}(\theta) = \{\psi_{ks}, \mu_{kts}, \omega_{ks}\}$, as well as weights W_s.[17,18] For each simulation draw s, we then calculate the conditional density at households' observed total healthcare spending and the probability of households' observed plan choices.

[17]Note that the mean vector of β_{kts} is a fixed function of θ and household demographics.

[18]We use the Matlab program *qnwnorm* to implement this method, with three points in each dimension of unobserved heterogeneity. The program can be obtained as part of Mario Miranda and Paul Fackler's CompEcon Toolbox; for more information, see Miranda and Fackler (2002).

Conditional Probability Density of Healthcare Spending. We have data on realized healthcare spending m_{kt} for each household and year. We aim to construct the distribution of healthcare spending for each household-year implied by the model and guess of parameters. We start by constructing individual-level health state distribution parameters μ_{it}, σ_{it}, and κ_{it} from θ and individual demographics, as described in Equation 7. We then construct household-level health state distribution parameters μ_{kts}, σ_{kt}, and κ_{kt} using the formulas in Equation 8 and the draws of $\beta_{kts}(\theta)$. The model predicts that upon realizing their health state l, households choose total healthcare spending m by trading off the benefit of healthcare utilization with its out-of-pocket cost. Specifically, accounting for the fact that negative health states may imply zero spending, the model predicts optimal healthcare spending $m^*_{jt}(l, \omega_{ks}) = \max(0\ ,\ \omega_{ks}(1 - c'_{jt}(m^*)) + l)$ if household k were enrolled in plan j in year t. Inverting the expression, the health state realization l_{kjts} that would have given rise to observed spending m_{kt} under moral hazard parameter ω_{ks} is given by

$$l_{kjts} : \begin{cases} l_{kjts} < 0 & m_{kt} = 0 \\ l_{kjts} = m_{kt} - \omega_{ks}(1 - c'_{jt}(m_{kt})) & m_{kt} > 0. \end{cases}$$

Household health state is distributed according to

$$l = \phi_f \tilde{l}$$
$$\log(\tilde{l} + \kappa_{kt}) \sim N(\mu_{kts}, \sigma^2_{kt}).$$

There are two possibilities to consider. First, if m_{kt} is equal to zero, the implied health state realization l_{kjts} is negative. Given monetary health state realization l_{kjts}, the implied "quantity" health state realization is equal to $\tilde{l}_{kjts} = \phi_f^{-1} l_{kjts}$, where f is the insurer offering plan j. Since $\phi_f > 0$, the probability of observing $l_{kjts} < 0$ is the probability of observing $\tilde{l}_{kjts} \leq \kappa_{kt}$. Second, if m_{kt} is greater than zero, it is useful to define $\lambda_{kjts} = \phi_f^{-1} l_{kjts} + \kappa_{kt}$, which itself is distributed lognormally (no shift). The density of m_{kt} in this case is given by the density of λ_{kjts}. Taken together, the probability density of total healthcare spending m conditional on plan, parameters, and household observables $\mathbf{X}_{kt}$ is given by $f_m(m_{kt}|c_{jt}, \beta_{kts}, \theta, \mathbf{X}_{kt}) = P(m = m_{kt}|c_{jt}, \beta_{kts}, \theta, \mathbf{X}_{kt})$, where

$$f_m(m_{kt}|c_{jt}, \beta_{ks}, \theta, \mathbf{X}_{kt}) = \begin{cases} \Phi\left(\frac{\log(\kappa_{kt}) - \mu_{kt}}{\sigma_{kt}}\right) & m_{kt} = 0, \\ \phi_f^{-1} \Phi'\left(\frac{\log(\lambda_{kjts}) - \mu_{kt}}{\sigma_{kt}}\right) & m_{kt} > 0, \end{cases}$$

and $\Phi(\cdot)$ is the standard normal cumulative distribution function. For a given guess of parameters, there are certain values of m_{kt} for which the probability density is zero. In order to rationalize the data at all possible parameter guesses, in practice we use a convolution of $f_m(m_{kt}|c_{jt}, \beta_{ks}, \theta, \mathbf{X}_{kt})$ and a uniform distribution over the range [-1e-75, 1e75].[19]

Probability of Plan Choices. We next calculate the probability of a household's observed plan choice. Given θ and β_{kts}, we simulate the distribution of health states l_{kjtsd} using $D = 30$ support points:

$$l_{kjtsd} = \phi_f\big(\exp(\mu_{kts} + \sigma_{kt} Z_d) - \kappa_{kt}\big),$$

where Z_d is a vector of points that approximates a standard normal distribution using Gaussian quadrature, and W_d (to be used soon) are the associated weights. We then calculate the privately optimal healthcare spending choice m_{kjtsd} associated with each potential health state realization.

Plans in our empirical setting are characterized by a deductible D, a coinsurance rate C, and an out-of-pocket maximum O. Marginal out-of-pocket costs $c'(m)$ equal 1 in the deductible region, c in the coinsurance region, and 0 in the out-of-pocket maximum region. Denote the boundary between the coinsurance region and the out-of-pocket maximum region (the "stop loss" level of total spending) by $A = C^{-1}(O - D(1 - C))$. Privately optimal spending falls into one of these three regions depending on the realization of the health state l and the moral hazard parameter ω. The relevant cutoff values for the health state are

$$\begin{aligned} Z_1 &= D - \omega(1 - C)/2, \\ Z_2 &= O - \omega/2, \\ Z_3 &= A - \omega(1 - C/2), \end{aligned}$$

where $Z_1 \leq Z_2 \leq Z_3$ so long as $O \geq D$ and $C \in [0, 1]$. There are two types of plans to consider. If D and A are sufficiently far apart (there is a sufficiently large coinsurance region), then only the cutoffs Z_1 and Z_3 matter, and it may be optimal to be in any of the three regions, depending on where the health state is relative to those two cutoff values. If D and A are close together, it will never be optimal to be in the coinsurance region (better to burn right though it and into the free healthcare of the out-of-pocket maximum region), and the cutoff Z_2 will determine whether the deductible or out-of-pocket maximum region is optimal. If the realized health state is negative, optimal spending will equal zero. In sum:

[19] We have experimented with varying these bounds and found that this does not affect parameter estimates as long as the uniform density is sufficiently small.

$$\text{If} \quad A - D > \omega/2: \qquad m^* = \begin{cases} \max(0,l) & l \le Z_1, \\ l + \omega(1-C) & Z_1 < l \le Z_3, \\ l + \omega & Z_3 < l; \end{cases}$$

$$\text{If} \quad A - D \le \omega/2: \qquad m^* = \begin{cases} \max(0,l) & l \le Z_2, \\ l + \omega & Z_2 < l. \end{cases}$$

A graphical example (of the case in which the coinsurance region is sufficiently large) is shown in Figure A.2b. All plans in our empirical setting have $A - D > \omega/2$ at reasonable values of ω.

With distributions of privately optimal total healthcare spending m^*_{kjtsd} in hand for each household, plan, year, and draw of β_{ks}, we can calculate households' expected utility from enrolling in each potential plan. We construct the numerical approximation to Equation 5 using the quadrature weights W_d:

$$U_{kjts} = -\sum_{d=1}^{D} W_d \cdot \exp\left(-\psi_k z_{kjts}(l_{kjtsd})\right),$$

where the monetary payoff z is calculated as in Equation 6. To avoid numerical issues arising from double-exponentiation, we estimate the model in certainty-equivalent units of U_{kjts}:

$$U^{CE}_{kjts} = \bar{z}_{kjts} - \frac{1}{\psi_k} \log\left(\sum_{d=1}^{D} W_d \cdot \exp\Big(-\psi_k\big(z_{kjts}(l_{kjtsd}) - \bar{z}_{kjts}\big)\Big)\right),$$

where $\bar{z}_{kjts} = \mathbb{E}_d[z_{kjts}(l_{kjtsd})]$. Another reason for estimating the model in certainty equivalents is that it becomes simple to denominate the logit error term in dollars rather than in utils. This ensures that our choice model is "monotone," in the sense that the probability of preferring a less-risky plan is everywhere increasing in risk aversion; see Apesteguia and Ballester (2018) for a full treatment of this issue.

Choice probabilities, conditional on β_{kts}, are given by the standard logit formula:

$$L_{kjts} = \frac{\exp(U^{CE}_{kjts}/\sigma_\epsilon)}{\sum_{i \in \mathcal{J}_{kt}} \exp(U^{CE}_{kits}/\sigma_\epsilon)}.$$

Likelihood Function. The numerical approximation to the likelihood of the sequence of choices and healthcare spending amounts for a given household is given by

$$LL_k = \sum_{j=1}^{J} d_{kjt} \sum_{s=1}^{S} W_s \prod_{t=1}^{T} f_m(m_{kt}|\theta, \beta_{kts}, c_{jt}, \mathbf{X}_{kt}) L_{kjts},$$

where $d_{kjt} = 1$ if household k chose plan j in year t and zero otherwise. The log-likelihood function for parameters θ is

$$LL(\theta) = \sum_{k=1}^{K} \log\left(LL_k\right).$$

C.3 Recovering Household-specific Types

We assume that household types $\beta_{kt}(\theta) = \{\psi_k, \mu_{kt}, \omega_k\}$ are distributed multivariate normal with observable heterogeneity in the mean vector, according to Equation 9. After estimating the model and obtaining $\hat{\theta}$, we want to use each household's observed outcomes (plan choices and healthcare spending amounts) to back out which type they are likely to be. Let $g(\beta|\hat{\theta})$ denote the population distribution of types. Let $h(\beta|\hat{\theta}, y)$ denote the density of β conditional on parameters $\hat{\theta}$ and a sequence of observed plan choices and healthcare spending amounts y. Using what Revelt and Train (2001) term the "conditioning of individual tastes" method, we recover households' posterior distribution of β using Bayes' rule:

$$h(\beta|\hat{\theta}, y) = \frac{p(y|\beta)g(\beta|\hat{\theta})}{p(y|\hat{\theta})}.$$

Taking the numerical approximations, $p(y|\hat{\theta})$ is simply the household-specific likelihood function LL_k for an observed sequence of plan choices and spending amounts; $g(\beta|\hat{\theta})$ is the quadrature weights W_s on each simulated point; and $p(y|\beta)$ is the *conditional* household likelihood function LL_{ks}:

$$LL_{ks} = \sum_{j=1}^{J} d_{kjt} \prod_{t=1}^{T} f_m(m_{kt}|\theta, \beta_{ks}, c_{jt}, \mathbf{X}_{kt}) L_{kjts}.$$

Taken together, the numerical approximation to each household's posterior distribution of unobserved heterogeneity is given by

$$h_{ks}(\beta|\hat{\theta}, y_k) = \frac{LL_{ks} \cdot W_s}{LL_k},$$

where $\sum_s h_{ks}(\beta|\hat{\theta}, y_k) = 1$.

For the purposes of examining total variation in types across households (accounting for both observed and unobserved heterogeneity), we assign each household the expectation of their type with respect to their posterior distribution.

We also use the household-specific distributions over types to calculated expected quantities of interest for each household. In particular, we calculate WTP_{kjt} and SS_{kjt} as

$$WTP_{kjt} = \sum_s h_{ks}(\beta|\hat{\theta}, y_k) WTP_{kjts},$$
$$SS_{kjt} = \sum_s h_{ks}(\beta|\hat{\theta}, y_k) SS_{kjts}.$$

Joint Distribution of Household Types. We investigate the distribution implied by our primary estimates in column 3 of Tables 3 and A.8. For each household, we first calculate the expectation of their type with respect to their posterior distribution of unobservable heterogeneity:

$$\psi_k = \sum_s h_{ks}(\beta|\hat{\theta}, y_k)\psi_{ks},$$
$$\omega_k = \sum_s h_{ks}(\beta|\hat{\theta}, y_k)\omega_{ks}.$$

In place of μ_{kt}, a more relevant measure of household health is the expected health state, i.e., expected total spending absent moral hazard. Using the expectation of a shifted lognormal variable and price parameter $\phi = 1$, the expected health state $\bar{l}_{kt}$ is given by

$$\bar{l}_{kt} = \sum_s h_{ks}(\beta|\hat{\theta}, y_k)(\exp(\mu_{kts} + \frac{\sigma_{kt}^2}{2}) - \kappa_{kt}).$$

To limit our focus to one type for each household, we look at $\bar{l}_{kt}$ for the first year each household appears in the data. Figure A.3 presents the joint distribution of household types along the dimensions of risk aversion, moral hazard parameter, and expected health state. We measure the expected health state on a log scale for readability.

References

Apesteguia, Jose, and Miguel A. Ballester. 2018. "Monotone Stochastic Choice Models: The Case of Risk and Time Preferences." *Journal of Political Economy*, 126(1): 74–106.

Bogin, Alexander, William Doerner, and William Larson. 2019. "Local House Price Dynamics: New Indices and Stylized Facts." *Real Estate Economics*, 47(2): 365–398.

Cobb, Barry, Rafael Rumí, and Antonio Salmerón. 2012. "Approximating the Distribution of a Sum of Log-normal Random Variables." *Proceedings of the 6th European Workshop on Probabilistic Graphical Models, PGM 2012*.

Dubin, Jeffrey A., and Daniel L. McFadden. 1984. "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption." *Econometrica*, 52(2): 345–362.

Einav, Liran, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen. 2013. "Selection on moral hazard in health insurance." *American Economic Review*, 103(1): 178–219.

Fenton, L. F. 1960. "The sum of log-normal probability distributions in scatter transmission systems." *IRE Transactions on Communication Systems*, 8: 57–67.

Ho, Kate, and Robin Lee. 2021. "Health Insurance Menu Design for Large Employers."

Manning, Willard G., Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, and Arleen Leibowitz. 1987. "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment." *The American Economic Review*, 77(3): 251–277.

Miranda, Mario J., and Paul L. Fackler. 2002. *Applied Computational Economics and Finance.* MIT Press.

Newhouse, Joseph. 1993. "Free for All? Lessons from the RAND Health Insurance Experiment." *Cambridge, MA. Harvard University Press.*

Revelt, David, and Kenneth Train. 1998. "Mixed Logit with Repeated Choices: Households' Choices of Appliance Efficiency Level." *The Review of Economics and Statistics*, 80(4): 647–657.

Revelt, David, and Kenneth Train. 2001. "Customer-Specific Taste Parameters and Mixed Logit: Households' Choice of Electricity Supplier." , (0012001).

Train, Kenneth. 2009. *Discrete Choice Methods with Simulation: Second Edition.* Cambridge University Press.

Table A.1. Plan Characteristics

Year	Plan	Actuarial Value	Avg. Employee Premium ($)	Full Premium ($)	Deductible ($)	OOP Max. ($)	Market Share
2008	Kaiser - 1	0.97	682	9,768	0	1,200	0.07
	Kaiser - 2	0.96	313	9,334	0	2,000	0.10
	Moda - 1	0.92	1,086	11,051	300	500	0.28
	Moda - 2	0.89	648	10,613	300	1,000	0.06
	Moda - 3	0.88	363	10,097	600	1,000	0.11
	Moda - 4	0.86	461	9,674	900	1,500	0.07
	Moda - 5	0.82	273	8,888	1,500	2,000	0.12
	Moda - 6	0.78	320	8,032	3,000	3,000	0.03
	Moda - 7	0.68	37	6,141	3,000	10,000	$<$0.01
	Providence - 1	0.96	1,005	10,645	900	1,200	0.14
	Providence - 2	0.95	933	10,563	900	2,000	0.02
2010	Kaiser - 1	0.96	701	11,586	0	2,400	0.17
	Kaiser - 2	0.95	582	11,231	0	3,000	0.03
	Moda - 1	0.89	3,876	15,794	600	1,200	0.10
	Moda - 2	0.86	2,867	14,579	600	1,500	0.01
	Moda - 3	0.85	1,833	13,300	600	1,800	0.17
	Moda - 4	0.84	897	11,904	900	2,000	0.12
	Moda - 5	0.82	528	10,890	1,500	2,000	0.21
	Moda - 6	0.78	311	9,795	3,000	3,000	0.09
	Moda - 7	0.75	106	7,472	3,000	10,000	0.02
	Providence - 1	0.91	4,702	16,680	1,200	1,200	0.04
	Providence - 2	0.89	4,314	16,245	1,800	1,800	0.01
2011	Kaiser - 1	0.95	520	11,051	0	2,400	0.16
	Kaiser - 2	0.92	348	10,126	300	4,000	0.04
	Moda - 1	0.86	3,414	15,622	600	4,500	0.06
	Moda - 2	0.84	1,009	12,391	900	6,000	$<$0.01
	Moda - 3	0.84	1,208	12,688	900	6,000	0.15
	Moda - 4	0.83	603	11,334	1,200	6,300	0.09
	Moda - 5	0.82	367	10,188	1,500	6,600	0.24
	Moda - 6	0.78	190	8,764	3,000	6,600	0.15
	Moda - 7	0.75	130	7,806	3,000	10,000	0.05
	Providence - 1	0.87	2,835	14,882	300	3,600	0.02
	Providence - 2	0.84	2,066	13,891	900	6,000	$<$0.01
2012	Kaiser - 1	0.95	1,478	13,408	0	2,400	0.18
	Kaiser - 2	0.93	843	12,278	450	4,000	0.04
	Moda - 1	0.87	5,677	18,514	600	4,500	0.06
	Moda - 2	0.85	2,164	14,299	900	6,000	0.01
	Moda - 3	0.85	2,995	15,359	900	6,000	0.12
	Moda - 4	0.84	1,899	13,902	1,200	6,300	0.06
	Moda - 5	0.83	1,082	12,670	1,500	6,600	0.22
	Moda - 6	0.79	501	11,139	3,000	6,600	0.17
	Moda - 7	0.76	148	8,395	3,000	10,000	0.11
2013	Kaiser - 1	0.95	1,815	14,203	0	3,000	0.20
	Kaiser - 2	0.94	998	12,895	600	4,400	0.03
	Moda - 1	0.87	6,537	19,675	600	6,000	0.03
	Moda - 2	0.85	3,069	15,765	1,050	7,200	0.08
	Moda - 3	0.84	1,152	13,157	1,500	7,800	0.22
	Moda - 4	0.82	692	12,212	2,250	8,400	0.06
	Moda - 5	0.80	493	11,427	3,000	9,000	0.11
	Moda - 6	0.78	344	10,480	3,750	12,000	0.05
	Moda - 7	0.77	151	8,574	3,000	10,000	0.13
	Moda - 8	0.76	224	9,474	4,500	15,000	0.05

Notes: The table shows the state-level master lists of plans available in 2008 and 2010–2013. The full premium reflects the premium negotiated by OEBB and the insurer; the one shown is for an employee plus spouse. The deductible and out-of-pocket maximum shown are for in-network services for a family household. This table is referenced in Section III.A.

Table A.2. Sample Construction

Criteria	2009	2010	2011	2012	2013
Individuals in membership file	161,502	162,363	156,113	156,042	157,799
Not eligible for coverage	7,370	8,265	8,422	8,719	8,388
Retiree, COBRA, or oldest member over 65	13,180	12,567	12,057	11,603	11,840
Partial year coverage	17,115	18,649	19,283	21,281	23,074
Covered by multiple plans	1,447	1,947	2,038	2,239	2,336
Opted out	3,241	4,205	4,321	4,576	4,529
Not in intact family	8,389	9,188	9,181	8,925	10,265
No prior year of data	6,175	3,947	2,455	3,104	3,702
Missing premium or contribution data	25,653	28,466	22,755	23,284	30,401
Final total	78,932	75,129	75,601	72,311	63,264

Notes: The table shows the counts of individuals dropped due to each sample selection criterion. Drops are made in the order in which criteria appear. All observations in 2008 are dropped because there is no year of prior data. This table is referenced in Section III.A.

Figure A.1. Example of Plan Cost-sharing Features Estimation

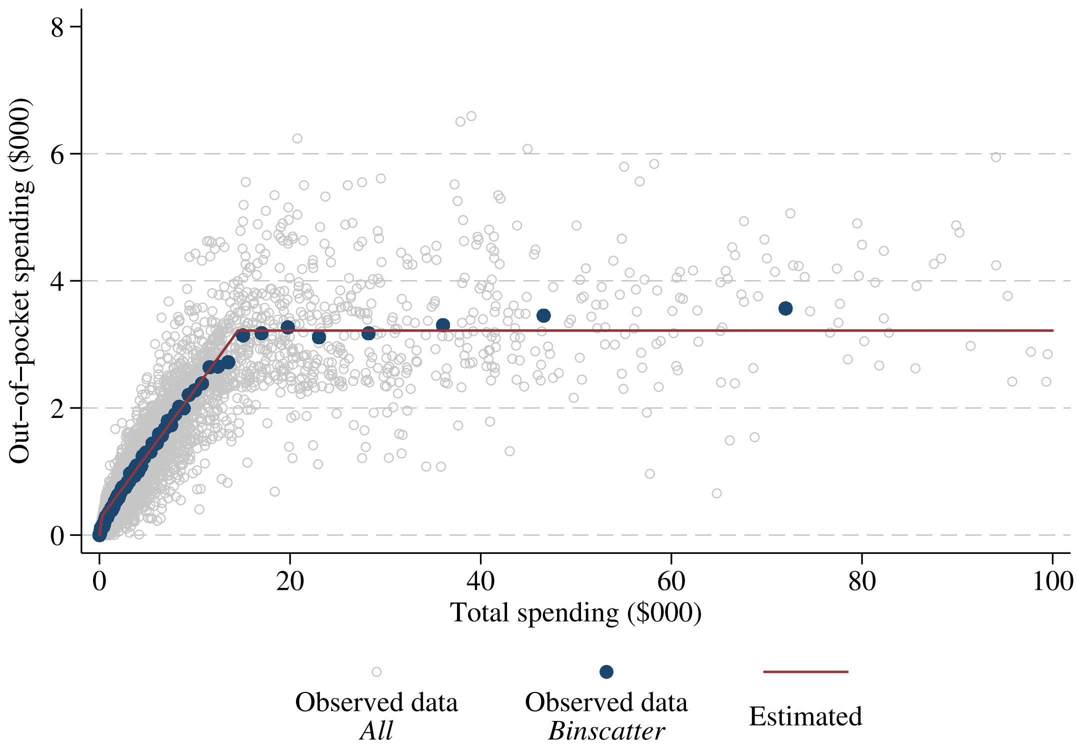

Notes: The figure shows the data used to estimate the cost-sharing features for plan Moda - 3 for individual households in 2012. Total healthcare spending is on the horizontal axis and out-of-pocket cost is on the vertical axis. Each gray dot represents a household, for a 20 percent random sample of households. The blue dots are a binned scatter plot of the gray data, using 100 points. The basic cost-sharing features of the plan (as observed in plan documents) are a deductible of $300, nonspecialist coinsurance rate of 20 percent, and in-network out-of-pocket maximum of $2,000. We estimate a best-fit cost-sharing function by finding the coinsurance rate and out-of-pocket maximum that minimizes the sum of squared errors between predicted and observed out-of-pocket spending. The estimated coinsurance rate is 20.5 percent and the estimated out-of-pocket maximum is $3,218. This figure is referenced in Appendix B.1.

Table A.3. Estimated Plan Characteristics

Year	Plan	Indiv.: Ded.	Coins.	OOP Max.	Fam.: Ded.	Coins.	OOP Max.
2009	Kaiser - 1	0	0.03	564	0	0.03	645
	Kaiser - 2	0	0.03	684	0	0.04	760
	Kaiser - 3	0	0.03	734	0	0.04	791
	Moda - 1	100	0.10	1,613	300	0.10	2,009
	Moda - 2	100	0.18	1,922	300	0.15	2,662
	Moda - 3	200	0.20	2,081	600	0.15	3,062
	Moda - 4	300	0.19	2,796	900	0.15	3,835
	Moda - 5	500	0.22	3,164	1,500	0.16	4,296
	Moda - 6	1,000	0.22	3,713	3,000	0.12	5,422
	Moda - 7	1,500	0.42	4,693	3,000	0.30	8,086
	Providence - 1	300	0.02	790	900	0.00	900
	Providence - 2	300	0.03	867	900	0.00	986
	Providence - 3	300	0.04	1,116	900	0.01	1,296
2010	Kaiser - 1	0	0.03	697	0	0.04	805
	Kaiser - 2	0	0.04	820	0	0.05	885
	Moda - 1	200	0.14	2,526	600	0.12	3,430
	Moda - 2	200	0.21	2,846	600	0.18	3,967
	Moda - 3	200	0.21	3,189	600	0.18	4,299
	Moda - 4	300	0.22	3,109	900	0.18	4,079
	Moda - 5	500	0.22	3,321	1,500	0.16	4,572
	Moda - 6	1,000	0.22	3,844	3,000	0.12	5,684
	Moda - 7	1,500	0.19	4,913	3,000	0.15	7,579
	Providence - 1	400	0.05	1,523	1,200	0.02	1,851
	Providence - 2	600	0.06	1,998	1,800	0.02	2,473
2011	Kaiser - 1	0	0.04	883	0	0.06	974
	Kaiser - 2	100	0.06	1,340	300	0.06	1,831
	Moda - 1	200	0.22	2,608	600	0.18	4,316
	Moda - 2	300	0.22	3,201	900	0.17	5,094
	Moda - 3	300	0.22	3,246	900	0.17	5,202
	Moda - 4	400	0.22	3,324	1,200	0.17	5,367
	Moda - 5	500	0.22	3,529	1,500	0.16	5,727
	Moda - 6	1,000	0.22	4,061	3,000	0.13	6,728
	Moda - 7	1,500	0.21	4,914	3,000	0.15	7,663
	Providence - 1	100	0.18	2,164	300	0.16	3,496
	Providence - 2	300	0.15	2,911	900	0.13	4,378
2012	Kaiser - 1	0	0.04	911	0	0.06	995
	Kaiser - 2	150	0.07	1,709	450	0.05	2,160
	Moda - 1	200	0.21	2,571	600	0.17	4,154
	Moda - 2	300	0.21	3,187	900	0.17	4,981
	Moda - 3	300	0.20	3,218	900	0.17	5,025
	Moda - 4	400	0.21	3,291	1,200	0.16	5,104
	Moda - 5	500	0.21	3,493	1,500	0.16	5,498
	Moda - 6	1,000	0.21	4,000	3,000	0.12	6,608
	Moda - 7	1,500	0.21	4,927	3,000	0.15	7,662
2013	Kaiser - 1	0	0.04	911	0	0.06	1,040
	Kaiser - 2	200	0.03	867	600	0.01	951
	Moda - 1	200	0.20	3,237	600	0.17	4,893
	Moda - 2	350	0.20	3,842	1,050	0.16	5,647
	Moda - 3	500	0.20	4,175	1,500	0.15	6,160
	Moda - 4	750	0.20	4,704	2,250	0.14	6,989
	Moda - 5	1,000	0.19	5,186	3,000	0.12	7,714
	Moda - 6	1,250	0.19	6,414	3,750	0.12	9,187
	Moda - 7	1,500	0.21	4,865	3,000	0.15	7,650
	Moda - 8	1,500	0.19	7,620	4,500	0.11	10,614

Notes: The table shows plan deductibles, estimated coinsurance rates, and estimated out-of-pocket maximums. The estimation procedure is described in Appendix B.1.

Table A.4. Plan Choice Logit Model

	2008	2009	2010	2011	2012	2013
Employee premium ($000)	-0.789	-0.674	-0.505	-0.372	-0.515	-0.490
	(0.017)	(0.014)	(0.008)	(0.010)	(0.008)	(0.008)
HRA/HSA contrib. ($000)	0.112		0.358	0.134	0.269	0.534
	(0.759)		(0.044)	(0.024)	(0.019)	(0.015)
Vision/dental contrib. ($000)	0.654	0.408	0.480	0.794	0.553	0.710
	(0.021)	(0.022)	(0.019)	(0.017)	(0.017)	(0.017)
Kaiser - 1	-0.771	-0.728				
	(0.026)	(0.030)				
Kaiser - 2	-1.287	-1.112	-0.846	-0.469	-0.375	-0.074
	(0.031)	(0.032)	(0.034)	(0.035)	(0.034)	(0.044)
Kaiser - 3		-1.563	-1.042	-0.985	-1.629	-1.820
		(0.384)	(0.056)	(0.051)	(0.048)	(0.058)
Moda - 1	$0.000^{\dagger}$	$0.000^{\dagger}$	$0.000^{\dagger}$	$0.000^{\dagger}$	$0.000^{\dagger}$	$0.000^{\dagger}$
Moda - 2	-1.113	-1.184	-0.911	-2.088	-2.578	-0.593
	(0.026)	(0.032)	(0.058)	(0.163)	(0.072)	(0.045)
Moda - 3	-1.226	-1.110	-0.518	-0.373	-0.389	-0.957
	(0.022)	(0.025)	(0.029)	(0.034)	(0.033)	(0.046)
Moda - 4	-1.751	-1.540	-1.356	-1.192	-1.554	-2.261
	(0.028)	(0.030)	(0.034)	(0.037)	(0.039)	(0.055)
Moda - 5	-1.951	-1.881	-1.341	-0.878	-0.999	-2.391
	(0.034)	(0.037)	(0.040)	(0.039)	(0.037)	(0.055)
Moda - 6	-2.785	-2.871	-2.205	-1.406	-1.917	-3.182
	(0.048)	(0.051)	(0.050)	(0.043)	(0.046)	(0.065)
Moda - 7	-4.391	-4.260	-3.388	-1.959	-3.007	-3.492
	(0.098)	(0.098)	(0.074)	(0.050)	(0.060)	(0.073)
Moda - 8						-3.679
						(0.068)
Providence - 1	0.001	0.048	0.135	-0.778		
	(0.019)	(0.028)	(0.038)	(0.053)		
Providence - 2	-0.600	-0.314				
	(0.043)	(0.049)				
Providence - 3		-0.048	-0.159	-0.939		
		(0.078)	(0.083)	(0.436)		
Number of observations	163,431	121,744	116,541	114,527	163,278	163,683

Notes: The table presents parameter estimates from the conditional logit model described by Equation B.1, presented separately for each year. The unit of observation is a household-plan. Moda - 1 (the highest coverage Moda plan) is the omitted plan. This table is referenced in Appendix B.2. †By normalization.

Table A.5. Plan Menu Generosity and Household Health

	2008	2009	2010	2011	2012	2013
Household risk score	-0.006	0.017	0.020	0.002	0.006	0.000
	(0.039)	(0.016)	(0.011)	(0.009)	(0.010)	(0.012)
Family type						
Employee alone	0.000†	0.000†	0.000†	0.000†	0.000†	0.000†
Employee + spouse	-1.389	-1.369	-1.498	-1.040	-1.626	-1.612
	(0.077)	(0.040)	(0.029)	(0.025)	(0.026)	(0.031)
Employee + child	-0.542	-0.634	-0.907	-0.616	-1.092	-0.937
	(0.084)	(0.053)	(0.039)	(0.031)	(0.031)	(0.037)
Employee + family	-1.792	-1.882	-1.804	-1.306	-2.147	-2.102
	(0.064)	(0.037)	(0.028)	(0.023)	(0.025)	(0.029)
Dependent variable mean	88.7	88.5	84.6	82.7	83.3	82.6
R^2	0.020	0.084	0.154	0.115	0.242	0.220
Number of observations	37,666	31,074	29,538	29,279	27,897	24,283

Notes: The table shows the relationship between plan menu generosity and household health. The unit of observation is the household. The dependent variable is plan menu generosity, as measured by predicted actuarial value conditional on choosing Moda, $\widehat{AV}_{d,Moda}$. This measure is calculated according to Equation B.2, and it is multiplied by 100 to increase the magnitude of estimated coefficients on household risk score. Household risk score is the mean risk score among all individuals in a household, and it has been z-scored such that the variable has a mean of zero and a standard deviation of one within each year. As we do not have data before 2008, the 2008 regression uses risk scores calculated using 2008 claims data. This table is referenced in Appendix B.2. †By normalization.

Table A.6. Explaining Plan Menu Generosity: 2008

	(1)	(2)	(3)	(4)
Household risk score	-0.006	0.016	0.011	0.025
	(0.039)	(0.039)	(0.038)	(0.040)
Family type				
Employee alone	0.000†	0.000†	0.000†	0.000†
Employee + spouse	-1.389	-1.374	-1.251	-1.085
	(0.077)	(0.083)	(0.083)	(0.085)
Employee + child	-0.542	-0.535	-0.478	-0.462
	(0.084)	(0.085)	(0.084)	(0.082)
Employee + family	-1.792	-1.819	-1.688	-1.437
	(0.064)	(0.071)	(0.071)	(0.074)
Part-time		-0.428	-0.448	-0.867
		(0.133)	(0.133)	(0.139)
Occupation type				
Admin.		-1.745	-1.883	-2.685
		(0.455)	(0.459)	(0.501)
Classified		-0.598	-0.469	-0.155
		(0.283)	(0.414)	(0.457)
Comm. coll. fac.		0.553	1.138	1.044
		(0.287)	(0.430)	(0.470)
Comm. coll. non-fac.		0.671	0.457	0.077
		(0.288)	(0.288)	(0.302)
Confidential		-2.759	-2.883	-3.133
		(0.855)	(0.856)	(0.915)
Licensed		0.001	1.645	1.628
		(0.278)	(0.459)	(0.505)
Substitute		-11.051	-9.312	-9.354
		(0.283)	(0.457)	(0.496)
Union affiliation				
AFT			0.251	-0.398
			(0.374)	(0.432)
IAFE			0.758	1.222
			(0.404)	(0.458)
OACE			2.671	1.617
			(0.389)	(0.449)
OEA			-1.799	-1.765
			(0.434)	(0.491)
OSEA			-0.086	-0.426
			(0.395)	(0.449)
District characteristics				
log(HPI)				-0.876
				(0.085)
Pct. Republican				-14.077
				(0.467)
Dependent variable mean	88.7	89.0	89.1	98.3
R^2	0.020	0.031	0.046	0.073
Number of observations	37,666	37,666	37,666	35,698

Notes: The table shows the relationship between plan menu generosity and household/employee characteristics. The unit of observation is the household. The dependent variable is plan menu generosity, as measured by predicted actuarial value conditional on choosing Moda, $\widehat{AV}_{d,Moda}$. This measure is calculated according to Equation B.2, and it is multiplied by 100 to increase the magnitude of estimated coefficients on household risk score. Household risk score is the mean risk score among all individuals in a household, and it has been z-scored such that the variable has a mean of zero and a standard deviation of one within each year. As we do not have data before 2008, the 2008 regression uses risk scores calculated using 2008 claims data. This table is referenced in Appendix B.2. †By normalization.

Table A.7. Estimates of Moral Hazard

	OLS *All*	IV *All*	IV *Individuals*	IV *Families*
	(1)	(2)	(3)	(4)
$\log(1 - AV_{j(k)})\mathbf{1}_{f(k)=Moda}$	-0.580	-0.269		
	(0.053)	(0.084)		
$\log(1 - AV_{j(k)})\mathbf{1}_{f(k)=Moda} \times Q_1$			-0.220	-0.415
			(0.290)	(0.131)
$\log(1 - AV_{j(k)})\mathbf{1}_{f(k)=Moda} \times Q_2$			-0.410	-0.235
			(0.189)	(0.088)
$\log(1 - AV_{j(k)})\mathbf{1}_{f(k)=Moda} \times Q_3$			-0.253	-0.218
			(0.136)	(0.090)
$\log(1 - AV_{j(k)})\mathbf{1}_{f(k)=Moda} \times Q_4$			-0.017	0.074
			(0.346)	(0.145)
R^2	0.19	0.19	0.44	0.37
Number of observations	35,146	35,146	8,962	26,184

Notes: The table shows the OLS and IV estimates of Equation B.5, describing the relationship between household total spending and plan generosity. The unit of observation is a household, and the dependent variable is log of 1 + total spending. In columns 3 and 4, coefficients can vary by household risk quartile Q_n, where Q_4 is the sickest households. Columns 1 and 2 are estimated on all households, while columns 3 and 4 are estimated only on individual or family households, respectively. All specifications also include insurer fixed effects and controls for household risk score and family structure. Standard errors (in parentheses) are clustered by household plan menu, of which there are 533 among individual households and 1,750 among family households. We can reject the hypothesis that the four coefficients are equal at the 10 percent level for families, but not for individuals. This table is referenced in Appendix B.3.

Figure A.2. Healthcare Spending Choice Example

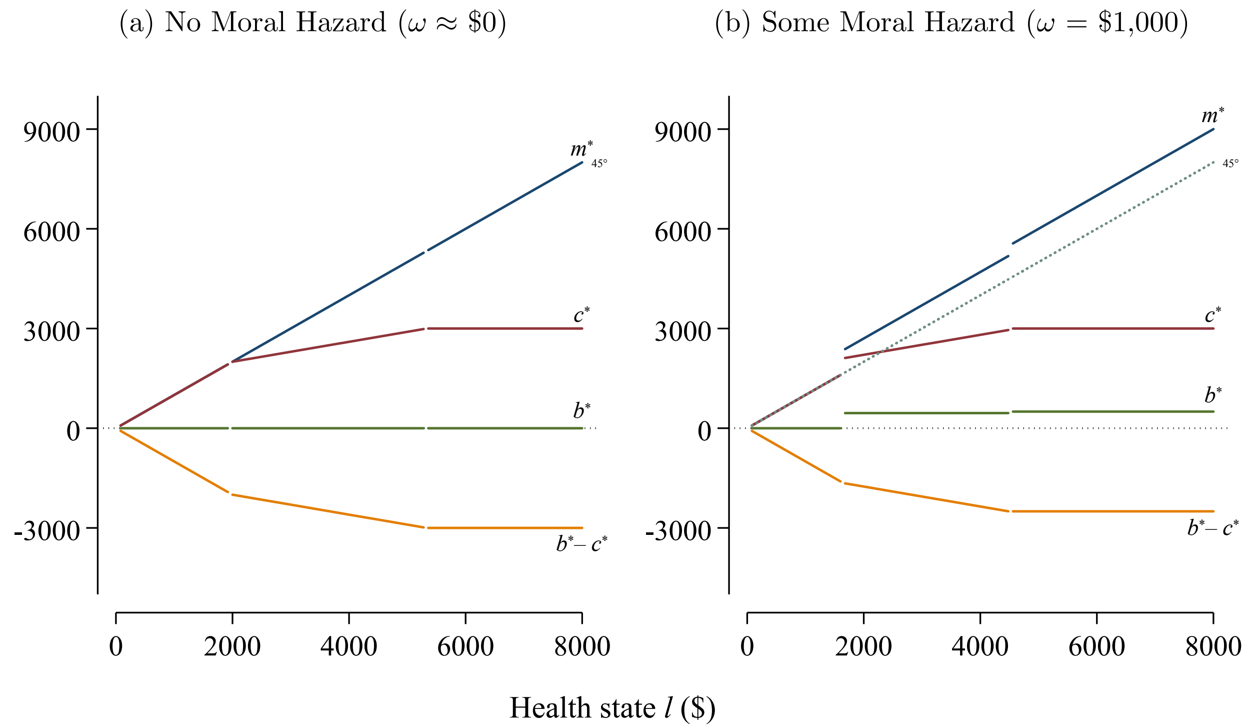

Notes: The figure shows optimal healthcare spending m^*, indirect benefit of optimal healthcare spending b^*, and the corresponding out-of-pocket cost c^* predicted by our parameterization of consumer preferences (Equation 4). The examples consider a contract with a deductible of \$2,000, a coinsurance rate of 30 percent, and an out-of-pocket maximum of \$3,000. Predicted behavior is shown under (a) no moral hazard and (b) under some moral hazard ($\omega = \$1{,}000$). The horizontal axis shows possible health state realizations l. Absent moral hazard (left panel), optimal healthcare spending is equal to the health state. The vertical axis also shows the net payoff from optimal healthcare utilization, $b^* - c^*$; this is the outcome over which households face a lottery. This figure is referenced at footnote 17 in the main text and footnote 4 in the Appendix.

Table A.8. Additional Parameter Estimates

Variable	(1) Parameter	(1) Std. Err.	(2) Parameter	(2) Std. Err.	(3) Parameter	(3) Std. Err.
Insurer fixed effects						
Providence * (Age−40)	−0.024	0.007	−0.023	0.007	−0.028	0.007
Providence * **1**[Children]	−0.681	0.151	−0.501	0.146	−0.595	0.147
Providence * Region 1	−2.114	0.144	−2.071	0.137	−1.649	0.138
Providence * Region 2	−2.658	0.185	−2.635	0.176	−2.179	0.176
Providence * Region 3	−1.877	0.207	−2.036	0.200	−1.409	0.193
Health state distributions						
κ	0.155	0.002				
κ * Risk Q_1			0.096	0.002	0.127	0.000
κ * Risk Q_2			0.224	0.002	0.155	0.001
κ * Risk Q_3			0.218	0.002	0.228	0.000
κ * Risk Q_4			0.128	0.042	0.418	0.041
κ * Risk Q_1 * Risk score			0.187	0.004	0.225	0.001
κ * Risk Q_2 * Risk score			0.140	0.002	0.019	0.002
κ * Risk Q_3 * Risk score			−0.060	0.001	0.002	0.001
κ * Risk Q_4 * Risk score			0.155	0.026	0.177	0.027
μ	0.590	0.005				
μ * **1**[Female 18–35]			0.125	0.017	0.088	0.018
μ * **1**[Age $<$ 18]			−0.113	0.017	−0.104	0.019
μ * Risk Q_1			1.405	0.137	1.872	0.154
μ * Risk Q_2			0.894	0.025	0.457	0.030
μ * Risk Q_3			0.815	0.008	0.504	0.009
μ * Risk Q_4			1.379	0.017	1.303	0.017
μ * Risk Q_1 * Risk score			3.590	0.185	4.875	0.210
μ * Risk Q_2 * Risk score			1.978	0.067	1.946	0.081
μ * Risk Q_3 * Risk score			0.894	0.019	1.053	0.022
μ * Risk Q_4 * Risk score			0.310	0.005	0.329	0.005
σ	1.174	0.002				
σ * Risk Q_1			1.626	0.006	1.748	0.007
σ * Risk Q_2			1.173	0.005	1.403	0.006
σ * Risk Q_3			1.060	0.003	1.215	0.004
σ * Risk Q_4			0.988	0.006	1.016	0.006

Notes: The table presents the parameter estimates that were not presented in Table 3. "Risk Q_n" is an indicator for an individual's risk quartile, where Q_4 is the sickest individuals. Higher risk scores correspond to worse predicted health. All parameters are measured in thousands of dollars. The insurer fixed effect of Moda is normalized to zero. This table is referenced in Section V.A.

Figure A.3. Joint Distribution of Household Types

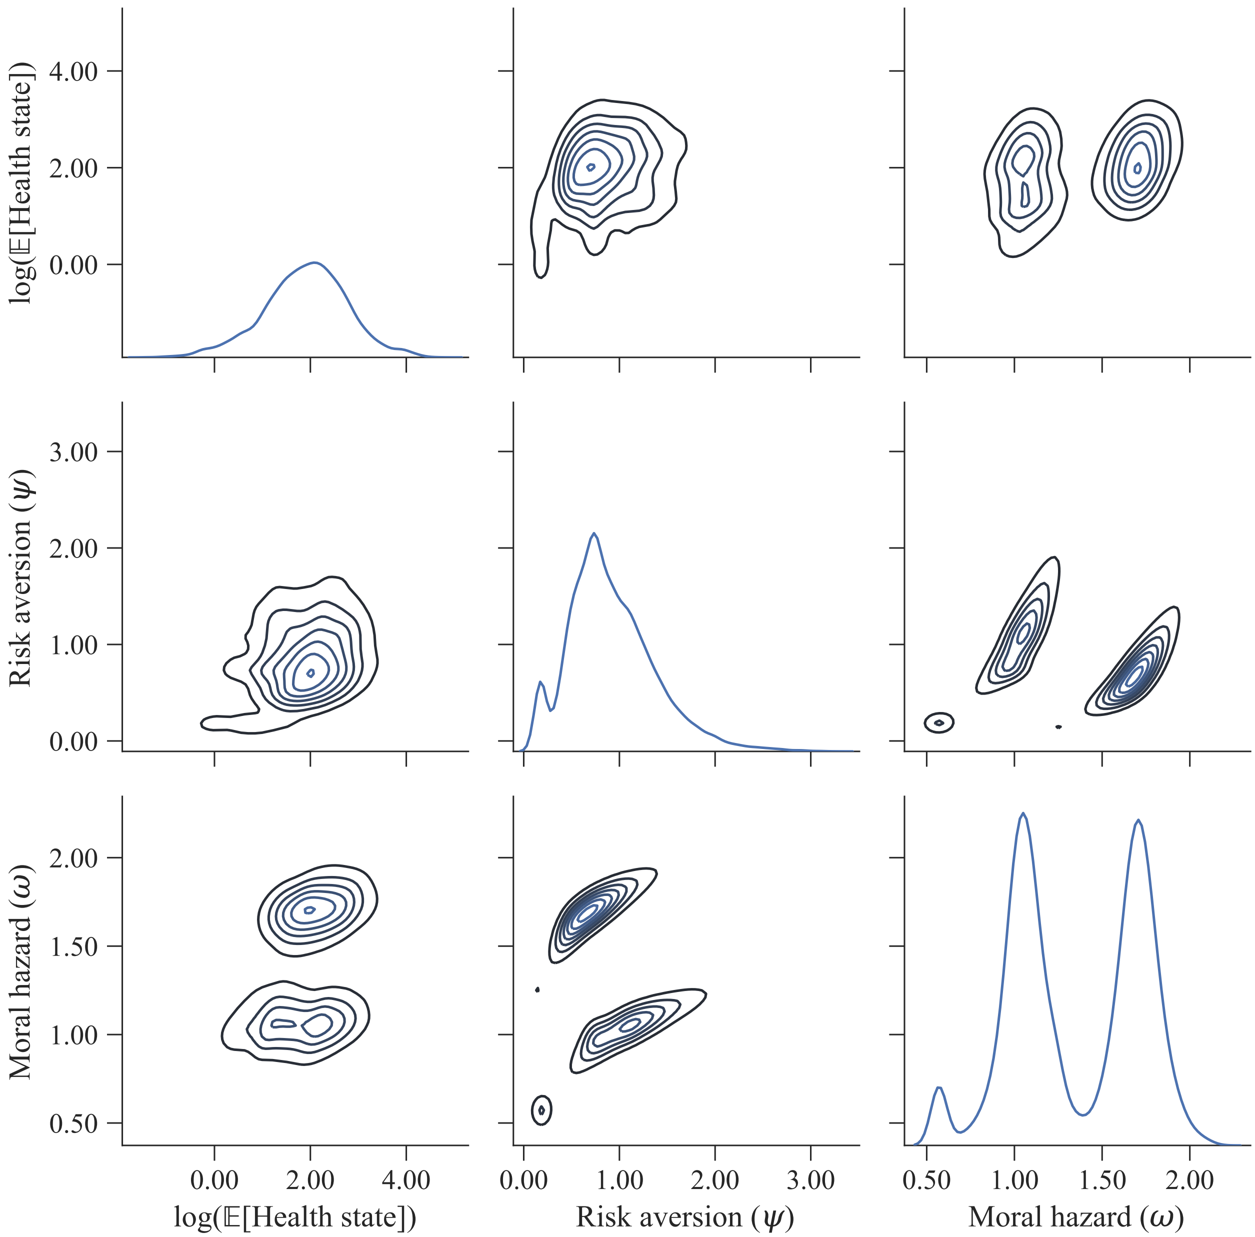

Notes: The figure shows the joint distribution of household types implied by parameter estimates in column 3 of Tables 3 and A.8. The diagonals show one-way distributions across households, and the off-diagonals show bivariate distributions. Households are ex post assigned a single type according to the procedure described in Appendix C.3. Because expected health state can vary over years within a household, this figure uses the first year a household appears in the sample. Expected health state is equivalent to a household's expected total spending absent moral hazard. This figure is referenced in Section V.A.

Figure A.4. Sets of Potential Contracts: Out-of-pocket Cost Functions

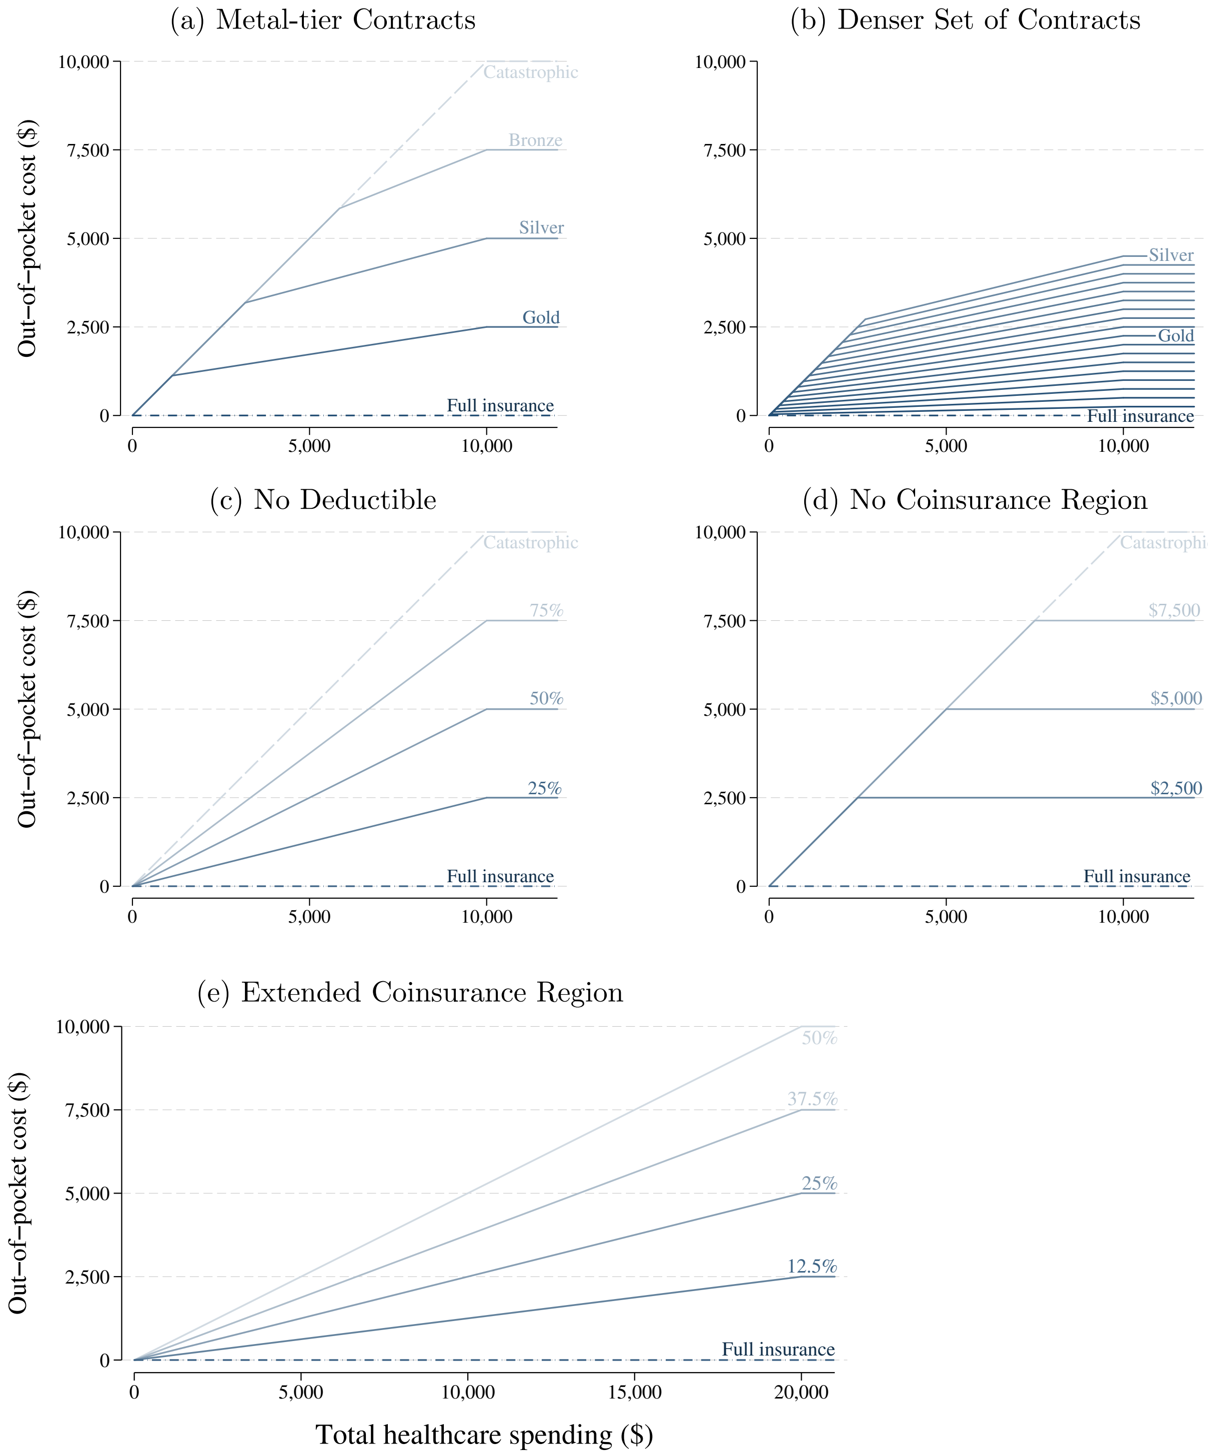

Notes: The figure shows out-of-pocket cost functions for five sets of potential contracts. Horizontal axes shows total healthcare spending, and vertical axes shows out-of-pocket cost. Panel (a) depicts our focal set of metal-tier contracts; panel (b) depicts a denser set of contracts with the same design. Panels (c)–(e) show alternative sets of potential contracts. Contract labels represent the varying feature: the coinsurance rate in panels (c) and (e) and the deductible in panel (d). Contracts are vertically differentiated and well-ordered by coverage level within each panel, but not necessarily across panels. See Appendix A.2 for these definitions. This figure is referenced in Sections V.B and V.C.

Figure A.5. Household Demographics by Willingness to Pay

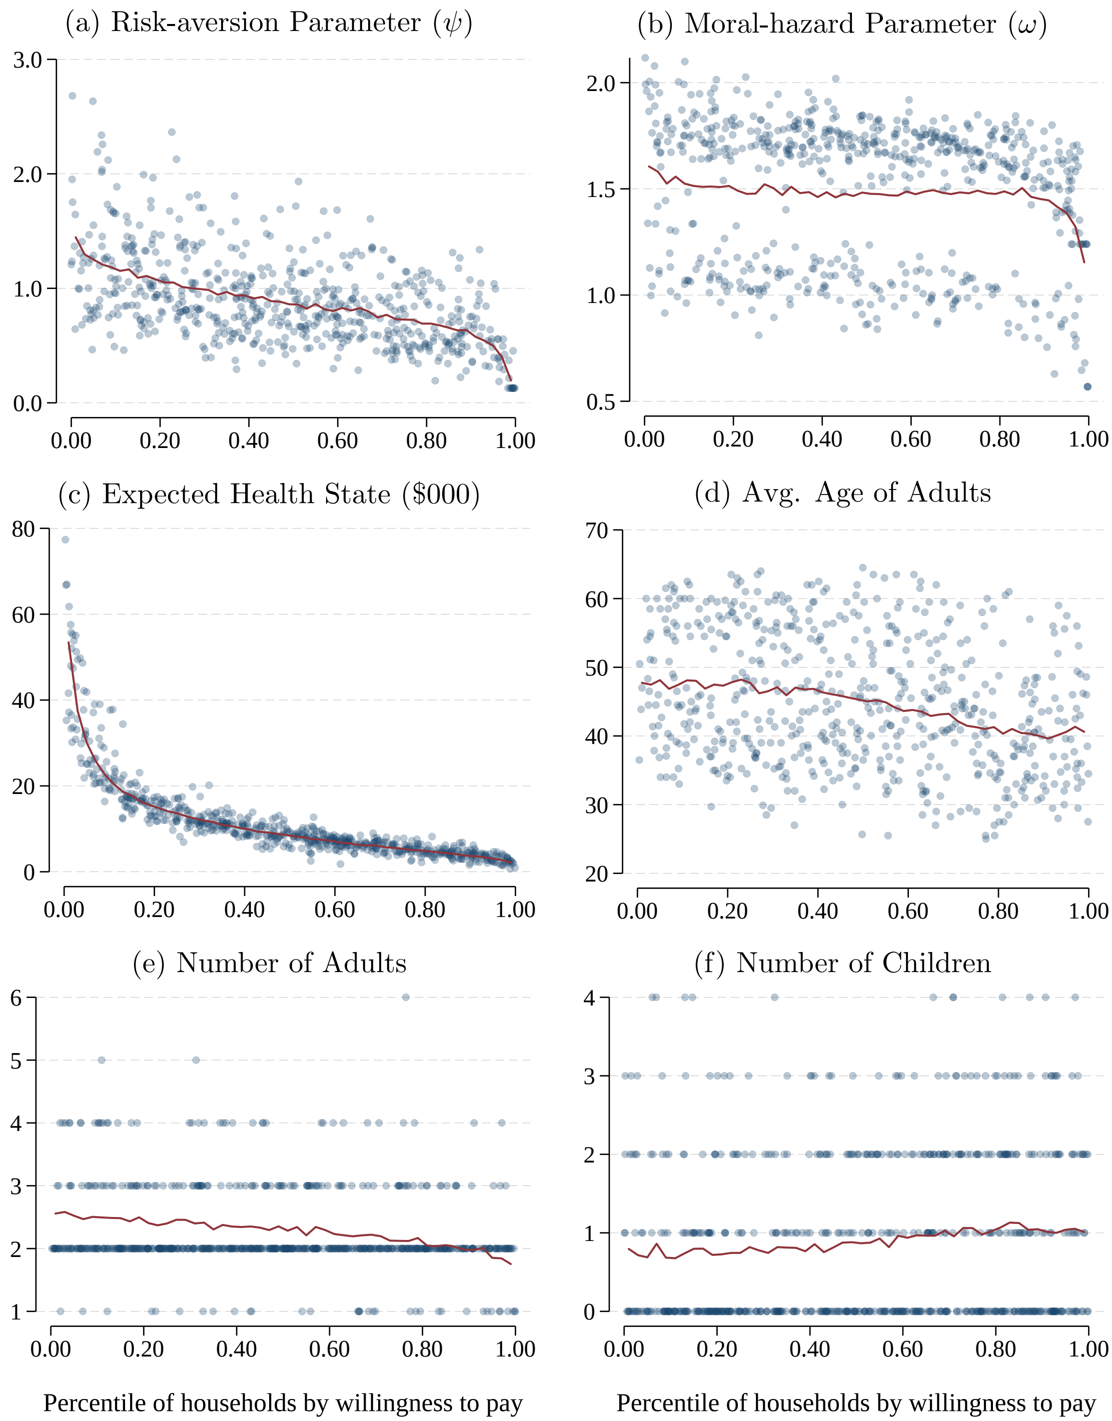

Notes: The figure shows the distribution across family households of (a) the risk aversion parameter, (b) the moral hazard parameter, (c) the expectation of the health state distribution, (d) the average age of adults in the household, (e) the number of adults in the household, and (f) the number of children in the household. An adult is defined as anyone 18 and older. Each dot represents a household, for a 2.5 percent random sample of households. The line in each panel is a connected binned scatter plot, representing the mean value of the vertical axis variable at each percentile of willingness to pay. This figure is referenced in Section V.B.

Figure A.6. Household Health State Distributions by Willingness to Pay

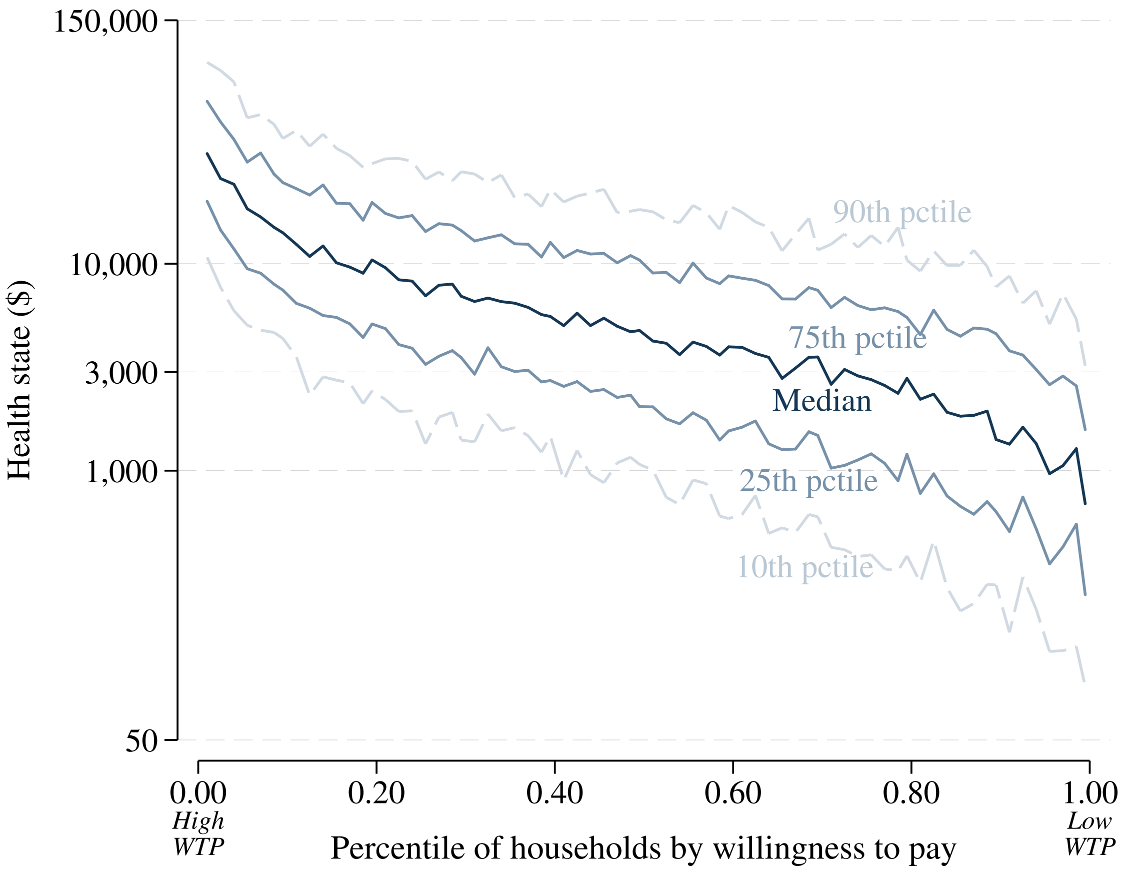

Notes: The figure shows the health state distributions faced by households at each percentile of willingness to pay. Health state distributions are represented by their 10th, 25th, 50th, 75th, and 90th percentiles. A health state realization is equal to total healthcare spending absent moral hazard. The vertical axis is on a log scale. This figure is referenced in Section V.B.

Figure A.7. Efficient Coverage Level by Willingness to Pay

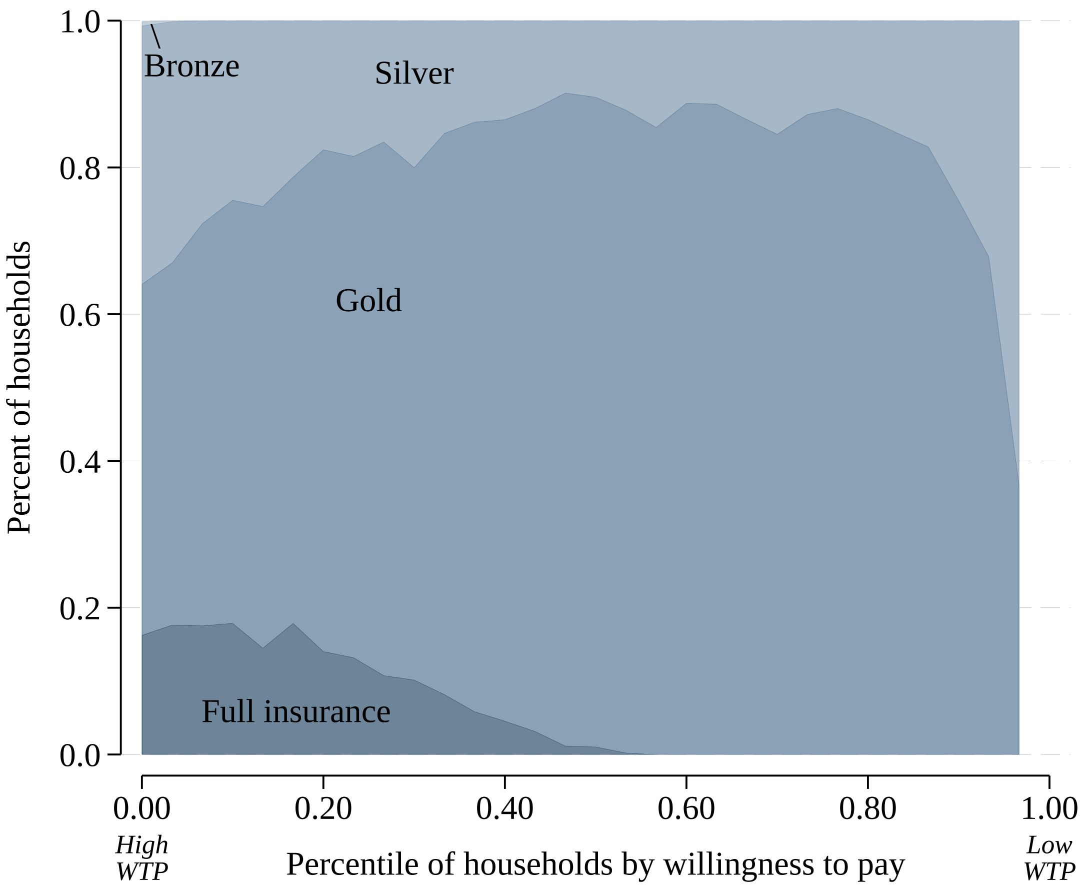

Notes: The figure shows the percentage of family households at each percentile of willingness to pay for whom each contract is optimal. Households are ordered on the horizontal axis according to their willingness to pay. Overall, full insurance is efficient for 6 percent of households, Gold for 75 percent of households, Silver for 19 percent of households, and Bronze for less than one percent of households. Coverage lower than Bronze is not efficient for any household. This figure is referenced in Sections V.B and VI.A.

Table A.9. Outcomes Under Alternative Sets of Potential Contracts

Allocation at First Best (*FB*) and under the Optimal Menu (*Opt*)											
	Metal-tier Contracts						No Deductible				
	Full	Gold	Silv.	Brnz.	Ctstr.		Full	25%	50%	75%	Ctstr.
FB:	0.06	0.75	0.19	<0.01	–	*FB:*	0.31	0.65	0.03	<0.01	–
Opt:	–	1.00	–	–	–	*Opt:*	–	1.00	–	–	–
	No Coinsurance Region						Extended Coins. Region				
	Full	\$2.5k	\$5.0k	\$7.5k	Ctstr.		Full	12.5%	25%	37.5%	50%
FB:	–	0.82	0.17	0.01	–	*FB:*	0.66	0.31	0.01	0.01	–
Opt:	–	1.00	–	–	–	*Opt:*	0.82	0.16	0.02	–	–

Notes: The table shows the percent of households allocated to each contract at the first best allocation (*FB*) and at the optimal feasible allocation (*Opt*), among alternative sets of potential contracts. *Metal-tier Contracts* are the primary set of contracts considered in the main text (and depicted in Fig. A.4a); *No Deductible* are a set of contracts that vary only in their coinsurance rate (see Fig. A.4c); *No Coinsurance Region* are a set of contracts between that vary only in their deductible (see Fig. A.4d); and *Extended Coins. Region* are a set of contracts that have no deductible and vary only in their coinsurance rate, and which have a stop-loss point of \$20,000, twice as high as the other contracts (see Fig. A.4e). This table is referenced in Section V.C.

Table A.10. Parameter Estimates from Full Sample (Including Kaiser)

Variable	Parameter	Std. Err.	Variable	Parameter	Std. Err.
Employee Premium ($000s)	−1.000†		Kaiser * (Age−40)	−0.067	0.006
OOP spending, $-\alpha^{OOP}$	−1.429	0.026	Kaiser * **1**[Children]	−1.832	0.141
HRA/HSA contributions, α^{HA}	0.286	0.023	Kaiser * Region 1	−4.790	0.135
Vision/dental contributions, α^{VD}	1.285	0.024	Kaiser * Region 2	−7.930	0.323
Plan inertia intercept, γ^{plan}	5.119	0.065	Providence * (Age−40)	−0.047	0.007
Plan inertia * **1**[Children], γ^{plan}	−0.154	0.040	Providence * **1**[Children]	−0.629	0.151
Kaiser insurer inertia	9.750	0.262	Providence * Region 1	−1.655	0.132
Moda/Prov. insurer inertia,γ^{ins}	0.392	0.232	Providence * Region 2	−2.259	0.186
Insurer inertia * Risk score, γ^{ins}	0.553	0.073	Providence * Region 3	−1.551	0.213
Moda-specific inertia, 2013	2.162	0.199	κ * Risk Q_1	0.157	0.000
Narrow net. plan, $\nu^{NarrowNet}$	−2.639	0.166	κ * Risk Q_2	0.204	0.000
Kaiser utiliz. multiplier, ϕ_K	0.853	0.008	κ * Risk Q_3	0.188	0.000
Providence utiliz. multiplier, ϕ_P	1.118	0.001	κ * Risk Q_4	0.146	0.016
Risk aversion intercept, $\boldsymbol{\beta}^{\psi}$	−0.872	0.109	κ * Risk $Q_{n<4}$ * Risk score	0.005	0.000
Risk aversion * **1**[Children], $\boldsymbol{\beta}^{\psi}$	−0.096	0.071	κ * Risk Q_4 * Risk score	0.259	0.013
Moral hazard intercept, $\boldsymbol{\beta}^{\omega}$	1.160	0.002	μ * **1**[Female 18–35]	0.097	0.015
Moral hazard * **1**[Children], $\boldsymbol{\beta}^{\omega}$	0.425	0.000	μ * **1**[Age < 18]	0.018	0.015
Std. dev. of private health info., σ_μ	0.184	0.004	μ * Risk Q_1	−0.399	0.019
Std. dev. of log risk aversion, σ_ψ	0.621	0.064	μ * Risk Q_2	0.326	0.010
Std. dev. of moral hazard, σ_ω	0.097	0.001	μ * Risk Q_3	0.449	0.008
Corr(μ, ψ), $\rho_{\mu,\psi}$	0.373	0.004	μ * Risk Q_4	1.245	0.014
Corr(ψ, ω), $\rho_{\psi,\omega}$	−0.252	0.032	μ * Risk $Q_{n<4}$ * Risk score	1.127	0.018
Corr(μ, ω), $\rho_{\mu,\omega}$	0.135	0.007	μ * Risk Q_4 * Risk score	0.339	0.004
Scale of idiosyncratic shock, σ_ϵ	2.519	0.028	σ * Risk Q_1	1.431	0.008
			σ * Risk Q_2	1.240	0.004
			σ * Risk Q_3	1.191	0.003
			σ * Risk Q_4	1.031	0.004
Number of observations: 451,268					

Notes: The table presents parameter estimates using the full sample of households. The specification corresponds to column 3 of Tables 3 and A.8, with two exceptions: (i) insurer inertia terms are estimated separately for Kaiser and for Moda/Providence, and (ii) the moral hazard parameter ω is estimated only among Moda/Providence plans, as opposed to among all three insurers. We note that though it would be interesting to also consider a Kaiser-specific ω, limited variation in coverage level among Kaiser plans prevents us from estimating it. Any Kaiser-specific effects of coverage level on utilization are absorbed into the utilization multiplier ϕ_K. Standard errors are derived from the analytical Hessian of the likelihood function. The model is estimated on an unbalanced panel of 44,562 households, 14 plans, and 5 years. "Risk Q_n" is an indicator for an individual's risk quartile, where Q_4 is the sickest individuals. Higher risk scores correspond to worse predicted health. All parameters are measured in thousands of dollars. The insurer fixed effect of Moda is normalized to zero, and the utilization multiplier for Moda (ϕ_M) is normalized to one. This table is referenced in Section V.C. †By normalization.

Figure A.8. Results from Full Sample Parameter Estimates (Including Kaiser)

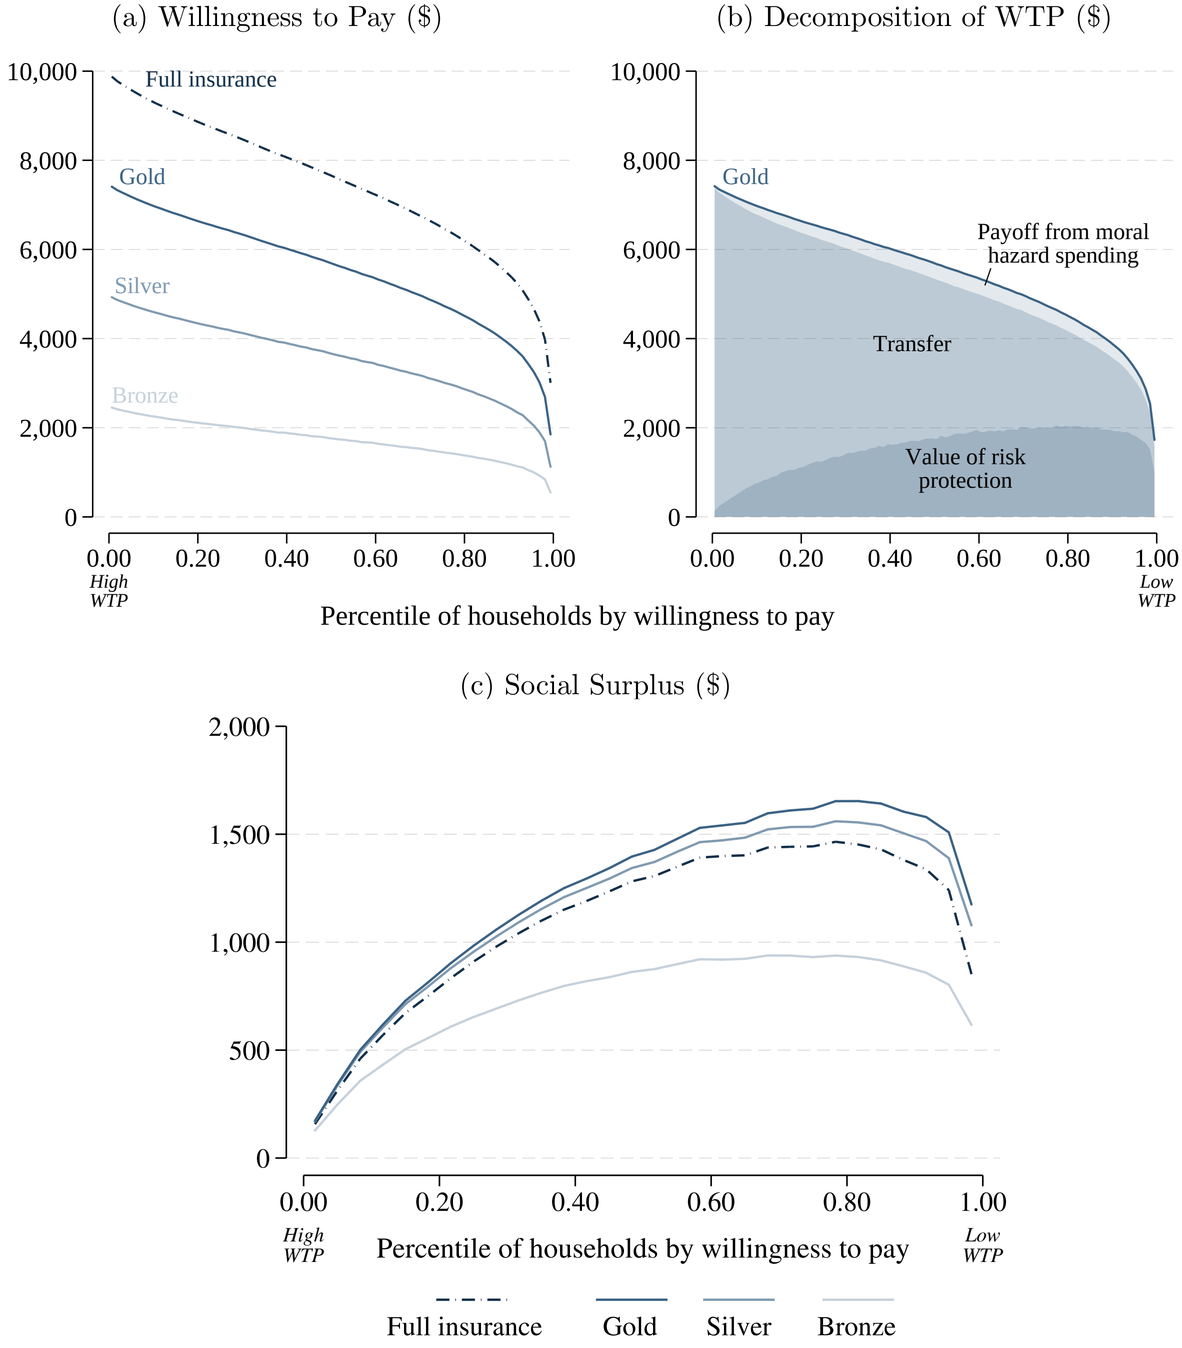

Notes: The figure shows the distribution across family households of (a) willingness to pay, (b) the decomposition of willingness to pay for the Gold contract, and (c) social surplus, using parameter estimates derived from the full sample of households (see Table A.10). The objects in all three panels are measured relative to the Catastrophic contract. Panel (a) consists of four connected binned scatter plots, with respect to 100 quantiles of households ordered by willingness to pay. Panel (b) consists of three connected binned scatter plots, with the area between each line shaded to indicate the component represented. Panel (c) consists of four connected binned scatter plots, with respect to 50 (to reduce noise) quantiles of households. This figure is referenced in Section V.C.

Table A.11. Outcomes Under Different Distributions of Consumer Types

		Outcomes at First Best (*FB*) and at the Optimal Menu (*Opt*), among:							
		Metal-tier contracts						Dense contracts	
Parameter Estimates		Full	Gold	Silv.	Brnz.	Ctstr.	*SS* ($)	Offer choice?	Δ *SS* ($)
Main estimates	*FB:*	0.06	0.75	0.19	<0.01	–	1,542		34
	Opt:	–	1.00	–	–	–	1,514	Yes	14
1. Double mean ω	*FB:*	–	0.29	0.64	0.07	–	1,091		42
	Opt:	–	–	1.00	–	–	1,069	Yes	4
2. Halve mean ω	*FB:*	0.39	0.61	<0.01	–	–	1,855		10
	Opt:	0.61	0.39	–	–	–	1,842	Yes	11
3. Double mean ψ	*FB:*	0.30	0.68	0.02	–	–	2,184		18
	Opt:	0.46	0.54	–	–	–	2,162	Yes	15
4. Halve mean ψ	*FB:*	–	0.35	0.63	0.02	<0.01	919		18
	Opt:	–	–	0.98	–	0.02	915	Yes	2
5. Increase var. ω	*FB:*	0.07	0.74	0.18	0.01	–	1,539		33
	Opt:	–	1.00	–	–	–	1,531	Yes	9
6. Increase var. ψ	*FB:*	0.13	0.64	0.21	0.02	<0.01	1,487		30
	Opt:	0.04	0.76	0.19	0.01	–	1,463	Yes	16
7. Fix F	*FB:*	0.06	0.83	0.11	–	–	1,410		17
	Opt:	–	1.00	–	–	–	1,407	Yes	6
8. Fix F and ω	*FB:*	0.16	0.67	0.17	–	–	1,457		14
	Opt:	0.14	0.68	0.18	–	–	1,456	Yes	12
9. Fix F and ψ	*FB:*	0.17	0.72	0.11	–	–	1,568		16
	Opt:	–	1.00	–	–	–	1,559	No	4

Notes: The table shows results under nine perturbations of our estimates, as well as under our main estimates (column 3 of Tables 3 and A.8). Two sets of results are shown. First, the table shows the percent of households assigned to each of the five metal-tier contracts (c.f. Figure A.4a) under the first best allocation (*FB*) and under the optimal feasible allocation (*Opt.*), as well as the social surplus (*SS*) achieved by those allocations, relative to allocating all households to the Catastrophic contrat. Second, the table indicates whether or not the optimal menu features a choice when considering a dense set of contracts (c.f. Figure A.4b), as well as the associated social surplus gains achieved by the the dense set contracts (Δ *SS*). The nine perturbation of estimates are as follows: (1) double the moral hazard parameter ω for all households; (2) halve ω for all households; (3) double the risk aversion parameter ψ for all households; (4) halve ψ for all households; (5) double the amount of unobserved heterogeneity in moral hazard σ_ω; (6) double the amount of unobserved heterogeneity in log risk aversion σ_ψ; (7) fix household health type F in the population; (8) fix both health F and the moral hazard parameter ω in the population; and (9) fix both health F and risk aversion ψ in the population. This table is referenced in Section V.C.

Figure A.9. Distribution of Consumer Surplus ($), Relative to "All Full Insurance"

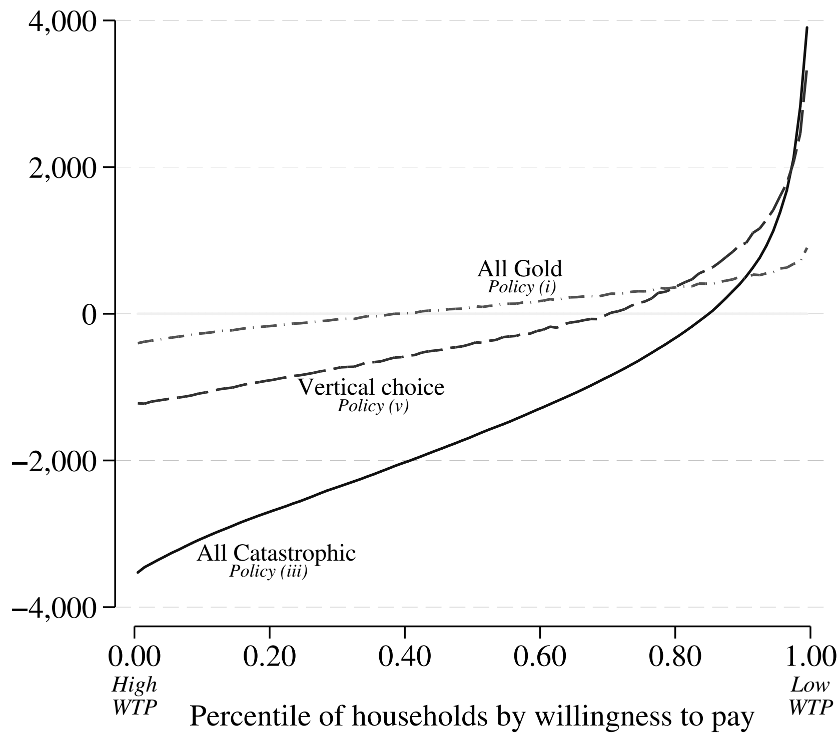

Notes: The figure shows the distribution of consumer surplus across households under three policies considered in Table 4. Households are arranged on the horizontal axis according to their willingness to pay. Consumer surplus equals marginal willingness to pay less marginal premium-plus-tax, *relative to the allocation of all households to full insurance.* That is, a policy of "All Full Insurance" would be represented by a horizontal line at zero. The premium-plus-tax that supports the single contract is $6,298 under "All Catastrophic," $10,619 under "All Gold," and $12,695 under "All full insurance." Premiums under "Vertical choice" are $7,059 for Full insurance, $4,594 for Gold, $2,173 for Silver, $375 for Bronze, $0 for Catastrophic, and a tax of $6,856. This figure is referenced in Section VI.B.