Economic Research Initiative on the Uninsured
CONFERENCE DRAFT

Health Savings Accounts, High-Deductible Policies, and the Uninsured: Simulating the Effects of HSA Tax Policy Using a Utility-Maximization Framework

Draft: Please do not cite or quote without permission.

James H. Cardon
cardon@byu.edu

Mark H. Showalter
showalter@byu.edu

Department of Economics
Brigham Young University
Provo, UT 84602

June, 2007

Economic Research Initiative on the Uninsured
University of Michigan
555 South Forest Street, 3rd Floor
Ann Arbor, MI 48104-2531

Abstract

The Bush administration has proposed aiding the uninsured population by offering tax credits and deductions to consumers outside the employment-based group market who use Health Savings Accounts (HSAs). It is unclear whether such proposals would actually decrease the number of uninsured. Some analysts have argued that these proposals could adversely affect the employment-based market, causing firms to offer less attractive policies to employees or to drop coverage altogether. But analysis is hindered by lack of data and theoretical uncertainty about consumers' valuation of HSAs and related products.

Our research will build on previous work to explore the impact of these proposals on the uninsured. A novel aspect of our approach is a calibration of underlying preference parameters that allow us to simulate how consumers will value products like high-deductible insurance and HSAs. We focus on three groups of particular policy interest: 1) the currently uninsured who do not have access to group coverage, 2) the currently uninsured who have access to group coverage, but choose to be uninsured, 3) the currently insured in group coverage.

1 Introduction

In December of 2003, Congress passed the Modernization of Medicare Act which had as a primary goal the addition of a prescription drug benefit to the Medicare program. But the legislation also included a provision establishing health savings accounts (HSAs). Such accounts, when coupled with a catastrophic high deductible health insurance policy (CHP), allow individuals to avoid income and payroll taxes for qualified medical expenditures.

The Bush administration has offered a number of proposals to make HSAs and high-deductible health insurance policies more affordable and attractive for the currently uninsured. For example, current proposals include the following: 1) create an income tax deduction for the premium on a HSA-qualified insurance policy, 2) create an income tax credit for HSA-qualified policies purchased outside the employment-based group market, and 3) increase the allowable annual HSA contribution amount. (State of the Union Address, January 31, 2006). The first two proposals would lower the cost of insurance by extending the tax subsidy to those not purchasing insurance through employers; and the third proposal was recently signed into law.

While these proposals would likely lower the cost for HSA-qualified insurance, it is unclear whether such proposals would actually decrease the number of uninsured. The tax-based policies noted above are designed to eliminate the tax advantage currently available to employment-based insurance. But employment-based insurance pools also economize on underwriting costs and help to reduce the potential for adverse selection. Some analysts have argued that eliminating that differential could adversely affect the employment-based market, causing firms to offer less attractive policies to employees, resulting in lower take-up rates, or to drop coverage altogether (Moon et al., 1996; Zabinski et al., 1999; Glied and Remler, 2005; Hoffman and Tolbert, 2006). The net effect could be a weakening of these pools and an increase in the number of uninsured.

We will explore the impact of HSAs and the various tax proposals on the uninsured population, including the effect of how adopting the tax deduction and credit policies for non-group insurance would affect the employment-based group market. The foundation for our work is a recently developed simulation model based on utility-maximizing, representative agents (Cardon and Showalter, 2007). From our previous work, it is apparent that the form of the HSA contract, which depends both on Federal regulations and employer choices, is important in determining whether the introduction of HSAs undermines traditional insurance pools and differentially affects individuals based on their health status. These issues will have important implications for determining the ultimate impact of these proposals on the uninsured population.

The main advantage of our framework over previous simulation models is the ability to determine consumers' valuation of products for which little or no data exists-e.g. HSA contributions and high-deductible insurance. This is accomplished through empirical calibration of "deep" structural parameters of consumer preferences and the risks associated with uncertain medical expenditures.

To illustrate, consider the problem of assessing the impact of various HSA-linked proposals. It is critical to have an estimate of consumers' price elasticity of high-deductible policies. But there are no generally accepted estimates of this elasticity due to the paucity of data on such policies. Thus previous work has been forced to use elasticity estimates based on more traditional insurance; but it is not clear these elasticities will be similar. Our methodology can adjust for such differences because it is based on consumers' underlying risk preferences.

We focus on three consumer groups of particular policy interest: 1) the currently uninsured who do not have access to group coverage, 2) the currently uninsured who have access to group coverage, but choose to be uninsured, 3) the currently insured in group coverage. For each group, we the Medical Expenditure Panel Survey to assess representative values for demographic variables such as income, age-specific health expenditures, and tax rates. For each group, we will then calibrate our model to match the observable characteristics. Then we model a variety of policy changes and assess the value of each change to consumers and their likelihood of changing from the status quo.

2 Related Simulation Approaches

Several methods have been used to estimate the impact of various health policy proposals. Glied et al. (2002) provide a useful summary of the literature, which we follow in this section.

The Elasticity Approach

This is the most widely used method. It typically uses individual-level data, and in highly simplified terms it can be described as having three steps: 1) estimate the impact of a policy change on prices and income for individuals; 2) compute the change in demand, given the Step (1) change, using elasticity estimates for insurance demand (or other behavioral responses); 3) aggregate across individuals, with the appropriate weighting, to get national estimates.

The Discrete Choice Approach

This method uses a binary choice regression framework where the dependent variable is 1 if a person has insurance, and 0 otherwise. Observable characteristics of the individual (age, income, gender, employment status, etc) and a measure of insurance price are used as explanatory variables. This methodology is similar to the elasticity approach, except that the elasticity estimate is embodied in the regression parameters and the functional form of the specification. The procedure is roughly as follows: 1) estimate the initial regression model which will specify the elasticity; 2) estimate the impact of the policy change on prices and incomes for individuals; 3) estimate the change in the probability of being insured using the data from Step (2) in the regression model of Step (1); 4) aggregate across individuals to get national estimates.[1]

The Matrix Approach

This method uses grouping of individuals, applies group-specific take-up elasticities, and then aggregates to get national estimates. The estimation steps include: 1) estimate the impact of the policy change on average prices and incomes for each group; 2) compute the change in demand for each group, using group-specific take-up elasticities; 3) aggregate across groups to get national estimates.

[1] An interesting recent example of this approach is Feldman et al. (2005) which estimates a conditional logit model of HSA takeup rates for employees in three large firms. They use the estimated model to simulate take-up rates in the group and non-group markets. Health Reimbursement Arrangements (HRA) are used as a proxy for HSA choice since HSAs were not an option for the employees in their data. Their simulations predict a 9% HSA take-up rate in the non-group market, but a very low take-up rate in the group market. A refundable tax-credit for HSAs is predicted to double the HSA take-up rate while lowering the number of uninsured by 8%.

The Reservation Price Approach

This category uses indirect approaches to measure how utility changes with a given policy option. The measured change in utility is then used to estimate behavioral responses. The two primary methodologies are Zabinski et al. (1999) and Pauly and Herring (2002). Zabinski et al. (1999) use a linear approximation to the change in utility, combined with income and price changes from a given policy, to estimate whether a given individual will switch from a 'standard' insurance policy to high-deductible insurance policy coupled with a medical savings account. Pauly and Herring use a revealed preference argument to estimate how individuals would respond to a tax credit for health insurance.

Both of these indirect approaches are tailored to answer specific questions and are not easily comparable to the first three general approaches.

Discussion

Remler et al. (2002) works through an example that shows that the first three methods give roughly the same results if the underlying data and elasticity estimates are the same. One important feature of the first three methods is that there is little scope for variation in insurance policies; to a first approximation, an individual is treated as having insurance or not; there is generally no attempt to account for variation in policy generosity and how consumers might value policy generosity.

But policy generosity is now at the center of the public debate. "Consumer directed health care" focuses on the advantages of high-deductible policies over more traditional insurance. But the framework and data used for currently available simulations are ill-suited for analyzing the impact of such policies. The estimation shortcomings are even more acute when trying to account for tax-subsidized financial accounts like health savings accounts, health reimbursement accounts, and flexible spending accounts. These accounts are explicitly designed to work in conjunction with insurance policies, but no data exist to allow researchers to understand how consumers view the tradeoff between money in a tax-preferred account and the generosity of insurance.

Of the reservation price approaches, it is not clear how the Pauly/Herring approach could handle policy innovations like health savings accounts. The Zabinski et al. (1999) article is an explicit attempt to measure the impact of HSAs, but the linear approximation approach embodies some very strong assumptions about the value of HSAs accounts which lead to some unrealistic predictions.[2]

In the next section we outline a framework for evaluating the impact of recent policy innovations such as health savings accounts and high-deductible policies. Our framework is explicitly based on a utility maximization assumption and this allows us to explore how consumers would react to new policy choices for which no data exists. The calibration of the model is done using observable data so the utility parameters have an empirical foundation, but they allow much greater flexibility in simulating policy options than is possible using current methods.

Our approach can be thought of as a combination of existing methodologies: We group consumers into identifiable demographic categories based on observable characteristics (e.g. income, age, employment status, etc) like the Matrix Approach. Each group is then modeled with a representative agent who maximizes utility, similar to the Reservation Price Approach, except we use a specific form of utility rather than an approximation to a change in utility. Within each demographic category, we then use a discrete choice framework to model the selection between being insured and uninsured, accounting for utility gain or loss given a particular policy proposal. However, insurance status is not completely determined by measured utility, and we allow for unobserved heterogeneity using techniques used in the industrial organization literature (Nevo, 2000).

3 Approach and Methodology

We begin by describing our utility-maximizing framework which will be applied to each demographic group. Our model accounts for the uncertainty inherent in making choices about insurance, and is dynamic: choices made in one period affect outcomes and utility in later periods.

[2]Their model predicts that even without the tax advantage, most consumer would choose to have a high-deductible policy and a health savings account. But we do not observe that outcome.

Preferences

Utility in each period is derived from the consumption of a composite consumption good, C, and Health. Health is determined by a random health state θ and "health services," X. θ denotes the random health status which determines the relative value of health expenditures. The utility function is a modest generalization of a standard constant elasticity of substitution utility function. For a single period it is

$$U(C,H) = \frac{C^\lambda - 1}{\lambda} + \gamma \frac{(X-\theta)^\lambda - 1}{\lambda},$$

where health $H = X - \theta$. The structure on utility implies that optimal X will not be less than θ. Note that the higher the value of θ, the higher will be marginal utility of health services in that period. Were this a single-good model, $1-\lambda$ would also be the coefficient of relative risk aversion. With two goods, however, the interpretation of $1-\lambda$ is not precisely the same; but we will treat it as an approximate measure of risk aversion when interpreting parameters. The parameter γ accounts for the relative value of health compared to C.

Dynamic Optimization Problem

For tractability we assume there are two periods denoted by subscript t: in period 1, a consumer knows her health status for that period, but health status, θ_2, in period 2 is unknown, although it follows a known distribution. An insurance policy for period 2 consists of a coinsurance rate, α_2, and a deductible, D_2. In period 1, the consumer chooses the following:

1. How much of good X to consume in period 1 (X_1).
2. How much to withdraw from the HSA balance available in period 1 to pay for period 1 health services (W_1).
3. How much of period 2 income to allocate to the HSA in period 2 (Z_2).

The HSA balance in period 2, M_2, is determined by current contributions, Z_2, and the return on unused balance from the previous period according to the equation $M_2 = Z_2 + (M_1 - W_1)(1+r)$, where r is the interest rate on HSA balances. The insurance policy will be associated with a premium, P_2, which includes a loading factor. Given these choices, the consumer in period 2 then chooses optimal health expenditures for period 2, X_2, conditional on the realized health state, θ_2. C_2 is determined by the budget constraint.

In period 1, the individual optimization problem is given by

$$\max_{\{W_1, X_1, Z_2\}} EU = \frac{C_1^\lambda - 1}{\lambda} + \gamma \frac{(X_1 - \theta_1)^\lambda - 1}{\lambda} + \beta E \left[\frac{C_2^\lambda - 1}{\lambda} + \gamma \frac{(X_2 - \theta_2)^\lambda - 1}{\lambda} \right] \quad (1)$$

subject to

Evolution of HSA balance:	$M_2 = Z_2 + (M_1 - W_1)(1 + r)$
Consumption Constraint:	$C_t = (Y_t - P_t - Z_t - T_t) - O(X_t) + W_t$
Out of Pocket Costs–I:	$O(X_t) = X_t$, for $X_t < D_t$
Out of Pocket Costs–II:	$O(X_t) = D_t + \alpha(X_t - D_t)$, for $X_t \geq D_t$
HSA Withdrawal Constraint:	$0 \leq W_t \leq M_t$

The values θ_1, M_1, Z_1, P_1, D_1, α_1 are known and given as of time 1. β is the one period discount rate, and T_t is the total tax, which is a function of the marginal tax rate, τ. Note that maximizing (1) assumes optimal choices for health expenditures as well as optimal usage of the HSA account for HSA policies in both periods. Note also that traditional insurance (and no insurance) are special cases of this model and that optimal behavior for these cases can be computed using the same method. The optimized value of (1) is EU^*.

The use of tax-preferred accounts like HSAs is intuitively a multi-period problem. Indeed, some observers have noted that HSAs could act as a useful retirement savings vehicle (Furman, 2006). We approximate a long-horizon dynamic model by allowing conversion of unused HSA balances to consumption after paying taxes plus a penalty, similar to the treatment of early IRA withdrawals. The justification comes from the theoretical framework found in Cardon and Showalter (2007): in an infinite horizon model, the expected marginal utility per dollar (accounting for taxes) will generally be equalized across X_t, C_t, and the discounted value of M_{t+1}, because the marginal dollar could be spent on any of those commodities. We therefore set the model to allow the representative agent to consume as C_2 any unused HSA balance in the second period, at a rate of \$1 of HSA convertible to \$ $(1-\tau)(1-Penalty)$ of C_2. We simulate the model for values for *Penalty* ranging from zero (which allows conversion to consumption after taxes are paid) to 10%, which matches the rules for IRA early withdrawal.[3].

Health State Distribution

Health type is determined by the distribution of θ. There are obviously a variety of ways that the distribution of θ could be modeled. Ideally, we would like to generate a distribution of expenditures that is similar to observed patterns of health expenditures, which is positively-skewed with a mass at zero.

Health expenditures are nearly linear in θ for the specified utility function, so by choosing a distribution for θ which is right-skewed and has a mass at zero, we obtain a similar distribution for health expenditures. We therefore model θ as coming from a mixture distribution of the following form:

$$\begin{aligned} \theta &= 0 \text{ with probability } p \\ &= \Theta \text{ with probability } 1-p \end{aligned}$$

where Θ is lognormal random variable with mean μ and variance σ^2.

Simulation Model

Due to the complexity of the problem, closed-form analytic solutions are not possible and we therefore proceed by developing a computer simulation model. One complication that arises in the simulation of equation (1) is the specification of the insurance contract. We assume the insurance contract arises out of a competitive process with proportional loading, so that the premium for an insurance contract with deductible D and coinsurance rate α is proportional to an actuarially fair contract. But this premium will depend upon the expected expenditures in period 2, which are

[3] The resulting consumption expenditures, C_2, will be too high relative to what we would find with an infinite horizon model (because the agent allocates to C_2 what would have been allocated to M_3), but expected utility, health expenditures, and HSA balances should be a better approximation than not allowing any conversion.

endogenous. We overcome this difficulty by using the following algorithm to compute a rational expectations equilibrium: 1) Choose a 2nd period insurance contract (α_2, D_2) and premium, P_2. 2) Maximize expected utility subject to that particular insurance contract. 3) If expected insured outlays exceed the premium, re-estimate with a higher premium; if the premium exceeds the expected outlays, re-estimate with a lower premium. Steps (1)-(3) are repeated until convergence of the premium. For group insurance, the pooled premium depends on the expected expenditures of consumers in the group, and the algorithm is modified accordingly to reflect group composition.

4 Setup and Calibration

Since the simulation model is programmed to determine optimal behavior for individual households, we simulate behavior for households of different types broken down by such factors as age, family size, sex, income, etc. For each "cell" of similar households we use data on health care utilization from the Medical Expenditure Panel Survey (MEPS) to obtain estimates of the unknown parameters of the utility function (λ and γ) and the distribution of health states (μ and σ^2)in the following way. We observe the sample mean and variance of expenditures for each cell, and we have estimates of the population price and income elasticities of health expenditures from the RAND Health Insurance Experiment. We obtain method of moments estimators for the 4 parameters by matching these sample moments with theoretical moments (which are functions of the parameters to be estimated) implied by the period utility function $U(C, H)$ above.[4] Using this method we can match the mean and variance of expenditures (measured in \$1,000s) as well as price and income sensitivity in a simple, flexible way. To estimate $p = Pr(\theta = 0)$ we simply use the frequency of zero expenditures for each cell. We will use estimates available from the literature on time preference for β.[5]

For this preliminary stage, we use 20 cells: 10 income deciles and an age indicator for whether the head of household is Young (age$\leq$ 40) or Old (age$>$ 40). Table 1 shows theoretical and sample

[4]The demand function for health expenditures is $X^* = \frac{Y+K\theta}{K+\alpha}$ This implies that the $E(X^*) = \frac{Y+KE(\theta)}{K+\alpha}$ and $Var(X^*) = (\frac{K}{K+\alpha})^2 Var(\theta)$. Price and income elasticities are the standard definitions evaluated at the mean of θ. See Table 1. Elasticities can vary across cells, but in preliminary work they were set equal for all cells at $\varepsilon_p = -0.15$ and $\varepsilon_Y = 0.2$. It is also feasible to estimate distinct $(\mu, \sigma^2, \lambda, \gamma)$ for each cell. We will experiment with which method provides the most reasonable estimates.

[5]Typical estimates range from 0.80 to 0.95.

moments and the resulting estimates. Note that the risk parameters vary across cells more than the utility parameters. We will add family size and additional age categories to improve the fit of the model in future versions.

Simulation

Each cell is characterized by a set of parameters for utility and risk, and a set of insurance options. For those not offered group insurance, the choice set is

{No insurance, Public Insurance, Non-group Private,}

while for those offered it is

{No insurance, Public Insurance, Non-group Private, Group Insurance}

The characteristics (price, cost-sharing, etc.) of these options will vary across cells according to employer choices (such as the decision to offer insurance and how much of the premium to pay) and public insurance eligibility.

Taking estimated (calibrated) parameters as given, we run the simulation model for cell i and option j, which yields a simulated utility, δ^*_{ij}. This represents the component of expected utility that can be explained by the simulation model. We repeat this for cells $i = 1, ..., I$ and for the range of options $j = 0, ..., J_i$ in the choice set for cell i, including option 0, which is no insurance of any type. To account for unobserved factors, including behavioral factors and within-cell heterogeneity, we include option-specific random taste shocks, a_{ij}. If this were a conventional discrete choice model, the variance of the shocks would need to be normalized and the coefficients of the model scaled to be in the same units. Here the scaling of the δ^*_{ij} from the simulation model are in the (arbitrary) units of the utility function. We assume that expected utility V^*_{ij} is a function of δ^*_{ij}, option- and cell-specific dummy variables Δ_{ij} and the random taste shocks:

$$V^*_{ij} = EU^*_{ij} + a_{ij} = \beta_i \delta^*_{ij} + \xi' \Delta_{ij} + a_{ij}, \quad \text{for} \quad i = 1, ..., N \ \text{ and } \ j = 0, ..., J_i. \tag{2}$$

The scale parameter β_i adjusts the scale of δ^*_{ij} so that it is in the same units as the (normalized) taste shocks. We allow the scaling to vary across cells. We then estimate the unknown parameters of (2) using a multinomial discrete choice model. Inclusion of a rich set of dummy variables enables us to account for cell and option specific factors that the simulation model has not captured. We are then able to match observed market shares fairly closely (figure?).

Finally, having simulated the market shares P_{ij} as the probability that consumer i chooses option j from data, we add the HSA option to the choice set, simulate $V^*_{i,HSA}$ for each cell, and recompute the probabilities. Comparison of these probabilities with the initial set allows us to predict the effects of HSA introduction.

Distributional Assumptions

If we assume the option-specific shocks a_{ij} are independent (over both i and j) and identically distributed Type 1 extreme-value random variables then model (2) is the conditional logit. The probability that individual i chooses option j is :

$$P_{ij} = \frac{e^{V^*_{ij}}}{\sum\limits_{k=0}^{J_i} e^{V^*_{ik}}}. \tag{3}$$

The advantage of the conditional logit model is the simple, closed form of the probabilities (3) and the straightforward way to add new options to the choice set. If cell i initially has J_i options but a new option $J_i + 1$ with utility V^*_{i,J_i+1} is added, then the new probabilities are easily computed as

$$P'_{ij} = \frac{e^{V^*_{ij}}}{\sum\limits_{k=0}^{J_i+1} e^{V^*_{ik}}}. \tag{4}$$

However, as is well known, the relative probability–or relative market shares–of any two options is independent of the addition of the new option, since the ratio reduces to

$$\frac{P'_{ij}}{P'_{ik}} = \frac{e^{V^*_{ij}}}{e^{V^*_{ik}}}$$

for any j and k. If options j and k are in the original choice set, independence of irrelevant alternatives (IIA) implies that their *relative* probabilities will remain unchanged regardless of the nature of the new option. This is a consequence of assuming that the taste shocks a_{ij} are independent.

Multinomial Probit Alternative

A more flexible alternative is the multinomial probit allowing for correlation of errors across options. This will reduce the influence of distributional assumptions substitution patterns. Given the set of δ^*_{ij}, we estimate (2) assuming $\mathbf{a_i} = (a_{i0}, a_{i2}, ..., a_{1,J_i}) \sim N(\mathbf{0}, \mathbf{\Sigma})$ on observed market shares to get estimates of the scale parameters (β_i), dummy variable coefficients, the baseline market shares P_{ij}, and the covariance matrix Σ. We then augment the covariance matrix to form $\hat{\Sigma}^A$ using plausible values for correlations between the original options and the new HSA options. We then use Monte Carlo simulation to estimate the new probabilities, P'_{ij}.

The simplest way to do this is as follows. For each cell i take a draw $\epsilon_{\mathbf{it}}$ from the multivariate normal random distribution $N(\mathbf{0}, \hat{\mathbf{\Sigma}}^{\mathbf{A}})$ to form

$$\hat{V}^*_{ijt} = \hat{\beta}_i \delta^*_{ij} + \xi' \widehat{\Delta}_{ij} + \epsilon_{ijt}$$

for each option in the new HSA-augmented choice set for that cell. The predicted choice for each draw is the option that yields the highest $\hat{V}$. Repeat this for T draws of $\epsilon_{\mathbf{i}}$ and compute new market shares P'_{ij} as the frequency with which each option is chosen.

5 Results

The first set of results is from the conditional logit model, given in Table 2. The top panel gives the results for the"not-offered" market, households that do not have the option of employer-provided insurance. Each row corresponds to a demographic category defined by age and income. The percent uninsured, publicly insured, and privately (non-group) insured predicted by the model are given in columns (1), (2), and (3), respectively. Columns (7)-(10) recompute the shares assuming a non-group "HSA policy" is introduced, which includes both a high-deductible policy (30 percent coinsurance, $3,000 deductible, loading factor of 0.40) and a tax-preferred health savings account. We see that HSAs tend to be more popular as income increases. For example, only 8.6 percent of the lowest decile young households choose an HSA while 35.4 percent of the highest decile households choose that option. This switch to HSA draws from all three previous options. For the lowest decile, the percent uninsured declines from 47.7 percent to 43.6 percent; the percent publicly insured declines from 45.6 percent to 41.6 percent; and "traditional" insurance falls from 6.7 percent to 6.1 percent.

The "old" display similar patterns: the HSA market share tends to increase with income, although it is not monotonic, and this share draws from all three original categories. Somewhat surprising is that there is not a marked difference in HSA shares between young and old. At the lower income deciles, market shares among the old tend to be larger than for the young, at upper deciles it reverses, but the differences don't tend to be particularly large.

Table 3 gives the results for individuals who have the option of employer-provided insurance. These households are modeled as having four options: no insurance, public insurance, non-group insurance, and group insurance. The simulation here adds two HSA policies, one as a non-group option and one as an alternative to the group policy. Both the HSA policies are identical except for the loading factors, which match with their respective non-HSA alternative.

The patterns are roughly similar to the results in the not-offered market. Adding in the HSA options draws heavily from the existing insurance options, but it also decrease the percentage uninsured, and the percentage on public insurance. Participation tends to increase with income, although the lowest decile for the old has an anomalously high participation rate.

Given the preliminary nature of these estimates, it would be a mistake to be dogmatic about their implications. But they are certainly suggestive that HSAs have the potential to decrease the ranks of the uninsured. Weighting by each income decile, the percentage uninsured declines from 58.7 percent to 52.4 for households not eligible for group insurance, and from 7.3 percent to 4.3 percent for those who are.

6 Policy Simulations and Extensions

This work is very preliminary, but we believe it is promising. In this section we outline the next steps in developing a policy-useful model for estimating the impact of various health-related proposals.

Our first extension will be to allow for non-independent error terms by using estimates from a multinomial probit model to compute market shares. As mentioned previously, the conditional logit model implies strong assumptions about the addition of new products. Multinomial probit relaxes those assumptions, but at the cost of a substantial increase in computational complexity.[6]

Once we are satisfied with the basic estimation strategy, we will then focus on firm-level decisions. In particular, we will explore the implications for pooling equilibria at the firm level. To incorporate firm decisions into our framework, our initial strategy will be to model the economy as having three types of firms: 1) large firms that offer insurance, 2) small firms that offer insurance, 3) small firms that do not offer insurance. Based on observable data in the MEPS, we will model each firm type as being composed of individuals with various demographic characteristics, each modeled as in our pure representative agent framework. Then we will use use existing data to construct a typical insurance contract (premium, coinsurance rate, deductible, loading, cost-sharing) for firms of Type 1 and 2 and proceed to estimate a baseline model given observed insurance status. Each firm type forms a separate risk pool, so the average premium will be applied to all members of the group who choose to be insured. Once a baseline is estimated, we will proceed with the simulation of introducing HSAs (or some other policy proposal) to estimate the impact on the overall pool, and on the outcomes of interest such as take-up rates and levels of insurance[7]. Of particular interest will be scenarios under which 'death spirals' might arise (Cutler and Zeckhauser, 1997).

Although our primary purpose is to estimate the effect of various proposals on the number of uninsured, our framework is sufficiently flexible to handle a variety of simulation tasks. For example, unlike most available methodologies discussed in our review section, our framework estimates health expenditures jointly with insurance choice. It would also be possible to estimate standard economic welfare effects, quantifying the gains and losses to various policies measured in terms of consumer surplus. Additional micro-level simulations concerning firm choice of optimal policies for various demographic compositions of employees are also feasible within this framework.

[6]We will actually be using simulated maximum likelihood due to the high number of options. Standard maximum likelihood estimation of multinomial probit is impractical for more than five alternatives

[7]The estimation strategy is similar to that used in Cardon and Showalter (2007), except this work will include several categories of consumers, not just two.

References

Cardon, James H. and Igal Hendel (1999). "Asymmetric Information in Health Markets: Evidence from the National Medical Expenditure Survey. *RAND Journal of Economics*. Vol. 32, No. 3, Autumn, 2001, pp.408-427.

Cardon, James H. and Mark H. Showalter (2001). An Examination of Flexible Spending Accounts. *Journal Of Health Economics*. Vol. 20, No. 6 (2001), pp. 935-954.

Cardon, James H. and Mark H. Showalter (2003). "Flexible Spending Accounts as Insurance." *The Journal of Risk and Insurance*. Vol.70, No. 1, 43-51.

Cardon, James H. and Mark H. Showalter (2007). "Insurance Choice and Tax-Preferred Health Savings Accounts Journal of Health Economics." *Journal of Health Economics*. Vol. 26, No. 2 (2007),pp. 373-399.

Cutler, David M., and Richard J. Zeckhauser. (1997) "Adverse Selection in Health Insurance". NBER Working Paper No. 6107.

Furman, Jason. (2006) "Expansion inf HSA Tax Breaks is Larger–and more Problematic–Than Previouly Understood." *Center on Budget and Policy Priorities.* Revised February 7, 2006. URL: http://www.globalaging.org/health/us/2006/BudgetHSA.pdf.

Glied, Sherry A. and Dahlia K. Remler (2005)."The Effect of Health Savings Accounts on Health Insurance Coverage." Commonwealth Fund Publication #811, 2005.

Glied, Sherry, Dahlia K. Remler and Joshua Graff Zivin (2002). "Inside the Sausage Factory: Improving Estimates of the Effects of Health Insurance Expansion Proposals." *The Milbank Quarterly*. Vol. 80, No. 4, 2002, pp.603-636.

Economic Report of the President, Chapter 10, "Health Care and Insurance." Council of Economic Advisers, February 2004.

Feldman, R., Parente, S.T., Abraham, J., Christianson, J., Taylor, R. (2005). "Health Savings Accounts: Early Estimates of National Take-up from the 2003 Medicare Modernization Act and Future Policy Proposals. November 18, 2005. *Health Affairs.*

Hamilton, Barton H. and James Marton (2006). "Employee Choice of Flexible Spending Account Participation and Health Plan," University of Kentucky working paper, February.

Nevo, Aviv (2000). "A Practitioners Guide to Estimation of Random Coefficients Logit Models of Demand, *Journal of Economics & Management Strategy,* 9(4), 513-548, 2000.

Moon, M. L. Nichols, and S. Wall (1996). "Medical Savings Accounts: A Policy Analysis." Urban Institute: http://www.urban.org/pubs/hinsure/msa.htm.

Pauly, M. and B. Herring (2000). "An efficient employer strategy for dealing with adverse selection in multiple-plan offerings: an MSA example." *Journal of Health Economics*, Vol. 19 (4) pp. 513-528.

Pauly, M. and B. Herring (2002). *Cutting Taxes For Insuring: Options and Effects of Tax Credits for Health Insurance.* Washington DC.: AEI Press.

Petrin, Amil (2002). "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy.* Vol. 110, No. 4, 2002, pp. 705-729.

Remler, Dahlia K., Joshua Graff Zivin, and Sherry A. Glied. (2002) "Modeling Health Insurance Expansions: Effects of Alternative Approaches." NBER Working paper 9130. http://www.nber.org/papers/w1930.

Zabinski, D., T. Selden, J. Moeller, and J. Banthin (1999), "Medical Savings Accounts: Microsimulation Results From a Model With Adverse Selection," *Journal of Health Economics*, Vol. 18, pp. 195-218.

Table 1: Method of Moments Estimation

Theoretical Moments	Sample Moments	Young	Old
$E(X^*) = \frac{Y+KE(\theta)}{K+\alpha}$	$\overline{X}$	5.6297	9.3854
$Var(X^*) = (\frac{K}{K+\alpha})^2 Var(\theta)$	S	10.4133	14.9242
$\varepsilon_p = \frac{\alpha\big((\alpha E(\theta)-I)\frac{dK}{d\alpha}-(I+KE(\theta))\big)}{(K+\alpha)(I+KE(\theta))}$	$\widehat{\varepsilon}_p^R$	-0.15	-0.15
$\varepsilon_I = \frac{I}{K+KE(\theta)}$	$\widehat{\varepsilon}_I^R$	0.2	0.2

Estimates

	Young	Old
$\widehat{\mu}$	4.7194	7.8454
$\widehat{\sigma}$	10.4743	15.0220
$\widehat{p}$	0.1055	0.0504
$\widehat{\lambda}$	0.0015	0.0017
$\widehat{\gamma}$	-0.3451	-0.3468

Notes: $X^* = \frac{Y+K\theta}{K+\alpha}$, where Y is income and $K = (\frac{\gamma}{\alpha})^{1-\lambda}$.

$\theta \sim Lognormal(\mu, \sigma^2)$

$\widehat{\varepsilon}_p^R$ and $\widehat{\varepsilon}_I^R$ are estimates from the RAND Health Insurance Experiment.

$U(C,H) = \frac{C^\lambda - 1}{\lambda} + \gamma\frac{(X-\theta)^\lambda - 1}{\lambda}$

Table 2--Estimates with and without Health Savings Accounts/High Deductible Policies
--No Group Insurance Available--

	NO HSA			HSA Available			
Income Decile	(1) Uninsured	(2) Public	(3) Non-Group	(4) Uninsured	(5) Public	(6) Non-Group	(7) Non-group HSA
"Young"							
1	0.477	0.456	0.067	0.436	0.416	0.061	0.086
2	0.487	0.473	0.039	0.463	0.449	0.037	0.051
3	0.704	0.229	0.067	0.645	0.210	0.062	0.084
4	0.729	0.203	0.068	0.671	0.187	0.062	0.080
5	0.731	0.134	0.135	0.624	0.114	0.115	0.147
6	0.746	0.106	0.149	0.630	0.089	0.126	0.155
7	0.645	0.152	0.203	0.515	0.122	0.162	0.201
8	0.681	0.060	0.260	0.518	0.045	0.198	0.239
9	0.555	0.076	0.369	0.385	0.053	0.256	0.305
10	0.463	0.064	0.472	0.299	0.042	0.305	0.354
"Old"							
1	0.643	0.304	0.053	0.559	0.265	0.046	0.131
2	0.473	0.471	0.056	0.437	0.436	0.052	0.075
3	0.612	0.316	0.072	0.557	0.288	0.066	0.090
4	0.654	0.278	0.068	0.600	0.255	0.062	0.084
5	0.671	0.195	0.134	0.569	0.166	0.114	0.152
6	0.709	0.148	0.144	0.598	0.125	0.121	0.156
7	0.615	0.198	0.187	0.496	0.159	0.151	0.194
8	0.672	0.078	0.250	0.512	0.059	0.190	0.239
9	0.558	0.090	0.352	0.391	0.063	0.247	0.299
10	0.471	0.076	0.454	0.307	0.049	0.296	0.347

Notes: Estimates from conditional logistic regression. Model includes dummies for expected utility, age, income decile, and interaction terms. Unit of observation is a household as defined in the MEPS.

Table 3--Estimates with and without Health Savings Accounts/High Deductible Policies
--Group Insurance Available--

	NO HSA				HSAs Available					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Income Decile	Uninsured	Public	Non-Group	Group	Uninsured	Public	Non-Group	Group	Non-group HSA	Group HSA
"Young"										
1	0.077	0.031	0.020	0.872	0.074	0.030	0.020	0.848	0.028	0.000
2	0.250	0.059	0.015	0.676	0.175	0.041	0.010	0.473	0.014	0.287
3	0.195	0.039	0.026	0.740	0.124	0.024	0.016	0.470	0.022	0.343
4	0.148	0.032	0.006	0.814	0.089	0.019	0.004	0.490	0.005	0.393
5	0.112	0.017	0.005	0.867	0.064	0.010	0.003	0.496	0.004	0.424
6	0.082	0.011	0.004	0.903	0.045	0.006	0.002	0.498	0.003	0.446
7	0.060	0.006	0.006	0.928	0.032	0.003	0.003	0.495	0.004	0.463
8	0.043	0.005	0.009	0.943	0.023	0.003	0.005	0.492	0.006	0.472
9	0.019	0.005	0.007	0.969	0.010	0.003	0.003	0.497	0.004	0.483
10	0.013	0.010	0.011	0.966	0.006	0.005	0.006	0.490	0.006	0.486
"Old"										
1	0.062	0.013	0.010	0.915	0.030	0.006	0.005	0.435	0.013	0.511
2	0.156	0.038	0.014	0.793	0.096	0.023	0.009	0.488	0.012	0.373
3	0.118	0.037	0.019	0.826	0.071	0.022	0.012	0.500	0.016	0.380
4	0.099	0.033	0.005	0.863	0.058	0.019	0.003	0.506	0.004	0.410
5	0.077	0.018	0.004	0.901	0.043	0.010	0.002	0.509	0.003	0.432
6	0.061	0.012	0.003	0.924	0.034	0.007	0.002	0.507	0.002	0.449
7	0.047	0.007	0.004	0.942	0.025	0.004	0.002	0.506	0.003	0.460
8	0.035	0.005	0.007	0.952	0.018	0.003	0.004	0.500	0.005	0.470
9	0.017	0.006	0.006	0.971	0.009	0.003	0.003	0.500	0.004	0.482
10	0.012	0.011	0.010	0.968	0.006	0.005	0.005	0.494	0.006	0.484

See explanation from Table 2

Figure 1: Actual and Predicted Market Shares