SUPPLEMENT TO "EQUILIBRIA IN HEALTH EXCHANGES: ADVERSE SELECTION VERSUS RECLASSIFICATION RISK"

By Ben Handel, Igal Hendel, and Michael D. Whinston

This Online Appendix has three sections. The first presents details of the choice model estimation algorithm, as well as additional estimates from our primary specification not included in the main text. The second describes our model for consumer self-insurance from savings and borrowing in detail. The third provides additional figures and tables referenced in the main text.

APPENDIX B: Choice Model Estimation Algorithm Details and Additional Results

This appendix describes the details of the choice model estimation algorithm. The corresponding section in the text provided a high-level overview of this algorithm and outlined the estimation assumptions we make regarding choice model fundamentals and their links to observable data. In addition, after the presentation of the estimation algorithm, we discuss further specification details and results for our primary choice model.

We estimate the choice model using a random coefficients simulated maximum likelihood approach similar to that summarized in Train (2009). The simulated maximum likelihood estimation approach has the minimum variance for a consistent and asymptotically normal estimator, while not being too computationally burdensome in our framework. Since we use panel data, the likelihood function at the family level is computed for a *sequence* of choices from t_0 to t_2, since inertia implies that the likelihood of a choice made in the current period depends on the choice made in the previous period. The maximum likelihood estimator selects the parameter values that maximize the similarity between actual choices and choices simulated with the parameters.

First, the estimator simulates Q draws from the distribution of health expenditures output from the cost model, F_{jkt}, for each family, plan, and time period. These draws are used to compute plan expected utility conditional on all other preference parameters. It then simulates S draws for each family from the distributions of the random coefficients γ_j and δ_j, as well as from the distribution of the preference shocks ϵ_k. We define the set of parameters θ as the full set of ex ante model parameters (before the S draws are taken):

$$\theta \equiv \left(\mu, \beta, \sigma_\gamma^2, \mu_\delta(A_j), \sigma_\delta(A_j), \alpha, \mu_{\epsilon_K}(A_j), \sigma_{\epsilon_K}(A_j), \eta_0, \eta_1 \right).$$

We denote θ_{sj} one draw derived from these parameters for each family, including the parameters constant across draws:

$$\theta_{sj} \equiv (\gamma_j, \delta_j, \alpha, \epsilon_{KT}, \eta_0, \eta_1).$$

Denote θ_{Sj} the set of all S simulated draws for family j. For each θ_{sj}, the estimator then uses all Q health draws to compute family-plan-time-specific expected utilities U_{sjkt} following the choice model outlined in earlier in Section 3. Given these expected utilities for each θ_{sj}, we simulate the probability of choosing plan k in each period using a smoothed Accept–Reject function with the form

$$\mathrm{Pr}_{sjt}\big(k = k^*\big) = \left(\frac{\frac{1}{-U_{sjk^*t}}(\cdot)}{\sum\limits_{K} \frac{1}{-U_{sjkt}}(\cdot)} \right)^{\tau} \Bigg/ \sum\limits_{k} \left(\frac{\frac{1}{-U_{s\hat{j}kt}}(\cdot)}{\sum\limits_{K} \frac{1}{-U_{sjkt}}(\cdot)} \right)^{\tau}.$$

This smoothed Accept–Reject methodology follows that outlined in Train (2009) with some slight modifications to account for the expected utility specification. In theory, conditional on θ_{sj}, we would want to pick the k that maximizes U_{jkt} for each family, and then average over S to get final choice probabilities. However, doing this leads to a likelihood function with flat regions, because for small changes in the estimated parameters θ, the discrete choice made does not change. The smoothing function above mimics this process for CARA utility functions: as the smoothing parameter τ becomes large, the smoothed Accept–Reject simulator becomes almost identical to the true Accept–Reject simulator just described, where the actual utility-maximizing option is chosen with probability 1. By choosing τ to be large, an individual will always choose k^* when $\frac{1}{-U_{jk^*t}} > \frac{1}{-U_{jkt}}$ $\forall k \neq k^*$. The smoothing function is modified from the logit smoothing function in Train (2009) for two reasons: (i) CARA utilities are negative, so the choice should correspond to the utility with the lowest absolute value, and (ii) the logit form requires exponentiating the expected utility, which in our case is already the sum of exponential functions (from CARA). This double exponentiating leads to computational issues that our specification overcomes, without any true content change since both models approach the true Accept–Reject function.

Denote any sequence of three choices made as k^3 and the set of such sequences as K^3. In the limit as τ grows large, the probability of a given k^3 will approach either 1 or 0 for a given simulated draw s and family j. This is because, for a given draw, the sequence (k_1, k_2, k_3) will either be the sequential utility-maximizing sequence or not. This implicitly includes the appropriate level of inertia by conditioning on previous choices within the sequential utility calculation. For example, under θ_{sj} a choice in period 2 will be made by a family j only if it is optimal conditional on θ_{sj}, other preference factors, and the inertia implied by the period 1 choice. For all S simulation draws, we compute the optimal sequence of choices for k with the smoothed Accept–Reject

simulator, denoted k_{sj}^3. For any set of parameter values θ_{Sj}, the probability that the model predicts k^3 will be chosen by j is

$$\widehat{P_j^{k^3}}\big(\theta, F_{jkt}, \mathbf{Z}_j^A, \mathbf{Z}_j^B, H_j, A_j\big) = \sum_{s\in S} \mathbf{1}\big[k^3 = k_{sj}^3\big].$$

Let $\widehat{P_j^{k^3}}(\theta)$ be shorthand notation for $\widehat{P_j^{k^3}}(\theta, F_{jkt}, Z_j^A, Z_j^B, H_k, A_j)$. Conditional on these probabilities for each j, the simulated log-likelihood value for parameters θ is

$$\mathrm{SLL}(\theta) = \sum_{j\in J} \sum_{k^3\in K^3} d_{jk^3} \ln \widehat{P_j^{k^3}}.$$

Here d_{jk^3} is an indicator function equal to 1 if the actual sequence of decisions made by family j was k^3. Then the maximum simulated likelihood estimator (MSLE) is the value of θ in the parameter space Θ that maximizes $\mathrm{SLL}(\theta)$. In the results presented in the text, we choose $Q = 100$, $S = 50$, and $\tau = 6$, all values large enough such that the estimated parameters vary little in response to changes.

B.1. *Specification for Inertia*

In the main text, we did not describe the details for our specification for consumer inertia. The model for inertia, which is similar to that in Handel (2013), specifics an inertial cost $\eta(\mathbf{Z}_j^B)$ that is linearly related to consumer characteristics and linked choices, $\mathbf{Z}_j^B$:

$$\eta\big(\mathbf{Z}_j^B\big) = \eta_0 + \eta_1 Z_{jt}^B.$$

The characteristics in $\mathbf{Z}_j^B$ include family status (e.g., single or covering dependents), income, several job status measures, linked choice of Flexible Spending Account (FSA), and whether the family has any members with chronic medical conditions (and, if so, how many chronic conditions total in the family).

B.2. *Additional Results*

In the interest of space, the text only presented the risk preference parameter estimates from our primary specification, since this was the key object of interest recovered there for our equilibrium analysis of insurance exchange pricing regulations. Here, for completeness, in Tables B.I and B.II we include the full set of estimates in the primary model for reference, including inertia parameters, PPO_{1200} random coefficients, and ε standard deviations. Overall, the parameters not discussed in the text have similar estimates to those in

TABLE B.I

THIS TABLE PRESENTS THE FIRST HALF OF THE FULL SET OF PRIMARY
CHOICE MODEL ESTIMATES[a]

Empirical Model Results

Parameter/Model	(1) Primary Model	(2) Parameter Standard Error
Risk Preference Estimates		
μ_γ—Intercept, β_0	$1.21 * 10^{-3}$	$5.0 * 10^{-5}$
μ_γ—$\log(\sum_{i \in j} \lambda_i)$, β_1	$-1.14 * 10^{-4}$	$9.8 * 10^{-6}$
μ_γ—age, β_2	$-5.21 * 10^{-6}$	$1.0 * 10^{-7}$
μ_γ—$\log(\sum_{i \in j} \lambda_i) * $age, β_3	$1.10 * 10^{-6}$	$1.3 * 10^{-7}$
μ_γ—Manager, β_4	$4.3 * 10^{-5}$	$5.2 * 10^{-5}$
μ_γ—Manager Ability, β_5	$1.4 * 10^{-5}$	$1.2 * 10^{-5}$
μ_γ—Nonmanager Ability, β_6	$7.5 * 10^{-6}$	$2.4 * 10^{-6}$
μ_γ—Population Mean	$4.39 * 10^{-4}$	–
μ_γ—Population σ	$6.63 * 10^{-5}$	–
σ_γ—γ Standard Deviation	$1.24 * 10^{-4}$	$3.5 * 10^{-5}$
Inertia Estimates		
η_0, Intercept	1,336	76
η_1, Family	2,101	52
η_1, FSA Enroll	−472	44
η_1, Income	96	15
η_1, Quantitative	6	27
η_1, Manager	162	34
η_1, Chronic Condition	108	24

[a]The set of estimates relevant for our analysis of exchange pricing regulation is presented and interpreted in much more detail in the main text. Standard errors are presented in Column 2.

Handel (2013), though the risk preference estimates differ here because they are linked explicitly to health risk to estimate correlations between those two micro-foundations.

APPENDIX C: SELF-INSURANCE MODEL

Section 6.3 describes our extension that allows for consumers to save and borrow to self-insure against health shocks. That section in the main text describes the key features of our model of saving and borrowing as well as the results from that model. In this section, we provide some additional details on this model and present a more formal treatment of it.

We allow for borrowing and saving by solving a finite horizon dynamic problem. To clarify notation and timing, we define the following terms:

- $W_t \equiv$ income in period t.
- $p_{it} \equiv$ price of policy i in period t.
- $m_t \equiv$ medical expenses in period t.

TABLE B.II

THIS TABLE PRESENTS THE SECOND HALF OF THE FULL SET OF PRIMARY
CHOICE MODEL ESTIMATES[a]

Empirical Model Results

Parameter/Model	(1) Primary Model	(2) Parameter Standard Error
PPO_{1200} Preferences		
μ_δ: Single	−2,504	138
σ_δ: Single	806	47
μ_δ: Family	−2,821	424
σ_δ: Family	872	48
Other		
α, High-Cost, PPO_{250}	−805	79
ε_{500}, σ_ε, Single	50	340
ε_{1200}, σ_ε, Single	525	180
ε_{500}, σ_ε, Family	141	56
ε_{1200}, σ_ε, Family	615	216

[a]The set of estimates relevant for our analysis of exchange pricing regulation is presented and interpreted in much more detail in the main text. Standard errors are presented in Column 2.

- $\lambda_t \equiv$ ACG health status realization for period t (realized in period $t-1$).
- $O_i(m) \equiv$ out-of-pocket expense for policy i with medical expenses outcome m.
- $S_t \equiv$ savings chosen in period t.
- $\overline{W}_t \equiv W_t + (1+r)S_{t-1}$ are funds available in period t.
- $c(m|i_t) = p_{i_t} + O_{i_t}(m)$ is the consumer's total medical expenses under policy i_t given m.

Timing:

In each period t, the consumer chooses an insurance policy, (λ_{t+1}, m_t) is realized, and then a savings decision, S_t, is made. Given λ_{t+1}, m_{t+1} is then drawn in period $t+1$ from a distribution $F_{t+1}(m_{t+1}|\lambda_{t+1})$. Thus, period t savings are decided after observing health expenses for period t and period $t+1$'s health status. This assumption reflects a fluid financial market where individuals can take a last minute loan if they were unlucky or deposit extra cash if they were healthier than expected.

Solving the model:

We start in period T and solve for optimal savings backward. In period T given realization λ_T and starting savings plus income $\overline{W}_T$, consumer expected utility is

$$-E\big[e^{-\gamma[\overline{W}_T - c(m_T|i_T)]}|\lambda_T\big] = -E\big[e^{-\gamma[W_T - c(m_T|i_T)]}|\lambda_T\big]e^{-\gamma(1+r)S_{T-1}}.$$

Given that $i_T^*(\lambda_T)$ is the consumer's policy choice at T when he has health status λ_T, expected period T utility is

$$-E\big[e^{-\gamma[W_T - c(m_T|i_T^*(\lambda_T))]}|\lambda_T\big]e^{-\gamma(1+r)S_{T-1}},$$

which is a function of λ_T and S_{T-1}. We can thus denote the value function in period T as a function of the state, $V_T(\lambda_T, S_{T-1})$. Optimal period $T-1$ saving S_{T-1} (saving for period T) solves

$$\max_{S_{T-1}, i_{T-1}} -E\big[e^{-\gamma[\overline{W}_{T-1} - c(m_{T-1}|i_{T-1})]}|\lambda_{T-1}\big]e^{\gamma S_{T-1}} + \delta V_T(\lambda_T, S_{T-1}),$$

which in turn delivers $V_{T-1}(\lambda_{T-1}, S_{T-2})$.

In this manner, we recursively solve the optimal savings level all the way backwards to period 1 for every possible history. Once we have $V_1(\lambda_1, 0)$, we compute the ex ante welfare of an unborn individual who does not yet know her future λ_1 as

$$W_0(W) = E_{\lambda_1}\big(V_1(\lambda_1, 0)\big).$$

The ex ante welfare depends on the income profile $W = [W_1, W_2, \ldots, W_T]$, on the initial distribution of types, and on the regulatory pricing regimes we want to evaluate. A pricing regime affects expected welfare through both the out-of-pocket expenses $O_i(m_t)$ as well as the premium paid, $p_i(\lambda_t)$. We translate the ex ante welfare difference between pricing regimes into yearly certainty equivalent values as in Section 5 in the main text.

C.1. *Computation*

To implement the dynamic problem, we need assumptions about the evolution of the state variable. Unlike the primary welfare analysis in the paper (which assumed a steady state population), the computation here requires transitions across health states (predictive ACG index) over time. Namely, at any point in time, we need to compute the expected evolution of the future uncertainty, to figure out optimal savings.

We estimate health state transitions using the observed transitions in our sample. So that we have enough sample size to nonparametrically estimate this transition matrix, we divide the population into 7 groups based on health status and compute a 7-by-7 transition matrix for each of 8 five-year age bins (25–30, 30–35, …). We assume that the estimated transition matrix for each five-year age bin reflects the transition probabilities for consumers in that five-year age bin transitioning to a given health status level for the next five-year age bin. Within each period, consumers experience five years of identical health claims in the insurance contract they chose for that period, appropriately discounted. For each age bin, health type, and regulatory pricing regime, we use the static

market equilibrium outcomes from our primary analysis $i^*(\lambda)$ and determine the actual choice each individual makes in each period, yielding her premia and out-of-pocket expenses.[54] We assume consumers have flat income profiles over time ($W_t = W$) (as in the first column of Table VI) in order to neutralize the other channels through which savings could impact welfare.

Given this setup, we solve the 8-period dynamic problem as described above. Once we recover the value function for an unborn individual (prior to age 25) for each possible realization of the initial health state, we compute the certainty equivalent of each regime x as

$$-\sum_t^{T=8} \beta^t e^{-\gamma CE_x} = -\sum_j^7 p_j e^{-\gamma V_1(\lambda_j, 0)}.$$

The results from this model are presented and discussed in Section 6.3.

APPENDIX D: ADDITIONAL ANALYSIS

This appendix contains several additional figures and tables discussed in the text. Figure D.1 presents the distribution of λ predicted for t_1, for all individuals in the data (including dependents) present during both t_0 and t_1. Predicted expected expenses are normalized by the average in this sample of \$4,878 (thus equal to 1 in this chart). The distribution presented is truncated at 5 times for this chart, but not in estimation/analysis. See Handel (2013) for additional detailed analysis of expected expenditures for employees at dependents at the firm we study.

Table D.I presents descriptive statistics for the pseudo-sample of individuals used in our insurance exchange simulations. The sample has a risk preference mean and standard deviation similar to those of the choice model estimation sample. Moreover, the distributions of income and health status are similar to those in the estimation sample and in the general population. The table just below in the text here illustrates that the simulation sample (as in our data overall) has a fairly uniform distribution of age between 25 and 65, consistent with our assumption of a steady state population in the welfare analysis. See

[54]Market outcomes are assumed to be the same as those in our primary equilibrium analysis. They thus do not account for a potential effect that borrowing and saving would have on consumer insurance choices. Accounting for these dynamic effects would likely push consumers more toward lower insurance, and thus likely not have a large impact on equilibrium outcomes. This reflects the goal of this section, which is to quantify the impact of savings on the welfare numbers, keeping other things (including static market equilibrium outcomes) equal. In that spirit, we keep the equilibrium prediction unchanged, as described in the paper for each pricing regime, and see how a representative individual's welfare would change if she is allowed to borrow or save.

FIGURE D.1.—This figure presents the distribution of λ predicted for t_1, for all individuals in the data (including dependents) present during both t_0 and t_1. Predicted expected expenses are normalized by the average in this sample of \$4,878 (thus equal to 1 in this chart). The distribution presented is truncated at 5 times for this chart, but not in estimation/analysis.

Section 3.6 for further details on the sample used in our counterfactual analyses.

Quantile	0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.95
Age	26	28	33	37	41	45	49	52	56	60	62

Table D.II shows average costs as a function of age 25 risk preferences, to illustrate the relationship between risk preferences and age that exists in our welfare framework. Following the choice model estimates, costs are negatively related to risk aversion conditional on age. See Section 5 in the main text for more details.

Table D.III supports the analysis in our age-based pricing extension in Section 6 in the main text. The table shows the compensation required to make an individual indifferent between a regime with health status quartile pricing for each age group, and another in which all individuals in each age band receive the 60 policy at its average cost for their age band (the result of pure age-based pricing). Once age is priced, health-based pricing, which appealed to individuals with steeply increasing income, is no longer preferred by those consumers. The benefit of health-based pricing is the reduction in adverse selection, and the postponement of premiums until later in life. With age-based pricing, the latter benefit is eliminated. The cost associated with reclassification risk then dominates the benefits of reducing adverse selection across the range of risk aversion types and for the different income path models studied.

Table D.IV presents the long-run welfare implications of allowing for insurer risk adjustment transfers, as specified in the HHS risk adjustment formula, described in Section 6 in the main text. Risk adjustment transfers partially reduce

TABLE D.I

THIS TABLE PRESENTS DESCRIPTIVE STATISTICS FOR
THE PSEUDO-SAMPLE OF INDIVIDUALS USED IN
OUR INSURANCE EXCHANGE SIMULATIONS[a]

Simulation Sample	
	Simulation Sample
N—Families	–
N—Individuals 25–65	10,372
Mean Age	44.5
Median Age	45
Gender (Male %)	45
Income	
Tier 1 (<$41K)	20%
Tier 2 ($41K–$72K)	40%
Tier 3 ($72K–$124K)	24%
Tier 4 ($124K–$176K)	8%
Tier 5 (>$176K)	8%
Predicted Mean Total Expenditures	
Mean	$6,099
25th Quantile	$1,668
Median	$3,654
75th Quantile	$8,299
90th Quantile	$13,911
95th Quantile	$18,630
99th Quantile	$34,008
Risk Preferences	
Mean μ_γ	$4.28 * 10^{-4}$
Standard Deviation μ_γ	$7.50 * 10^{-5}$

[a]The sample has risk preference means and standard deviations that are similar to those of the choice model estimation sample. Moreover, the distributions of income and health status are similar to those in the estimation sample and general population.

TABLE D.II

AVERAGE COSTS AS A FUNCTION OF AGE 25
RISK PREFERENCES[a]

	Average Costs at Various Ages Conditional on Age 25 Risk Aversion		
γ	30–35	45–50	55–60
0.0002	5,586	7,196	10,857
0.0003	4,212	6,390	10,319
0.0004	3,100	5,687	9,767
0.0005	2,328	4,911	9,271
0.0006	1,775	4,373	8,813

[a]Following the choice model estimates, costs are negatively related to risk aversion conditional on age.

TABLE D.III

LONG-RUN WELFARE COMPARISON BETWEEN THE TWO PRICING REGULATIONS OF
(I) PRICING BASED ON HEALTH STATUS QUARTILES BY AGE (x = "HB4 + age")
AND (II) PRICING BASED ON JUST AGE (x' = "age")[a]

	Welfare Loss From Health Status-Quartile Age-Based Pricing ($/Year)		
γ	$y_{\text{HB4+age,age}}(\gamma)$ Fixed Income	$y_{\text{HB4+age,age}}(\gamma)$ Non-Manager Income Path	$y_{\text{HB4+age,age}}(\gamma)$ Manager Income Path
0.0002	2,111	2,129	1,100
0.0003	2,911	2,028	920
0.0004	3,707	1,842	778
0.0005	4,510	1,646	1,353
0.0006	5,137	1,612	1,876

[a]The results presented are based on the RE outcomes for each of the two pricing regulations. As before, the assumed discount rate is $\delta = 0.975$.

the extent of adverse selection under pure community rating, improving consumer welfare.

Figure D.2 presents an additional calibration of the framework developed in Section 2 that highlights the tradeoff between adverse selection and reclassification risk, as a function of the fraction of health risk information known by consumers at the time of contracting. This is similar to a figure in that section, but calibrated so that consumers face more health risk ($R = 30,000$). Unraveling occurs at higher ϕ when R is greater (larger variance of medical expenditures), reflecting the fact that with greater variance consumers are more reluctant to choose a low coverage plan. As a result, in the figure in the appendix there is a smaller range of ϕ over which health-based pricing is better than community rating.

TABLE D.IV

LONG-RUN WELFARE IMPLICATIONS OF INSURER RISK ADJUSTMENT REGULATION
(TRANSFERS BASED ON THE HHS RISK ADJUSTMENT FORMULA)

|	Welfare Benefit of Risk Adjustment Transfers: RE ($/Year)		
γ	$y_{\text{PCR,risk-adj}}(\gamma)$ Fixed Income	$y_{\text{PCR,risk-adj}}(\gamma)$ Non-Manager Income Path	$y_{\text{PCR,risk-adj}}(\gamma)$ Manager Income Path
0.0001	316	261	106
0.0002	327	202	27
0.0003	336	139	18
0.0004	349	84	0
0.0005	368	36	38
0.0006	386	23	72

FIGURE D.2.—Adverse selection versus reclassification risk, $R = 30,000$. X curve: market share of low coverage plan; dashed curve: certainty equivalent with pure community rating; solid curve: certainty equivalent with perfect health-based pricing.

Dept. of Economics, University of California at Berkeley, Berkeley, CA 94720, U.S.A.; handel@berkeley.edu,

Dept. of Economics, Northwestern University, Evanston, IL 60208, U.S.A.; igal@northwestern.edu,

and

Dept. of Economics and Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.; whinston@mit.edu.