Online Supplement

Equilibria in Health Exchanges: Adverse Selection vs. Reclassification Risk

Ben Handel[*], Igal Hendel[†], and Michael D. Whinston[‡]

January 7, 2015

Abstract

This Online Supplement has four sections. The first investigates Wilson equilibria as an alternative to Riley equilibria and Nash equilibria. The second describes in detail our cost model that predicts consumer health risk. The third describes our extension when an individual mandate is not fully enforced, and consumers can opt out of the market. The fourth describes our extension where we reweight our main sample using weights derived from the nationally representative MEPS survey, and reproduce our main positive and normative analyses.

[*]Department of Economics, UC Berkeley; handel@berkeley.edu

[†]Department of Economics, Northwestern University; igal@northwestern.edu

[‡]Department of Economics and Sloan School of Management, M.I.T.; whinston@mit.edu

1 Supplement: Wilson Equilibria

1.1 Characterization of Wilson Equilibria

A price configuration $P = (P_H, P_L)$ is a Wilson equilibrium (WE) if there is no deviation by an entrant to a price pair that is strictly profitable once any offers are withdrawn that make losses after the deviation.[1] We will say that a deviation from price configuration P that is strictly profitable after any such withdrawals is a "profitable Wilson deviation."[2] Note that no policy L offers will ever be withdrawn after a deviation, because a reduction in P_H can never cause a P_L offer to make losses (since a reduction in P_H lowers AC_L).

We establish the following result, which we use to identify WE in our data:

Proposition S1. *Let $(\underline{P}_H^{BE}, \underline{P}_L^{BE})$ be the break-even price configuration associated with $\underline{\Delta P}^{BE}$, and let $\Delta P^w \in Arg\max_{\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]} \Pi(\underline{P}_L^{BE} + \Delta P, \underline{P}_L^{BE})$. If $\Delta AC(\underline{\theta}) > \underline{\theta}$, then the break-even price configuration (P_H^w, P_L^w) associated with price difference ΔP^w is a WE, and is the unique WE whenever $\Delta P^w = Arg\max_{\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]} \Pi(\underline{P}_L^{BE} + \Delta P, \underline{P}_L^{BE}) \in (\underline{\theta}, \underline{\Delta P}^{BE})$. If instead $\Delta AC(\underline{\theta}) < \underline{\theta}$, then the unique WE outcome has all consumers purchasing policy H at price $P_H^* = \underline{AC}_H$.*

We establish Proposition S1 through a series of lemmas. First, we identify some properties that any WE must satisfy:

Lemma S1. *If $P^w = (P_H^w, P_L^w)$ is a WE price configuration, then*

(a) $\Pi(P_H^w, P_L^w) = 0$;

(b) $\Pi_H(P_H', P_L^w) \leq 0$ *for all* $P_H' \leq P_H^w$*;*

(c) $\Delta P^w = (P_H^w - P_L^w) \leq \underline{\Delta P}^{BE}$*, the lowest break-even ΔP with positive sales of policy L.*

Proof. (a) If $\Pi(P_H^w, P_L^w) < 0$, then some firm would be better off dropping its offers, while if $\Pi(P_H^w, P_L^w) > 0$ then an entrant could profit by offering $(P_H^w - \varepsilon, P_L^w - \varepsilon)$ for sufficiently small $\varepsilon > 0$. (b) If this is violated at P_H', then $\Pi_H(P_H' - \varepsilon, P_L^w) > 0$ for sufficiently small $\varepsilon > 0$. A entrants' offering of $P_H' - \varepsilon$ would be a profitable Wilson deviation. (c) This is immediate if $\underline{\Delta P}^{BE} = \overline{\theta}$. So suppose that $\underline{\Delta P}^{BE} < \overline{\theta}$ and that $\Delta P^w > \underline{\Delta P}^{BE}$, which implies that there are positive sales of policy L at P^w. Since both policies break even at $\underline{\Delta P}^{BE}$, and $\Pi_L(P_H^w, P_L^w) \geq 0$ by parts (a) and (b), it must be that the break-even price configuration associated with $\underline{\Delta P}^{BE}$, $(\underline{P}_H^{BE}, \underline{P}_L^{BE})$, has $\underline{P}_L^{BE} = AC_L(\underline{\Delta P}^{BE}) < AC_L(\Delta P^w) \leq P_L^w$. Since $\underline{P}_L^{BE} < P_L^w$ and $\underline{\Delta P}^{BE} < \Delta P^w$, we also have $\underline{P}_H^{BE} < P_H^w$. So an entrant's offer of $(\underline{P}_H^{BE} + \varepsilon, \underline{P}_L^{BE} + \varepsilon)$ for sufficiently small $\varepsilon > 0$ is a profitable Wilson deviation. $\square$

[1] Note that since at least one of P_H and P_L is undercut by any profitable entrant deviation, there is no ambiguity about which polices to withdraw in the event that one of the offers in the price configuration makes losses.

[2] Note that, in principle, a NE need not be a WE, as a profitable Wilson deviation may not be profitable if no policies are withdrawn.

Consider the following problem:

$$
\begin{array}{lc}
\min_{(P_H,P_L)} & P_L \\
\text{s.t.} & \text{(i) } \Pi(P_H,P_L)=0 \\
 & \text{(ii) } \Pi_H(P'_H,P_L)\le 0 \text{ for all } P'_H\le P_H \\
 & \text{(iii) } P_H-P_L\in[\underline{\theta},\underline{\Delta P}^{BE}]
\end{array}
\tag{1}
$$

Lemma S2. *Any $P^*=(P^*_H,P^*_L)$ that solves problem (1) is a WE price configuration.*

Proof. We construct an equilibrium in which all prices $P\ge P^*$ are offered by multiple firms and each firm has an equal share of sales of both policies. Thus, all active firms earn zero, and we need only consider deviations by entrants.

To begin, it follows from constraint (ii) of problem (1), and the fact that L offers are never withdrawn, that there is no profitable Wilson deviation in which an entrant makes sales only of the H policy (which would require a price $\widehat{P}_H<P^*_H$).

Next, there is no profitable Wilson deviation in which an entrant makes sales only of policy L. Suppose there were and let the deviation price be $\widehat{P}_L<P^*_L$. If everyone buys policy L at prices $(P^*_H,\widehat{P}_L)$ then no policy H offers will be withdrawn and $\widehat{P}_L>\overline{AC}_L$. But then prices $(P^*_H,\overline{AC}_L)$ would be feasible in problem (1) and attain a lower value of P_L than P^*_L, contradicting P^* being a solution. Suppose instead that some consumers still buy policy H at prices $(P^*_H,\widehat{P}_L)$. Then $\Pi_H(P^*_H,\widehat{P}_L)<0$, which implies that offer P^*_H will be withdrawn, as will every P_H up to the lowest $\overline{P}_H$ above P^*_H such that $\Pi_H(\overline{P}_H,\widehat{P}_L)=0$. The entrant's profit will therefore be $\Pi_L(\overline{P}_H,\widehat{P}_L)$. However, it cannot be that $\Pi_L(\overline{P}_H,\widehat{P}_L)>0$: if so then we have $\Pi(\overline{P}_H,\widehat{P}_L)>0$. But this would imply that there is an $\delta>0$ such that price pair $(\overline{P}_H-\delta,\widehat{P}_L-\delta)$ is feasible in problem (1) and achieves a lower P_L than P^*_L, a contradiction to P^* solving problem (1).[3]

Finally, suppose that there is a profitable Wilson deviation for an entrant offering $\widehat{P}=(\widehat{P}_H,\widehat{P}_L)$, in which the entrant makes sales of both policies. Then since offers for policy L are never withdrawn, $\widehat{P}_L\le P^*_L$. We first argue that $\Pi_H(P_H,\widehat{P}_L)\le 0$ for all $P_H\le\widehat{P}_H$. If $\widehat{P}_H<P^*_H$, then this follows because P^* satisfies constraint (ii) and $\widehat{P}_L\le P^*_L$. If, instead, $\widehat{P}_H>P^*_H$, then it follows because the entrant can make sales of the H policy only if $\Pi_H(P_H,\widehat{P}_L)<0$ for all $P_H<\widehat{P}_H$, so that rivals' offers are withdrawn. Next, observe that if $\Pi_H(\widehat{P}_H,\widehat{P}_L)\le 0$ and $\Pi(\widehat{P}_H,\widehat{P}_L)>0$, then for some $\delta>0$ price pair $(\widehat{P}_H-\delta,\widehat{P}_L-\delta)$ is feasible in problem (1) and achieves a lower P_L than P^*_L, a contradiction to P^* solving problem (1). $\square$

To solve for the Wilson equilibrium, we examine a relaxed version of problem (1). For $\Delta P\in[\underline{\theta},\overline{\theta}]$, we first define $P^{BE}_L(\Delta P)$ by

$$[P^{BE}_L(\Delta P)-AC_L(\Delta P)]F(\Delta P)+[P^{BE}_L(\Delta P)+\Delta P-AC_H(\Delta P)][1-F(\Delta P)]=0,$$

[3] This δ would set $\Pi(\overline{P}_H-\delta,\widehat{P}_L-\delta)=0$, and would satisfy constraint (ii) of problem (1) since $\Pi_H(P_H,\widehat{P}_L-\delta)\le 0$ for all $P_H\le\overline{P}_H$.

and $P_H^{BE}(\Delta P) \equiv P_L^{BE}(\Delta P) + \Delta P$. Note that $P_L^{BE}(\Delta P)$ and $P_H^{BE}(\Delta P)$ are continuous functions. Note as well that, for $\Delta P \in [\underline{\theta}, \overline{\theta}]$, $[P_L^{BE}(\Delta P) - AC_L(\Delta P)] \gtreqless 0$ if and only if $\Delta AC(\Delta P) \gtreqless \Delta P$.[4]

We will consider the relaxed problem

$$\min_{\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]} \quad P_L^{BE}(\Delta P) \tag{2}$$

Note that in problem (2) the constraint set is closed and bounded, and the objective function is continuous, so a solution exists. In Lemma S3, we show the equivalence of this problem when $\Delta AC(\underline{\theta}) > \underline{\theta}$ to the problem of finding the profit-maximizing multi-policy Nash deviation from price configuration $(\underline{P}_H^{BE}, \underline{P}_L^{BE})$:

$$\max_{\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]} \quad \Pi(\underline{P}_L^{BE} + \Delta P, \underline{P}_L^{BE}) \tag{3}$$

Lemma S3. *Suppose that $\Delta AC(\underline{\theta}) > \underline{\theta}$. Then $Arg\min_{\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]} P_L^{BE}(\Delta P) = Arg\max_{\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]} \Pi(\underline{P}_L^{BE} + \Delta P, \underline{P}_L^{BE})$.*

Proof. Letting $\delta(\Delta P) \equiv \underline{P}_L^{BE} - P_L^{BE}(\Delta P)$, we have

$$\begin{aligned} \Pi(\underline{P}_L^{BE} + \Delta P, \underline{P}_L^{BE}) &= \Pi(P_L^{BE}(\Delta P) + \Delta P + \delta(\Delta P), P_L^{BE}(\Delta P) + \delta(\Delta P)) \\ &= \Pi(P_L^{BE}(\Delta P) + \Delta P, P_L^{BE}(\Delta P)) + \delta(\Delta P) \\ &= \underline{P}_L^{BE} - P_L^{BE}(\Delta P), \end{aligned}$$

so for any ΔP and $\Delta P'$ we have

$$\Pi(\underline{P}_L^{BE} + \Delta P, \underline{P}_L^{BE}) - \Pi(\underline{P}_L^{BE} + \Delta P', \underline{P}_L^{BE}) = P_L^{BE}(\Delta P') - P_L^{BE}(\Delta P).$$

□

Thus, the solution to the relaxed problem (3) is exactly the $\Delta P \leq \underline{\Delta P}^{BE}$ that maximizes the multi-policy deviation profits from $\underline{\Delta P}^{BE}$. By Lemma 1, any solution to problem 3 for which there is a price configuration (P_H, P_L) with $P_H - P_L = \Delta P$ that is feasible in problem 1 is a WE. This is the case whenever the solution to problem 3 is $\Delta P = \underline{\Delta P}^{BE}$.

The usefulness of the relaxed problems (2) and (3) also stems from the following result:

Lemma S4. *Suppose that $\Delta AC(\underline{\theta}) > \underline{\theta}$ and that $\Delta P^* = \arg\min_{\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]} P_L^{BE}(\Delta P)$. Then the price configuration $(P_H^{BE}(\Delta P^*), P_L^{BE}(\Delta P^*))$ is the unique solution to problem (1).*

[4] This follows because

$$\Delta AC(\Delta P) \gtreqless \Delta P \Leftrightarrow P_L^{BE}(\Delta P) - AC_L(\Delta P) \gtreqless P_H^{BE}(\Delta P) - AC_H(\Delta P).$$

Proof. By Lemma S3, we need only show that $(P_H^{BE}(\Delta P^*), P_L^{BE}(\Delta P^*))$ is feasible in problem (1). By construction $(P_H^{BE}(\Delta P^*), P_L^{BE}(\Delta P^*))$ satisfies constraints (i) and (iii) of problem (1). We therefore need only show that $(P_H^{BE}(\Delta P^*), P_L^{BE}(\Delta P^*))$ satisfies constraint (ii). Observe that when $\Delta AC(\underline{\theta}) > \underline{\theta}$, at any $\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]$ we have $\Delta AC(\Delta P) > \Delta P$. This implies that for all $\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]$

$$\Pi_H(P_H^{BE}(\Delta P), P_L^{BE}(\Delta P)) \leq 0.$$

Since $P_L^{BE}(\Delta P^*) \leq P_L^{BE}(\Delta P)$ for all $\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]$ by virtue of ΔP^* being the solution to problem (2), we therefore have $\Pi_H(P_H^{BE}(\Delta P), P_L^{BE}(\Delta P^*)) \leq 0$ for all $\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]$. Continuity of $P_H^{BE}(\Delta P)$ in ΔP then implies that

$$\Pi_H(P_H, P_L^{BE}(\Delta P^*)) \leq 0 \text{ for all } P_H \in [\underline{AC}_H, P_H^{BE}(\Delta P^*)].$$

Since we also have that

$$\Pi_H(P_H, P_L^{BE}(\Delta P^*)) \leq 0 \text{ for all } P_H \leq \underline{AC}_H,$$

$(P_H^{BE}(\Delta P^*), P_L^{BE}(\Delta P^*))$ satisfies constraint (ii) of problem (1). □

We next show that, when $\Delta AC(\underline{\theta}) \neq \underline{\theta}$, a unique solution P^* to problem (1) is the *only* WE whenever $\Delta P^* \in (\underline{\theta}, \underline{\Delta P}^{BE})$.

Lemma S5. *Suppose that $\Delta AC(\underline{\theta}) > \underline{\theta}$, there is a unique solution P^* of problem (1), and that $\Delta P^* \in (\underline{\theta}, \underline{\Delta P}^{BE})$. Then P^* is the unique WE price configuration.*[5]

Proof. Lemma S1 shows that any WE price configuration must satisfy the constraints of problem (1). We next argue that when P^* is the unique solution to problem (1), any price configuration $\widetilde{P} = (\widetilde{P}_H, \widetilde{P}_L)$ that satisfies the constraints but is not a solution cannot be a WE price configuration. By definition, $P_L^* < \min\{\underline{AC}_H - \underline{\theta}, \widetilde{P}_L\}$. [6]

If $(P_H^*, P_L^*) << (\widetilde{P}_H, \widetilde{P}_L)$ then at price configuration $(\widetilde{P}_H, \widetilde{P}_L)$ an entrant has a profitable Wilson deviation to $(P_H^* + \varepsilon, P_L^* + \varepsilon)$ for small $\varepsilon > 0$. So, for the rest of the proof, suppose instead that $P_H^* \geq \widetilde{P}_H$, which also implies that $\Delta\widetilde{P} < \Delta P^*$ since $P_L^* < \widetilde{P}_L$.

We will show that

$$\Pi_H(P_H, P_L^*) < 0 \text{ for all } P_H \in (\widetilde{P}_H, P_H^*] \tag{4}$$

which will imply that at $\widetilde{P}$ an entrant has a profitable Wilson deviation offering prices $(P_H^* + \varepsilon, P_L^* + \varepsilon)$ for $\varepsilon > 0$ such that $\Pi_H(P_H', P_L^* + \varepsilon) < 0$ for all $P_H' \in [\widetilde{P}_H, P_H^* + \varepsilon]$ and $P_L^* + \varepsilon < \widetilde{P}_L$, which results in all H policy offers in $[\widetilde{P}_H, P_H^* + \varepsilon]$ being withdrawn.

Condition (4) follows immediately if $P_H^* \leq \underline{AC}_H$, since $P_H - P_L^* < \Delta P^*$ for all $P_H \leq P_H^*$ implies that there are positive sales of policy H at price configuration (P_H, P_L^*). So suppose henceforth that $P_H^* > \underline{AC}_H$.

[5] We conjecture, but have not proven, that the result extends to cases in which $\Delta P^* = \underline{\Delta P}^{BE}$.

[6] The inequality $P_L^* < \underline{AC}_H - \underline{\theta}$ holds because the price configuration $(\underline{AC}_H, \underline{AC}_H - \underline{\theta})$, which results in all consumers choosing policy H, is feasible in problem (1), but is not the solution.

Because $\Delta AC(\underline{\theta}) > \underline{\theta}$, we have $\Delta AC(\Delta P) > \Delta P$ for all $\Delta P \in (\underline{\theta}, \underline{\Delta P}^{BE})$, which implies that $\Pi_H(P_H^{BE}(\Delta P), P_L^{BE}(\Delta P)) < 0$ for all $\Delta P \in (\underline{\theta}, \underline{\Delta P}^{BE})$. Moreover, continuity of $P_H^{BE}(\cdot)$ implies that for each $P_H \in [\underline{AC}_H, P_H^*]$, there is a $\Delta P' \in (\underline{\theta}, \underline{\Delta P}^{BE})$ such that $P_H^{BE}(\Delta P') = P_H$. Thus, we have

$$\Pi_H(P_H, P_L^*) < 0 \text{ for all } P_H \in (\underline{AC}_H, P_H^*] \tag{5}$$

since there are positive sales of policy H at price configuration (P_H, P_L^*) and

$$P_H = P_H^{BE}(\Delta P') < AC_H(P_H^{BE}(\Delta P') - P_L^{BE}(\Delta P')) < AC_H(P_H^{BE}(\Delta P') - P_L^*) = AC_H(P_H - P_L^*),$$

[the first inequality follows because $\Pi_H(P_H^{BE}(\Delta P'), P_L^{BE}(\Delta P')) < 0$ and the last inequality follows because P^* being the solution to problem (1) implies that $P_L^* < P_L^{BE}(\Delta P')$]. If $\widetilde{P}_H > \underline{AC}_H$, then this establishes (4).

Finally, suppose that $\widetilde{P}_H \leq \underline{AC}_H$. Then

$$\Pi_H(P_H, P_L^*) < 0 \text{ for all } P_H \in [\widetilde{P}_H, \underline{AC}_H] \tag{6}$$

since $P_H - P_L^* < \Delta P^*$ implies that there are positive sales of policy H at price configuration (P_H, P_L^*) and $P_H - P_L^* > \overline{\theta}$ if $P_H = \underline{AC}_H$ (since $P_L^* < \underline{AC}_H - \underline{\theta}$) implies that there are positive sales of policy L at price configuration $(\underline{AC}_H, P_L^*)$. Then (5) and (6) together imply (4). □

Finally, for the case where $\Delta AC(\underline{\theta}) < \underline{\theta}$ we have the following result:

Lemma S6. *Suppose that $\Delta AC(\underline{\theta}) < \underline{\theta}$. Then the unique WE outcome has all consumers purchasing policy H at price $P_H^* = \underline{AC}_H$.*

Proof. In Lemma 4 in the main text, we established that all-in-90 is a RE when $\Delta AC(\underline{\theta}) < \underline{\theta}$ by arguing that when $P_H^* = \underline{AC}_H$ and $P_L^* > \underline{AC}_H - \underline{\theta}$, there was no profitable single-policy deviation in P_L. Since the exit of policy H cannot make policy L profitable (it would raise AC_L), there are also no profitable Wilson single-policy deviations in P_L. It is also immediate that there are no profitable Wilson deviations in only P_H or in both prices. The proof of uniqueness follows much as in the proof of Proposition 1 in the main text: Now any WE P^{**} must have $P_L^{**} \geq AC_L(\Delta P^{**})$ by Lemma S1. A single policy deviation to $\widehat{P}_H = P_L^{**} + \underline{\theta}$ attracts all consumers to policy H. Since $\widehat{P}_H > AC_L(\Delta P^{**}) + \underline{\theta}$ $> \underline{AC}_L + \underline{\theta} > \underline{AC}_H$ it is profitable absent any exit of policy L, and policy L will not exit since it does not make losses (and could not make policy H unprofitable even if it did exit). □

Remark 1. *Proposition S1 implies that [provided that $\Delta AC_H(\underline{\theta}) \neq \underline{\theta}$] any NE is a WE, and that whenever WE and RE outcomes coincide, they are also a NE outcome. Another implication of our discussion is that WE outcomes weakly Pareto dominate RE outcomes. In particular, when $\Delta AC(\underline{\theta}) < \underline{\theta}$ and $\Delta P^w < \underline{\Delta P}^{BE}$, we have $P_L^w < \underline{P}_L^{BE}$ (by Lemma S4) and $P_L^w < \underline{P}_L^{BE}$ (since $P_L^w < \underline{P}_L^{BE}$ and $\Delta P^w < \underline{\Delta P}^{BE}$).*

Wilson Equilibria: Community Rating and Health Status-based Pricing (Quartiles)						
Market	P_{60}	S_{60}	AC_{60}	P_{90}	S_{90}	AC_{90}
Full Population	4,006	83.7	2,477	7,105	16.3	14,961
Quartile 1	302	60.2	290	1,502	39.8	1,519
Quartile 2	1,307	64.7	1,155	3,307	35.3	3,586
Quartile 3	4,443	70.0	3,337	7,193	30.0	9,648
Quartile 4	9,704	73.6	7,259	13,204	26.4	20,007

Table S1: Equilibrium results for Wilson solution concept for (i) pure community rating (no pre-existing conditions) and (ii) health-based pricing with quartiles.

Welfare Loss from Health-Status-based Pricing (Quartiles) in Wilson Equilibrium ($/year)			
γ	$y_{HB4,no-pre}(\gamma)$ Fixed Income	$y_{HB4,no-pre}(\gamma)$ Non-Manager Income path	$y_{HB4,no-pre}(\gamma)$ Manager Income Path
0.0002	2,101	1,390	-468
0.0003	2,577	1,592	-682
0.0004	2,964	1,711	-950
0.0005	3,277	1,628	-1,076
0.0006	3,506	1,923	-1,050

Table S2: Long-run welfare based on the Wilson Equilibrium results. Compares the two pricing regulations of (i) pricing based on health status quartiles (x = "$HB4$") and (ii) pure community rating / no pre-existing conditions (x' = "$no - pre$").

1.2 Empirical Results for Wilson Equilibria

We identify WE using Proposition S1, focusing on our baseline case of a 90 and a 60 policy. When $\Delta AC(\underline{\theta}) > \underline{\theta}$ (which is the case in our data), the price difference that maximizes the profit from a multi-policy deviation from $(\underline{P}_{90}^{BE}, \underline{P}_{60}^{BE})$, the break-even price configuration associated with $\underline{\Delta P}^{BE}$, is a WE.[7] Table S1 shows the equilibria with community rating and with health status quartile pricing. Wilson equilibrium policies break even in total, but they do so allowing the policy L to cross-subsidize policy H. The cross-subsidization can be seen by comparing the prices to the average costs for each policy. We see that in every population the WE has a positive share of consumers purchasing the 90 policy, in contrast to the RE/sp-NE of Section 4.

Table S2 shows welfare results for WE, which are of a similar order of magnitude to those for RE.

[7] As noted above, this is the unique Wilson equilibrium when $\Delta P^w \in (\underline{\theta}, \underline{\Delta P}^{BE})$. We conjecture, but have not proven that the same is true if $\Delta P^w \in \{\underline{\theta}, \underline{\Delta P}^{BE}\}$.

2 Supplement: Cost Model Setup and Estimation

This appendix describes the details of the cost model, which is summarized at a high-level in Section 3, and similar to that used in Handel (2013). The output of this model, F_{jkt}, is a family-plan-time-specific distribution of predicted out-of-pocket expenditures for the upcoming year. This distribution is an important input into the empirical choice model, where it enters as a family's predictions of its out-of-pocket expenses at the time of plan choice, for each plan option. We predict this distribution in a sophisticated manner that incorporates (i) past diagnostic information (ICD-9 codes) (ii) the Johns Hopkins ACG predictive medical software package (iii) a non-parametric model linking modeled health risk to total medical expenditures using observed cost data and (iv) a detailed division of medical claims and health plan characteristics to precisely map total medical expenditures to out-of-pocket expenses. The level of precision we gain from the cost model leads to more credible estimates of the choice parameters of primary interest (e.g., risk preferences and health risk). Crucially, the cost model output is also used to predict consumer expected average costs for the upcoming year, λ, which is used to determine plan costs (as a function of who selects which plans) in our equilibrium analyses.

In order to predict expenses in a precise manner, we categorize the universe of total medical claims into four mutually exclusive and exhaustive subdivisions of claims using the claims data. These categories are (i) hospital and physician services (ii) pharmacy (iii) mental health and (iv) physician office visits. We divide claims into these four specific categories so that we can accurately characterize the plan-specific mappings from total claims to out-of-pocket expenditures since each of these categories maps to out-of-pocket expenditures in a different manner. We denote this four dimensional vector of claims $\mathbf{C}_{it}$ and any given element of that vector $C_{d,it}$ where $d \in D$ represents one of the four categories and i denotes an individual (employee or dependent). After describing how we predict this vector of claims for a given individual, we return to the question of how we determine out-of-pocket expenditures in plan k given $\mathbf{C}_{it}$.

Denote an individual's past year of medical diagnoses and payments by ξ_{it} and the demographics age and sex by ζ_{it}. We use the ACG software mapping, denoted A, to map these characteristics into a predicted mean level of health expenditures for the upcoming year, denoted θ:

$$A : \xi \times \zeta \rightarrow \theta$$

In addition to forecasting a mean level of total expenditures, the software has an application that predicts future mean pharmacy expenditures. This mapping is analogous to A and outputs a prediction κ for future pharmacy expenses.

We use the predictions θ and κ to categorize similar groups of individuals across each of four claims categories in vector in $\mathbf{C}_{it}$. Then for each group of individuals in each claims category, we use the actual ex post realized claims for that group to estimate the ex ante distribution for each individual under the assumption that this distribution is identical for all individuals within the cell. Individuals are categorized into cells based on different metrics for each of the four elements of $\mathbf{C}$:

Pharmacy:	κ_{it}
Hospital / Physician (Non-OV):	θ_{it}
Physician Office Visit:	θ_{it}
Mental Health:	$C_{MH,i,t-1}$

For pharmacy claims, individuals are grouped into cells based on the predicted future mean pharmacy claims measure output by the ACG software, κ_{it}. For the categories of hospital / physician services (non office visit) and physician office visit claims individuals are grouped based on their mean predicted total future health expenses, θ_{it}. Finally, for mental health claims, individuals are grouped into categories based on their mental health claims from the previous year, $C_{MH,i,t-1}$ since (i) mental health claims are very persistent over time in the data and (ii) mental health claims are generally uncorrelated with other health expenditures in the data. For each category we group individuals into a number of cells between 8 and 10, taking into account the tradeoff between cell size and precision. The minimum number of individuals in any cell is 73 while almost all cells have over 500 members. Thus, since there are four categories of claims, each individual can belong to one of approximately 10^4 or 10,000 combination of cells.

Denote an arbitrary cell within a given category d by z. Denote the population in a given category-cell combination (d, z) by I_{dz}. Denote the empirical distribution of ex-post claims in this category for this population $\hat{G_{I_{dz}}}(\cdot)$. Then we assume that each individual in this cell has a distribution equal to a continuous fit of $\hat{G_{I_{dz}}}(\cdot)$, which we denote G_{dz}:

$$\varpi : \hat{G_{I_{dz}}}(\cdot) \rightarrow G_{dz}$$

We model this distribution continuously in order to easily incorporate correlations across d. Otherwise, it would be appropriate to use $G_{I_{dz}}$ as the distribution for each cell.

The above process generates a distribution of claims for each d and z but does not model correlation over D. It is important to model correlation across claims categories because it is likely that someone with a bad expenditure shock in one category (e.g., hospital) will have high expenses in another area (e.g., pharmacy). We model correlation at the individual level by combining marginal distributions G_{idt} $\forall$ d with empirical data on the rank correlations between pairs (d, d').[8] Here, G_{idt} is the distribution G_{dz} where $i \in I_{dz}$ at time t. Since correlations are modeled across d we pick the metric θ to group people into cells for the basis of determining correlations (we use the same cells that we use to determine group people for hospital and physician office visit claims). Denote these cells based on θ by z_θ. Then for each cell z_θ denote the empirical rank correlation between claims of type d and type d' by $\rho_{z_\theta}(d, d')$.

[8] It is important to use rank correlations here to properly combine these marginal distribution into a joint distribution. Linear correlation would not translate empirical correlations to this joint distribution appropriately.

Then, for a given individual i we determine the joint distribution of claims across D for year t, denoted $H_{it}(\cdot)$, by combining i's marginal distributions for all d at t using $\rho_{z_\theta}(d, d')$:

$$\Psi : G_{iDt} \times \rho_{z_{\theta_{it}}}(D, D') \to H_{it}$$

Here, G_{iDt} refers to the set of marginal distributions $G_{idt} \forall d \in D$ and $\rho_{z_{\theta_{it}}}(D, D')$ is the set of all pairwise correlations $\rho_{z_{\theta_{it}}}(d, d') \forall (d, d'^2$. In estimation we perform Ψ by using a Gaussian copula to combine the marginal distribution with the rank correlations, a process which we describe momentarily.

The final part of the cost model maps the joint distribution H_{it} of the vector of total claims C over the four categories into a distribution of out of pocket expenditures for each plan. For each of the three plan options we construct a mapping from the vector of claims C to out-of-pocket expenditures X_k:

$$\Omega_k : C \to X_k$$

This mapping takes a given draw of claims from H_{it} and converts it into the out-of-pocket expenditures an individual would have for those claims in plan k. This mapping accounts for plan-specific features such as the deductible, co-insurance, co-payments, and out-of-pocket maximums described in the text. We test the mapping Ω_k on the actual realizations of the claims vector C to verify that our mapping comes close to reconstructing the true mapping. Our mapping is necessarily simpler and omits things like emergency room co-payments and out of network claims. We constructed our mapping with and without these omitted categories to insure they did not lead to an incremental increase in precision. We find that our categorization of claims into the four categories in C passed through our mapping Ω_k closely approximates the true mapping from claims to out-of-pocket expenses. Further, we find that it is important to model all four categories described above: removing any of the four makes Ω_k less accurate. See Handel (2013) for figures describing this validation exercise with the data used in this paper.

Once we have a draw of X_{ikt} for each i (claim draw from H_{it} passed through Ω_k) we map individual out-of-pocket expenditures into family out-of-pocket expenditures. For families with less than two members this involves adding up all the within family X_{ikt}. For families with more than three members there are family level restrictions on deductible paid and out-of-pocket maximums that we adjust for. Define a family j as a collection of individuals i_j and the set of families as J. Then for a given family out-of-pocket expenditures are generated:

$$\Gamma_k : X_{i_j,kt} \to X_{jkt}$$

To create the final object of interest, the family-plan-time specific distribution of out of pocket expenditures $F_{jkt}(\cdot)$, we pass the claims distributions H_{it} through Ω_k and combine families through Γ_k. $F_{jkt}(\cdot)$ is then used as an input into the choice model that represents each family's information set over future medical expenses at the time of plan choice. Eventually, we also use H_{it} to calculate total plan cost when we analyze counterfactual plan pricing based on the average cost of enrollees.

We note that the decision to do the cost model by grouping individuals into cells, rather then by specifying a more continuous form, has costs and benefits. The cost is that all individuals within a given cell for a given type of claims are treated identically. The benefit is that our method produces local cost estimates for each individual that are not impacted by the combination of functional form and the health risk of medically different individuals. Also, the method we use allows for flexible modeling across claims categories. Finally, we note that we map the empirical distribution of claims to a continuous representation because this is convenient for building in correlations in the next step. The continuous distributions we generate very closely fit the actual empirical distribution of claims across these four categories.

Cost Model Identification and Estimation. The cost model is identified based on the two assumptions of (i) no moral hazard / selection based on private information and (ii) that individuals within the same cells for claims d have the same ex ante distribution of total claims in that category. Once these assumptions are made, the model uses the detailed medical data, the Johns Hopkins predictive algorithm, and the plan-specific mappings for out of pocket expenditures to generate the final output $F_{jkt}(\cdot)$. These assumptions, and corresponding robustness analyses, are discussed at more length in the main text and in Handel (2013).

Once we group individuals into cells for each of the four claims categories, there are two statistical components to estimation. First, we need to generate the continuous marginal distribution of claims for each cell z in claim category d, G_{dz}. To do this, we fit the empirical distribution of claims $G_{I_{dz}}$ to a Weibull distribution with a mass of values at 0. We use the Weibull distribution instead of the lognormal distribution, which is traditionally used to model medical expenditures, because we find that the lognormal distribution overpredicts large claims in the data while the Weibull does not. For each d and z the claims greater than zero are estimated with a maximum likelihood fit to the Weibull distribution:

$$\max_{(\alpha_{dz},\beta_{dz})} \Pi_{i\in I_{dz}} \frac{\beta_{dz}}{\alpha_{dz}}(\frac{c_{id}}{\alpha_{dz}})^{\beta_{dz}-1}e^{-(\frac{c_{id}}{\alpha_{dz}})^{\beta_{dz}}}$$

Here, $\hat{\alpha_{dz}}$ and $\hat{\beta_{dz}}$ are the shape and scale parameters that characterize the Weibull distribution. Denoting this distribution $W(\hat{\alpha_{dz}},\hat{\beta_{dz}})$ the estimated distribution $\hat{G_{dz}}$ is formed by combining this with the estimated mass at zero claims, which is the empirical likelihood:

$$\hat{G_{dz}}(c) = \begin{cases} G_{I_{dz}}(0) & \text{if } c = 0 \\ G_{I_{dz}}(0) + \frac{W(\hat{\alpha_{dz}},\hat{\beta_{dz}})(c)}{1-G_{I_{dz}}(0)} & \text{if } c > 0 \end{cases}$$

Again, we use the notation $\hat{G_{iDt}}$ to represent the set of marginal distributions for i over the categories d: the distribution for each d depends on the cell z an individual i is in at t. We combine the distributions $\hat{G_{iDt}}$ for a given i and t into the joint distribution H_{it} using a Gaussian copula method for the mapping Ψ. Intuitively, this amounts to assuming a parametric form for correlation across $\hat{G_{iDt}}$ equivalent to that from a standard normal distribution with correlations equal to empirical rank correlations $\rho_{z_{\theta_{it}}}(D, D')$ described in the previous section. Let $\Phi^i_{1|2|3|4}$ denote the standard multivariate normal distribution with pairwise correlations $\rho_{z_{\theta_{it}}}(D, D')$ for all pairings of the four claims categories D. Then an individual's joint distribution of non-zero claims is:

$$\hat{H_{i,t}}(\cdot) = \Phi_{1|2|3|4}(\Phi_1^{-1}(\hat{G_{id_1t}}), \Phi_2^{-1}(\hat{G_{id_2t}}), \Phi_3^{-1}(\hat{G_{id_3t}}), \Phi_4^{-1}(\hat{G_{id_4t}}))))$$

Above, Φ_d is the standard marginal normal distribution for each d. $\hat{H}_{i,t}$ is the joint distribution of claims across the four claims categories for each individual in each time period. After this is estimated, we determine our final object of interest $F_{jkt}(\cdot)$ by simulating K multivariate draws from $\hat{H}_{i,t}$ for each i and t, and passing these values through the plan-specific total claims to out of pocket mapping Ω_k and the individual to family out of pocket mapping Γ_k. The simulated $F_{jkt}(\cdot)$ for each j, k, and t is then used as an input into estimation of the choice model.

Table S3 presents summary results from the cost model estimation for the final choice model sample, including population statistics on the ACG index θ, the Weibull distribution parameters $\hat{\alpha_{dz}}$ and $\hat{\beta_{dz}}$ for each category d, as well as the across category rank correlations $\rho_{z_{\theta_{it}}}(D, D')$. These are the fundamentals inputs used to generate F_{jkt}, as described above, and lead to accurate characterizations of the overall total cost and out-of-pocket cost distributions (validation exercises which are not presented here).

Final Sample				
Cost Model Output				
	Overall	PPO_{250}	PPO_{500}	PPO_{1200}
Individual Mean (Median)				
Unscaled ACG Predictor				
Mean		1.42	0.74	0.72
Median		0.83	0.37	0.37
Pharmacy: Model Output				
Zero Claim Pr.	0.35 (0.37)	0.31 (0.18)	0.40 (0.37)	0.42 (0.37)
Weibull α	1182 (307)	1490 (462)	718 (307)	596 (307)
Weibull β	0.77 (0.77)	0.77 (0.77)	0.77 (0.77)	0.77 (0.77)
Mental Health				
Zero Claim Pr.	0.88 (0.96)	0.87 (0.96)	0.90 (0.96)	0.90 (0.96)
Weibull α	1422 (1295)	1447 (1295)	1374 (1295)	1398 (1295)
Weibull β	0.98 (0.97)	0.99 (0.97)	0.98 (0.97)	0.98 (0.97)
Hospital / Physician				
Zero Claim Pr.	0.23 (0.23)	0.21 (0.23)	0.26 (0.23)	0.26 (0.23)
Weibull α	2214 (1599)	2523 (1599)	1717 (1599)	1652 (1599)
Weibull β	0.58 (0.55)	0.59 (0.55)	0.55 (0.55)	0.55 (0.55)
$(> \$40,000)$ Claim Pr.	0.02 (0.01)	0.02 (0.01)	0.01 (0.01)	0.01 (0.01)
Physician OV				
Zero Claim Pr.	0.29 (0.20)	0.26 (0.20)	0.33 (0.46)	0.34 (0.46)
Weibull α	605 (553)	653 (553)	517 (410)	529 (410)
Weibull β	1.15 (1.14)	1.15 (1.14)	1.15 (1.14)	1.14 (1.14)
Correlations				
Rank Correlation Hospital-Pharm.	0.28 (0.34)	0.26 (0.32)	0.31 (0.34)	0.32 (0.34)
Rank Correlation Hospital-OV	0.73 (0.74)	0.72 (0.74)	0.74 (0.74)	0.74 (0.74)
Rank Correlation Pharm.-OV	0.35 (0.41)	0.33 (0.37)	0.38 (0.41)	0.39 (0.41)

Table S3: This table describes the output of the cost model in terms of the means and medians of individual level parameters, classified by the plan actually chosen. These parameters are aggregated for these groups but have more micro-level groupings, which are the primary inputs into our cost projections in the choice model. Weibull α, Weibull β, and Zero Claim Probability correspond to the cell-specific predicted total individual-level health expenses as described in more detail in this supplement.

3 Supplement: Exchange Participation

Our market analysis in the main text assumed full participation in the market. This could result from, for example, a legally enforced individual mandate (as in the ACA) with a large penalty, or, alternatively, an employer requiring all workers to remain in the insurance pool of a private exchange. In reality, such a requirement may be difficult to enforce, or the penalty for not purchasing insurance may be small, leading to a scenario where certain consumers, especially healthy ones, may prefer to opt out of the market.

To understand the role of mandated participation, we investigate the case where individuals can opt out of the exchanges should their expected utility from being uninsured be higher than joining their favorite insurance plan in the market. Uninsured means that the consumer pays zero premium and pays for the total cost of their health expenses. We again focus on the case of a 90% policy and a 60% policy in the market. We find equilibria allowing individuals to opt out without any penalty.[9]

Recall that equilibria without age-based pricing unraveled to all-in-60. The column "Better-off In" in the "Community Rating" section of Table S4 shows the percentage of each age group (and of the population as a whole) that is better off insured at the equilibrium premium of \$4,068 than remaining uninsured. For example, 44.2% ($= 100 - 55.8$) of 25 to 30 year old individuals prefer to opt out as their expected utility from non-insurance is higher than being pooled with the whole population.

Naturally, those that prefer to opt out are younger, healthier and less risk averse. The expected costs of insuring consumers who prefer to decline coverage is \$3,107 versus \$5,107 for those that prefer to participate. The average risk aversion coefficient of those that prefer to participate is $4.26 * 10^{-4}$ versus $4.03 * 10^{-4}$ for those that prefer to decline coverage.

Allowing healthier individuals to opt out increases the cost of covering the remaining pool, which in turn draws more people out of the pool. The process stops with a RE premium of \$5,339 when no more individuals want to drop out (that is, the RE for the remaining pool has $P_{60} = \$5,339$). The equilibrium without the mandate involves full unraveling to 60, with 74.3% of the population voluntarily covered. The column "No Mandate: Participation" under "Community Rating" shows participation by age in the non-mandate equilibrium.

We can also compute the welfare impact of removing the mandate. Those individuals that remain covered, 74.3% of the population, suffer a loss equal to the premium increase \$1,271 ($= 5,339 - 4,068$). Comparing the certainty equivalent of remaining uninsured versus participation in the exchange for the 25.7% of the population that opts out, we find that they are better off by \$1,972, on average. Thus, removing the mandate entails a welfare loss of \$434.3 [$= 0.743(1,271) - 0.257(1,972)$] per person. On the right side of Table S4 we show the corresponding numbers for age-based pricing. As we saw in Section 6.2 in the main text, all the equilibria under the mandate (with no opting out) for the

[9] More concretely, we find the equilibrium with the mandate, and eliminate from the sample those individuals that are better off uninsured. We then iterate finding equilibria and eliminating the worse off consumers, until all buyers want to remain in the market.

Implications of Individual Mandate						
	Community Rating		Age-Based Pricing			
	Mandate:	No Mandate:	Mandate:		No Mandate:	
Ages	Better-Off In	Participation	Premium	Better-Off In	Premium	Participation
All	78.3%	74.3%	-	80.7%	-	77.0%
25-30	55.8%	50.6%	1,786	70.1%	2,732	63.1%
30-35	59.6%	54.1%	2,215	70.0%	3,409	62.5%
35-40	68.7%	62.2%	2,542	75.9%	3,476	70.8%
40-45	75.1%	70.9%	3,242	77.7%	4,233	74.5%
45-50	82.5%	79.3%	4,103	82.9%	4,976	80.6%
50-55	90.6%	87.2%	5,038	88.6%	5,714	86.9%
55-60	94.7%	92.5%	6,304	92.1%	6,927	89.9%
60-65	95.8%	93.9%	7,259	91.6%	7,959	90.2%

Table S4: Implications of the individual mandate for equilibrium prices and market participation.

different age groups involve unravelling to 60. At the equilibrium premium, reported in the "Mandate: Premium" column, only some of the population would voluntarily participate in the exchange. Column "Mandate: Better-off In," shows that the share that prefers to participate is an increasing share in age. Older individuals are more likely to benefit from participation, but the differences across ages are less pronounced once age is priced.

For each age, as individuals opt out, the cost of coverage increases. The column "No Mandate: Premium" reports the equilibrium premia for each age group absent a mandate. It is substantially higher than under the mandate, especially so for younger cohorts for whom the mandate is binding for a larger proportion of individuals. In a similar fashion we can use the model to study the participation level for different subsidy or penalty levels (analysis available upon request).

4 Supplement: Population Re-Weighting

The analysis in the main text uses health choice and utilization data from a large firm with approximately 10,000 employees and 20,000 covered lives. While these data have a lot of depth on dimensions that are essential to model health risk and risk preferences, they represent a specific population working for a specific large employer. Our results thus represent the case of exchange design as if this population were the population of interest. This could correspond closely to the case where either (i) this large employer (or a similar one) sets up a private exchange or (ii) our population represents a population of general interest for a public exchange (such as the ACA state exchanges). While our analysis thus far is clearly relevant for (i), and conceptually relevant for (ii), it is also likely that our sample is not the same as the sample of interest for policymakers setting up state insurance exchanges under the ACA.

To provide a rough sense of how our results could change under a population more similar to that enrolling in state insurance exchanges under the ACA, we extend the analysis by applying our framework to a more externally relevant sample from the Medical Expenditures Panel Survey (MEPS), which was specifically created to study medical care decisions for a nationally representative population. Column 1 in Table S5 contains the summary statistics for the entire MEPS population during the years we focus on (2004-2008) with no sample cuts (N = 166,539). We analyze exchange equilibria and welfare outcomes using an "ACA relevant" sample composed of individuals in the MEPS data who are (i) between the ages of 25-65 and (ii) either uninsured or covered by a plan on the individual market (N = 21,856). This sample is similar in spirit to the sample that actually enrolls in the state insurance exchanges proposed under the ACA (which contain few people who already have access to existing public or employer-sponsored insurance). We note that, in addition to this "ACA relevant" MEPS sample, we also perform our equilibrium and welfare analysis for a second, broader, sample composed of all individuals in MEPS between the ages of 25 and 65, including those with employer sponsored or public insurance (Column 2 in Table S5, N = 81,733). For the remainder of this section, we focus on the "ACA relevant" MEPS sample, our primary sample of interest.

Our analysis matches individuals in the employer data used throughout our analysis to the MEPS "ACA relevant" population and creates a new simulation sample with demographic weights similar to the MEPS sample but with detailed health and risk preference data from our estimates.[10] We match individuals in our data to those in the MEPS data based on three demographics: age, income, and gender. To do this, we probabilistically model cells of age, gender, and income in the MEPS sample, and then draw randomly from individuals in those bins in our data with weights proportional to the MEPS cell weights. We note that, before we construct the MEPS cell weights, we incorporate the survey sample weights in the MEPS data, which are intended to correct for sampling and response issues. Table S6 describes the non-parametric age, income, and gender cell multivariate cell weights for

[10] We bring in the cost data from our data set because it is more detailed on the health risk dimension and our setting provides more precise plan characterizations, with which it is possible to estimate risk preferences.

this MEPS sample.[11] [12]

For the uninsured / individual market MEPS reweighted sample, we reproduce our earlier equilibrium and welfare analysis for the cases of (i) pure community rating and (ii) health status- based pricing for health status quartiles in the market setup where insurers can offer either 90% or 60% insurance contracts. Table S8 presents the main results for this sample, and can be directly compared to Table 4 from the main text. The comparison yields several important insights. First, the equilibrium premia and market shares are similar in this MEPS re-weighted sample and our main analysis: the market fully unravels to all-in-60 for the case of pure community rating. Under health-based pricing, in both cases the healthiest quartile has substantial market share in both 60 and 90: in our main analysis 64.8% in this quartile choose 60% coverage, while 57.5% do in the exchange-relevant MEPS re-weighted sample. Interestingly, while no consumers from the second healthiest quartile enroll in 90% coverage in our primary analysis, in the MEPS re-weighted sample 30.4% do. Thus, under our framework, if the exchanges are comprised of only uninsured individuals and those that would have been on the individual market, there will be higher insurance rates for the within-exchange population under health status-based pricing. For both our primary and MEPS analysis, the market unravels for the two sickest quartiles. Finally, and importantly, we note that the population expense levels are very similar between our main sample and the re-weighted MEPS sample: if all enroll in 60, the average costs in the former are \$4,051 while in the latter they are \$3,901.[13] Overall, the analysis of MEPS data in this section suggests that, at a first pass, our main results are not substantially changed when applied to a sample that more closely reflects the demographic profile of individuals who will sign up for the ACA state exchanges. [14] Table S9 and S10 present, respectively, the demographics and equilibrium results for the broader sample of all individuals in MEPS between the ages of 25 and 65, including those with employer sponsored or public insurance within-sample.

[11] We note that in this analysis, we do not match our sample to MEPS using health expenditure data (conditional on the other demographics) since our sample has more detailed medical information on consumers. However, the analysis and tables below show that average costs conditional on demographic bins are similar in our data and in the MEPS data.

[12] We note that Table S5 presents the data "as is." In our analysis, we use MEPS sample weights, which re-weight this "as is" population to correct for survey sampling bias. In addition, as in our main analysis, we assume that the market is purely an individual market: there could be multiple people from one family in each sample represented in Table S5.

[13] Though the means are similar, the exchange-relevant MEPS sample is more heavily skewed in both directions, with more very healthy and more very sick individuals.

[14] The market unravelling we find under community rating (with or without age-based pricing) is somewhat consistent with experience in the Massachusetts exchange, where most buyers opted for the Bronze (60%) plan in the early years of this ACA-like exchange [see, e.g., Ericson and Starc (2013)].

	Entire MEPS (1)	All Ind. 25-65 (2)	25-65 Unins/Ind (3)
N - Individual-Year Obs.	166,539	81,733	21,856
N - Individuals in Panel	105,353	51,922	13,804
N - Family-Year Obs.	58,647	-	-
N - Families in Panel	36,317	-	-
Avg. Family Members	2.90	-	-
Age-Individual			
Mean	33.82	43.15	42.6
10th Qtile	5	28	27
25th Qtile	14	34	32
Median	32	43	42
75th Qtile	51	52	52
90th Qtile	66	59	60
Gender-Individual			
Male %	47.7%	46.6%	50.2%
Total Income-Family-Year* **			
Mean	53613	64058	42746
10th Qtile	9240	12733	8000
25th Qtile	19000	26000	17068
Median	39080	50000	31114
75th Qtile	72375	85584	54995
90th Qtile	115086	131080	89600
Wage Income-Family-Year**			
Mean	44583	59945	38882
10th Qtile	0	7348	300
25th Qtile	8000	24000	14280
Median	32000	48300	30000
75th Qtile	65000	83753	52000
90th Qtile	104438	124996	82680
Region-Individual			
Northeast	14.5%	15.0%	10.1%
Midwest	19.2%	19.6%	15.0%
South	38.3%	38.7%	46.3%
West	26.9%	26.8%	28.7%

Table S5: This table describes demographic data for key samples of interest in the MEPS data, for the pooled data from 2004-2008. A more detailed description of each column's sample is contained in the text.

*In individual samples, a given family's income may count twice since two individuals can be from same family.

MEPS Weights Incorporated					
All 25-65 Sample					
Age Bucket / Fam. Wages	0-$35,000	$35,000-$70,000	$70,000-$105,000	$\geq$ 105,000	**Total**
25-29	4.1%	4.5	2.7	1.9	13.1%
30-34	3.3%	4.4	2.6	1.9	12.3%
35-39	3.5%	4.2	2.8	2.3	12.9%
40-44	3.6%	4.5	3.0	2.8	13.9%
45-49	3.5%	4.2	3.0	3.1	13.9%
50-54	3.5%	3.8	2.8	2.9	13.1%
55-59	3.8%	3.2	2.3	2.3	11.7%
60-64	4.4%	2.3	1.3	1.2	9.2%
Total	29.7%	31.1%	20.5%	18.4%	100%
% Male by Income*	45.6%	49.9%	50.3%	51.4%	

25-65 Unins./ Private					
Age Bucket / Fam. Wages	0-$35,000	$35,000-$70,000	$70,000-$105,000	$\geq$ 105,000	**Total**
25-29	7.4%	5.0	1.9	1.6	15.9%
30-34	6.0%	4.4	1.3	0.7	12.4%
35-39	6.4%	3.5	1.1	0.6	11.6%
40-44	6.1%	4.0	1.4	0.8	12.2%
45-49	6.2%	3.1	1.6	0.9	10.8%
50-54	5.9%	2.9	1.1	0.9	10.8%
55-59	7.0%	2.5	1.1	0.8	11.4%
60-64	10.1%	2.3	0.8	0.8	14.0%
Total	55.1%	27.7%	10.3%	7.1%	100%
% Male by Income*	51.4%	56.2%	55.4%	56.8%	

Table S6: This table describes the discrete age probabiliities for different age / gender / income categories for (i) all individuals in MEPS, age 25-65, and (ii) all uninsured / individual market insured individuals in MEPS, age 25-65. These weights incoporate MEPS sample weights as well, as an additional weighting factor. *Percentages of gender across age are essentially constant conditional on income, which is why those figures are not presented here.

MEPS Weights Incl.							
All 25-65 Sample							
Age Bucket / Quantile	10th	25th	50th	75th	90th	95th	**Mean**
25-29	0 (0)	0 (203)	125 (843)	620 (2833)	2109 (7638)	4155 (12007)	**997 (2820)**
30-34	0 (0)	0 (241)	224 (940)	922 (3179)	2815 (9040)	5582 (13122)	**1376 (3146)**
35-39	0 (0)	0 (239)	331 (925)	1314 (2928)	3499 (8158)	6333 (13595)	**1696 (3126)**
40-44	0 (0)	25 (258)	450 (967)	1669 (2955)	4513 (7844)	9099 (13843)	**2235 (3544)**
45-49	0 (0)	115 (365)	703 (1342)	2425 (3827)	6423 (9143)	12125 (15505)	**3016 (3838)**
50-54	0 (90)	221 (563)	1114 (1860)	3385 (4744)	8562 (10683)	16271 (17135)	**4187 (4551)**
55-59	0 (102)	410 (781)	1837 (2437)	4953 (5820)	11929 (13615)	21069 (22741)	**5315 (6129)**
60-64	71 (255)	707 (1109)	2337 (2906)	5916 (6771)	15261 (14493)	27033 (24997)	**6790 (6666)**

25-65 Unins./ Private							
Age Bucket / Quantile	10th	25th	50th	75th	90th	95th	**Mean**
25-29	0 (0)	0 (0)	0 (166)	173 (758)	819 (2959)	1824 (5502)	**391 (952)**
30-34	0 (0)	0 (0)	0 (180)	254 (852)	1062 (3234)	2024 (6095)	**608 (1322)**
35-39	0 (0)	0 (0)	0 (174)	328 (1024)	1650 (3187)	3164 (5748)	**744 (1223)**
40-44	0 (0)	0 (0)	50 (308)	750 (1459)	2929 (3966)	4500 (6908)	**1381 (2449)**
45-49	0 (0)	0 (0)	120 (425)	857 (1846)	3108 (4566)	6719 (9658)	**2089 (1967)**
50-54	0 (0)	0 (144)	340 (798)	1576 (2866)	5590 (7462)	11851 (12952)	**2474 (3085)**
55-59	0 (0)	24 (176)	1076 (1312)	3565 (3996)	9290 (9990)	16419 (19459)	**3898 (4941)**
60-64	0 (60)	449 (732)	1966 (2398)	5166 (5730)	13749 (12017)	24157 (21839)	**6003 (6043)**

Table S7: This table describes the expenditure quantiles for (i) all individuals in MEPS age 25-65 (top panel) and (iii) all uninsured / individual market insured individuals in MEPS, age 25-65 (bottom panel). Female numbers presented in parantheses, male numbers are not.

MEPS Unins. Weighted: Equilibria without Pre-existing Conditions						
Equilirium Type	P_{60}	S_{60}	AC_{60}	P_{90}	S_{90}	AC_{90}
RE	3,901	100.0	3,901	–	0	–
NE			Does not exist			

MEPS Unins. Weighted: Equilibria with Health Status-based Pricing (Quartiles)							
Market	Equilibrium Type	P_{60}	S_{60}	AC_{60}	P_{90}	S_{90}	AC_{90}
Quartile 1	RE	311	57.5	311	1,476	42.5	1,476
Quartile 2	RE	1,128	69.6	1,128	3,228	30.4	3,228
Quartile 3	RE	4,121	100.0	4,121	-	0	-
Quartile 4	RE	9,751	100.0	9,751	-	0	-

Table S8: This table presents the analogous table to Table 4 (in the main text) on equilibrium outcomes, applied to the sample reweighted by characteristics of the uninsured / individual coverage MEPS, described in the text. The top presents the equilibrium results for the case of pure community rating (no pricing of pre-existing conditions) and the bottom for the case where insurers can price based on health status quartiles.

	Entire MEPS (1)	All Ind. 25-65 (2)	25-65 Unins/Ind (3)
Family-Year: Coverage Type*			
Private (Employer or Ind.)	66.3%	73.3%	41.0%
Medicaid (someone)	30.7%	33.4%	45.4%
Medicare (someone)	29.01%	14.0%	16.4%
Uninsured** (someone)	26.7%	35.0%	84.7%
Only Public in Fam	22.5 %	15.1%	0%
Always Offered Employer (someone)	48.8 %	62.1%	–
Offered Employer Sometimes (someone)	62.0%	76.1%	–
Family Member Emp. Always	69.7%	84.7%	76.2%
Family Member Emp. Once	77.5%	92.3%	87.4%
Individual-Year: Coverage Type*			
Private (Employer or Ind.)	54.5%	64.0%	16.8%
Medicaid	25.4%	12.4%	0.72%
Medicare	13.4%	3.9% 1	.25%
Uninsured**	16.6%	22.3%	83.2%
Only Public	27.6%	12.7%	0%
Always Offered Employer	21.3 %	38.9%	–
Offered Employer Sometimes	32.5%	55.0%	–
Individual Emp. Always	37%	65.4%	37.5%
Individual Emp. Once	48%	78.3%	48.0%

Table S9: This table describes insurance coverage, expenditures, and other statistics in the MEPS data for the pooled data from 2004-2008. A more detailed description of each column's sample is contained in the text.
*Coverage type reflects whether a family ever had this kind of coverage (for any member) throughout the year, so these numbers add to more than 100%.
**Uninsured variable occurs when none of other coverage types are held, and the family is uninsured for whole year.

MEPS Weighted: Equilibria without Pre-existing Conditions							
	Equilirium Type	P_{60}	S_{60}	AC_{60}	P_{90}	S_{90}	AC_{90}
	RE	3,852	100.0	4,051	–	0	–
	NE			Does not exist			

MEPS Weighted: Equilibria with Health Status-based Pricing (Quartiles)							
Market	Equilibrium Type	P_{60}	S_{60}	AC_{60}	P_{90}	S_{90}	AC_{90}
Quartile 1	RE	321	60.2	321	1,521	39.8	1,521
Quartile 2	RE	1,445	100.0	1,445	-	0	-
Quartile 3	RE	4,239	100.0	4,239	-	0	-
Quartile 4	RE	9,347	100.0	9,347	-	0	-

Table S10: This table presents the analogous table to Table 4 (in the main text) on equilibrium outcomes, applied to the sample reweighted by characteristics of the MEPS full population, as described in the text. The top presents the equilibrium results for the case of pure community rating (no pricing of pre-existing conditions) and the bottom for the case where insurers can price based on health status quartiles.