

**Modeling Health Insurance Expansions:
Effects of Alternate Approaches**

Dahlia K. Remler

Joshua Graff Zivin

Sherry A. Glied

Department of Health Policy and Management
Mailman School of Public Health
Columbia University

This draft: April 2002

Preliminary

Please do not distribute or cite without authors' permission

The authors thank the Robert Wood Johnson Foundation for funding; Bowen Garrett for many extremely helpful discussions; Adam Atherly, Linda Bilheimer, Danielle Ferry, Jonathan Gruber, John Holahan and participants in an RWJ-sponsored conference for helpful comments and/or discussions about modeling; and Danielle Ferry for excellent research assistance.

Abstract:

Comparison of policy proposals to increase the number of people with health insurance often focus on predictions of the numbers of individuals that will be insured and the costs of the proposals. Several different simulation modeling methods are used to predict these effects, making the predictions difficult to compare. In this paper, we do the following: (1) We categorize and describe the different approaches used; (2) we explain the conceptual and theoretical relationships between the methods; (3) we demonstrate empirically an example of the (quite restrictive) conditions under which all approaches can yield quantitatively identical predictions; and (4) we empirically demonstrate conditions under which the approaches diverge and the quantitative extent of that divergence. All modeling approaches implicitly make assumptions about functional form that impose restrictions on unobservable heterogeneity. Those assumptions can dramatically affect the quantitative predictions made.

I. Introduction

Estimates of the costs and consequences of many types of public policy proposals play an important role when policymakers decide whether to proceed with an option. They also figure centrally in the deliberations that go into the development of policy options. Unfortunately, different estimates of the same, or similar, policies can yield widely divergent results, depending on which model is used to make the estimates. Furthermore, it is often difficult for even a well-trained observer to understand why the results diverge so much. Health insurance expansion proposals are a particularly important and complex example (Glied, Remler and Graff Zivin 2001).

Simulation models need to capture whatever features of the insurance purchasing process might affect the costs of an expansion. Since we do not understand (and could not practically model, even if we did) all features of the purchasing decision, every modeler must make simplifying assumptions to carry out the simulation task. Foremost among these simplifying assumptions is the framework of the model itself. One reason for the opacity of the simulation process – and perhaps for the observed diversity in findings -- is that there are many different, accepted, frameworks for simulation modeling. Since different models use different frameworks, it is hard to see what drives the difference in their results. Moreover, the modeling approaches themselves may introduce distinct biases that generate the variation in outcomes.

The purposes of this paper are to (1) explain conceptually how the most widely used modeling approaches are related to one another; (2) to show how the quantitative values of the parameters in the various methods can be directly compared; (3) to empirically illustrate the (quite restrictive) conditions under which all approaches yield identical results; and (4) to empirically demonstrate some common conditions under which the approaches diverge and the quantitative extent of that divergence. We will show that, theoretically, all the methods are tightly related, but that in many circumstances, the choice of functional form, and by extension, decisions about how to model unobserved heterogeneity, can have substantial effects on results. The paper is organized as follows. Section II will formally demonstrate the relationships among reservation prices, elasticities and take-up rates. Section III briefly describes current modeling approaches and the conditions for their equivalence. Section IV empirically illustrates the

conditions under which the approaches are equivalent and under which they diverge. Section V concludes.

II. Heterogeneity and Measuring Insurance Uptake

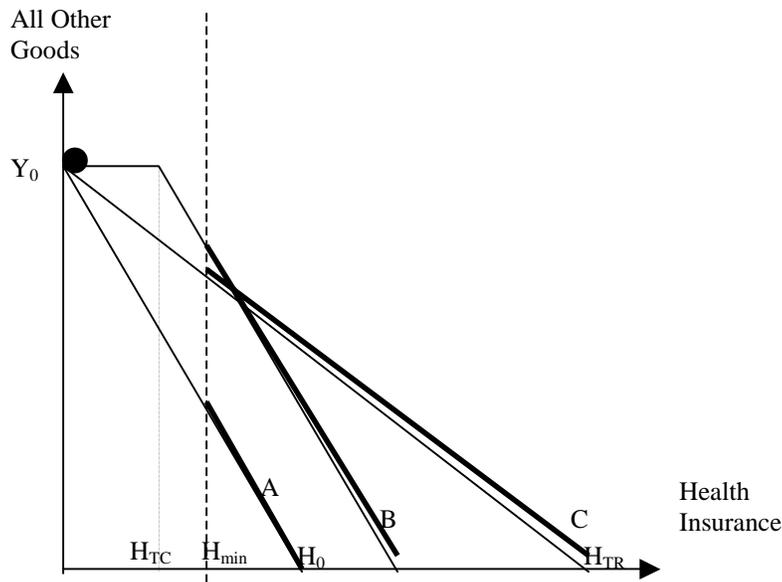
The underlying decision to enroll in any given insurance program is an individual level discrete choice, in which one of several insurance options is selected. For ease of exposition, we will abstract from issues of different forms of insurance (ESI, Medicaid, non-group, etc.), different qualities of insurance, and issues of family structure. We consider only single individuals with a single insurance option and measure whether or not individuals have health insurance, not the quantity of that health insurance.

This dichotomous approach differs from the continuous utility maximization approach typically used in public policy modeling. In the continuous approach, the effect of different types of health insurance expansions (proportional vs. flat tax credits, for example) is illustrated by comparing the effects of linear and kinked rotations of the budget constraint. Since we are modeling participation, this detail is generally irrelevant.¹ We illustrate this with a diagram below.

Figure 1 shows the budget constraints for three health insurance regimes: no favorable treatment of health insurance, a tax deduction and a tax credit, A, B, and C, respectively. Health insurance is plotted on the x-axis with all other goods on the y-axis. The purchase price for a minimum benefits health insurance policy, H_{\min} , is shown as a dotted line. The accessible parts of the budget constraints are shown by the bold lines and the mass-point at $(0, Y_0)$, where all the uninsured cluster. The effect of an expansion program on participation is determined by how many people move from the mass-point of zero health insurance to any place beyond H_{\min} . In this case, the form of the subsidy does not matter: a tax-deduction and a tax-credit of equal size will have the same effect. Rather than a budget constraint and indifference curve framework, discrete choice problems such as these are best represented through a reservation price framework, as we describe below.

Figure 1

¹ The analysis below assumes that the value of any tax credit is less than or equal to the cost of the minimum health insurance package.



Consider a population of observationally identical individuals. The population is homogeneous in all observable variables, such as income, employment status, etc. We assume that everyone faces a single price for health insurance. The population is heterogeneous only in unobservable variables, such as unobservable aspects of health status, attitudes towards risk, opportunities for free care, non-health economic needs, family support, etc. Such heterogeneity causes observationally identical individuals to make different decisions when faced with the same health insurance purchase choice. Thus, some members of our population are insured at the market price while others are not.

The most intuitive way to represent such heterogeneity is through heterogeneity in the reservation price. The reservation price is the maximum amount that a given individual would be willing to pay for health insurance. The distribution of reservation prices in a population determines the number insured under any conceivable policy regime. This distribution is sufficient to fully characterize the market level demand for health insurance in this population.

Clearly, the population distribution of reservation prices can take on many shapes. While we illustrate the distribution functions graphically using the uniform distribution, as shown in Figure 2, all discussion and formulas will be completely general, unless otherwise stated. If the market price is above the maximum reservation price in the population, r_{\max} , no one is insured. If the market price is below r_{\min} , everyone in the population is insured. At market prices between

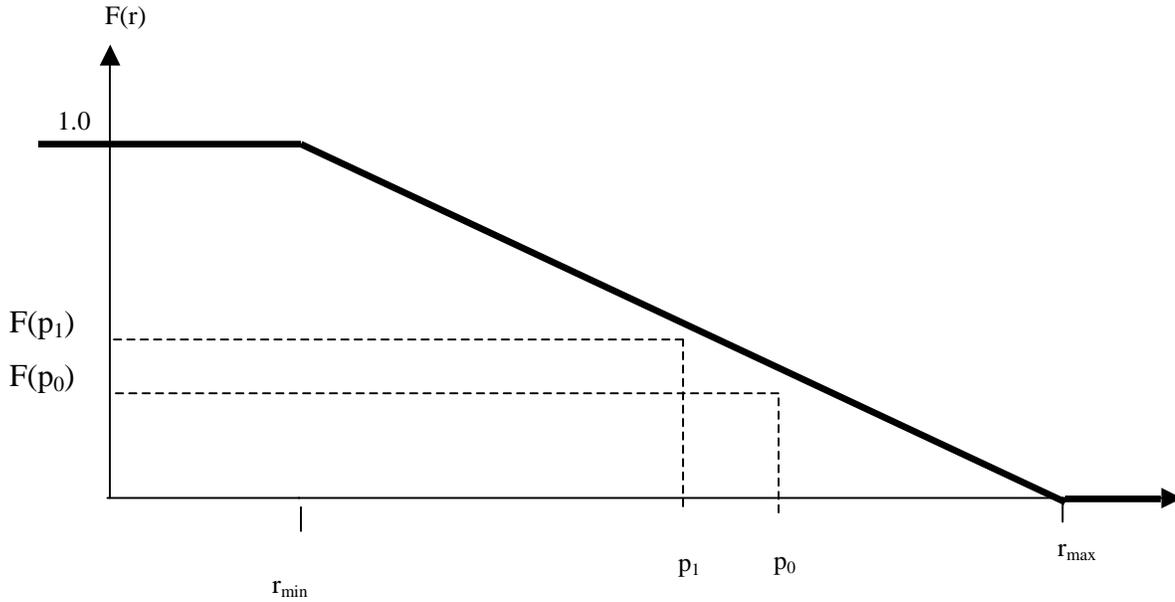
r_{\min} and r_{\max} , the share of the population for whom the market price p_0 is below their reservation price will be insured.

Figure 2: Uniform distribution of reservation prices.



Another way of representing the same distribution is through the cumulative distribution function, $F(r)$. Since at prices below r_{\min} , everyone is insured, the cumulative distribution function takes on the value 1. At prices above r_{\max} it is zero, since no one is insured. For any given market price, the cumulative distribution function reveals the share of the population that will purchase insurance. Coupled with the total number of people in the population, this determines the number of insured. Thus, the cumulative distribution function is, in essence, the market demand function. Lowering the price from p_0 to p_1 raises the share of the population that is insured from $F(p_0)$ to $F(p_1)$.

Figure 3: Cumulative distribution function of reservation prices, for a uniform distribution.



The distribution of reservation prices can be used to derive an elasticity measure. It is important to note, however, that while health insurance acquisition is an individual discrete choice, elasticity, as used in health insurance modeling, is a market-level concept. The demand elasticity represents a change in the number of individuals insured, not a change in the quantity of insurance purchased, as would be described by the elasticity of an individual demand curve. In this context, the elasticity should be thought of as a participation elasticity, which is only meaningful when applied to a population. A point elasticity, the percentage change in the number of insured per percentage change in price at a given price, can be expressed as:

$$\text{Point elasticity: } \varepsilon = \frac{p}{F(p)} \frac{dF(p)}{dp} = \frac{pf(p)}{F(p)}$$

It is often convenient to use a semi-elasticity, which is used in some of the literature (e.g. Gruber and Poterba 1994). A semi-elasticity does not use the percentage change in the population insured but rather the *percentage point* change in the number insured. The point semi-elasticity is:

$$\text{Point Semi-elasticity: } \varepsilon_s = p \frac{dF(p)}{dp} = pf(p)$$

In some empirical work and simulations, demand responsiveness is measured using a discrete price change. Since lowering price to p_1 increases the share of the entire population who takes-up insurance by $F(p_1) - F(p_0)$ percentage points²,

$$\text{Discrete Elasticity} \quad \varepsilon = \frac{(F(p_1) - F(p_0)) / F(p_0)}{(p_1 - p_0) / p_0}$$

$$\text{Discrete Semi-elasticity} \quad \varepsilon_s = \frac{F(p_1) - F(p_0)}{(p_1 - p_0) / p_0}$$

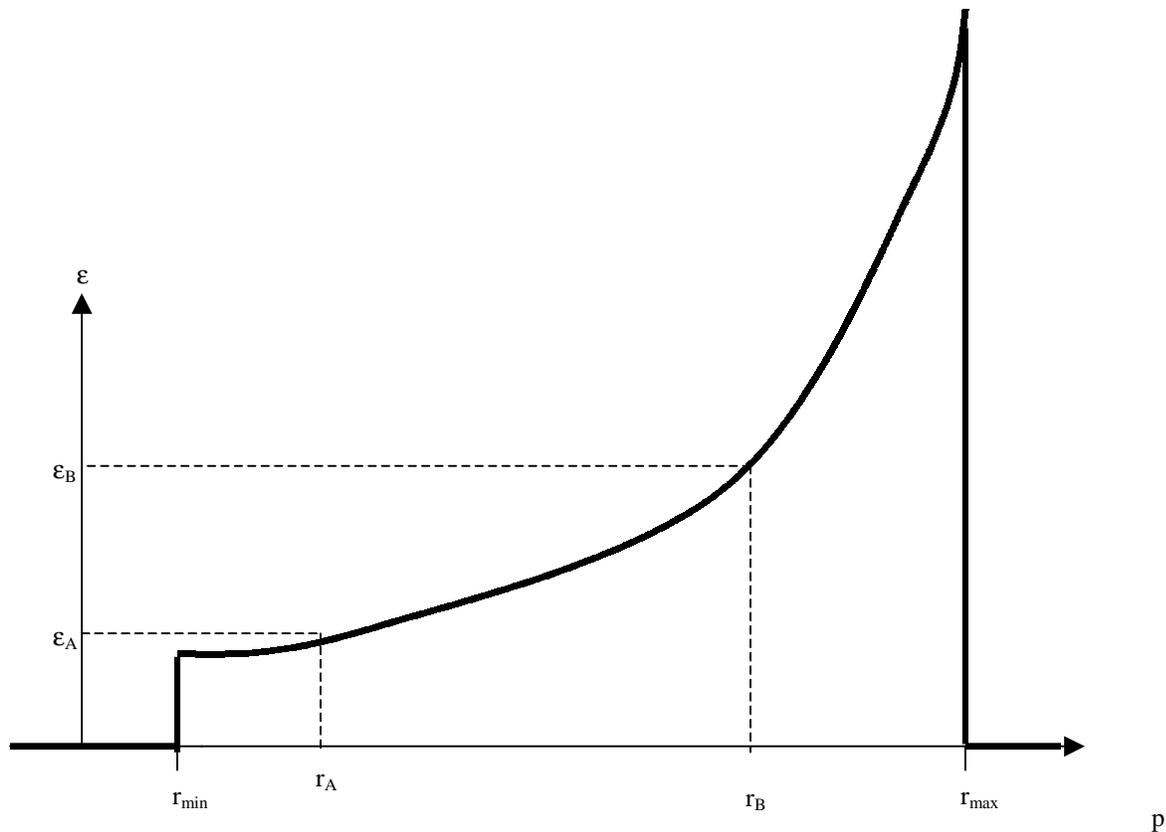
The elasticity varies with price as we move through the reservation price distribution. The extent and form of this variation depends on the shape of the distribution function. Depending on that shape, discrete change elasticities can be strikingly different from the point elasticity at the starting price, as we will illustrate in the empirical section. Further, we would expect very different elasticities at the floor where almost no one is insured and at the ceiling where almost everyone is insured. When almost everyone who is observationally identical has taken up insurance, it seems highly likely that the remaining few uninsured are different along some important unobserved dimensions that influence the participation decision.

Taking an elasticity estimated at the low part of the distribution and applying it to those at the high point of the distribution can generate significant errors. For example, in the case of a uniform distribution, the elasticity is zero at prices below r_{\min} , and above r_{\max} , since either everyone or no one is insured and nothing further changes. Between those two prices, the uniform distribution elasticity is $\frac{p}{r_{\max} - p}$, as illustrated in Figure 4³. As drawn, an elasticity measure derived from prices in the vicinity of r_A but applied to individuals with reservation prices in the vicinity of r_B would be off by roughly a factor of three.

² It would be possible and preferable in real empirical work, of course, to define arc elasticity versions of the discrete elasticities. Since that issue is not central to our points, we avoid the arc elasticities to keep the formulas simpler.

³ Due to the unrealistic abrupt changes at the edges of the distribution, the elasticity actually goes to infinity at r_{\max} .

Figure 4: Possible (point) elasticity as a function of price for a uniform distribution of reservation prices.



The elasticity is not the only measure of consumer responsiveness. Some simulation methods use take-up rates instead. The take-up rate is the share of the initially uninsured population who take up insurance in response to some policy or price change. Imagine that our homogeneous population was initially ineligible for a free program and faced a non-zero market price for insurance. For simplicity, consider the effect of making the population eligible for free health insurance, resulting in a 100% price reduction.⁴ In that case, the semi-elasticity, denoted ϵ_s , is $F(0) - F(p_0)$, or the fraction of the *entire* population who take-up insurance. The take-up rate is the fraction of the *uninsured* population newly insured, $(F(0) - F(p_0)) / (1 - F(p_0))$.

⁴ Non-financial costs of insurance are not included in the price and would need to be incorporated elsewhere.

In general, the take-up rate for an arbitrary price change is

$$\tau = \frac{\epsilon_s}{1 - F(p_0)} \left(\frac{p_1 - p_0}{p_0} \right) = \frac{\epsilon}{1 - F(p_0)} (F(p_0)) \left(\frac{p_1 - p_0}{p_0} \right)$$

If the initial group is all uninsured, the semi-elasticity and the take-up rate are identical. Applying an elasticity computed from a total population (including both the insured and uninsured) to an uninsured population, as if it were a take-up rate, creates a scaling error⁵. To illustrate the importance of scaling the measured elasticity correctly when applying it to the uninsured population only, consider the following simple example.

Imagine a simple controlled experiment in which there are 10,000 people in the experimental group of whom 8000 are initially insured. Suppose the experiment reduces the price of health insurance faced by this group by 10%. In response, 1000 of those initially uninsured purchase coverage. In this case, the estimated semi-elasticity is equal to the change in probability of having insurance = (9000-8000)/10,000 = .1 divided by the price change of .1, yielding a semi-elasticity = .1 / .1 = 1.0. (The corresponding elasticity is 1.25.)

Now consider what happens if this semi-elasticity is applied to the initially uninsured population only, as would be the case with a take-up rate. Applying the 1.0 semi-elasticity to the 10% price change over a population of 2000 initially uninsured people yields an estimate that 200 people will gain new coverage. In the example, however, 1000 people gained coverage, 5 times as many as we estimate. Using a semi-elasticity as if it were a take-up rate will lead to inconsistencies. In order to be applied to a population that is entirely uninsured, the semi-elasticity must be scaled by the inverse of the share of the population that was initially uninsured, in this case 1/5.

On theoretical grounds, there is no reason to prefer elasticities over take-up rates or vice versa. However, if predictions across studies are to be compared, transparent and consistent descriptions of how parameters are employed is critical. More generally, in order to compare simulation results, we need to understand the origins of differences and the circumstances under which different modeling approaches would yield identical results.

⁵ Some researchers use the term take-up rate to refer to an average take-up rate, the share of all people who participate in a program. This rate is then applied only to the uninsured. The previously insured, albeit with a different form of insurance, have a different take-up rate (a crowd-out rate) applied to them. If not everyone in the initial population had been uninsured, then a similar scaling error occurs. In our simple model with only one form of insurance, the concept of crowd-out is not relevant. As such, we define the take-up rate as the share of the uninsured who take-up insurance, since that is how it will be applied.

III. Modeling Approaches in Practice

Methods of modeling the effects of health insurance reform fall roughly into four broad categories. Note that all implementations are not always pure forms of these categories and that hybrids exist (Blumberg and Nichols 2000⁶). The first approach, which we call the elasticity approach (EA), applies price elasticities to data on current prices and current insurance patterns. This approach is the most familiar and widely used. In some cases (for example, Gruber 2000 and Baumgardner 1998), price elasticities are obtained from data sources other than the one used for the microsimulation, while in other cases (for example, RWJF 2000), the elasticities are estimated from the same data used for the microsimulation. The EA uses individual-level data in determining prices, incomes and other parameters, and is thus considered a microsimulation approach. However, while the prices and incomes are individual-specific in the EA, the elasticity is not.

The second approach, which we call the matrix approach (MA), applies take-up rates to groups of individuals defined by particular characteristics, such as income, family-size, and so on. This approach has primarily been used for Medicaid take-up calculations (e.g. Holahan, Uccello and Feder).

We call the third approach the discrete choice approach (DCA). In this method, a discrete choice model is estimated and used to predict the effects of policy changes (e.g., Blumberg and Nichols; Pauly and Herring 2000a and 2000b). While only Pauly and Herring interpret such models as reservation prices, all discrete choice models implicitly define a cumulative distribution function of reservation prices. Thus, these estimated discrete choice models are conceptually akin to the insurance demand model we described in the previous section.

The final approach is the most recent. It compares health insurance options using reservation prices estimated through a utility-based structural model (Pauly and Herring, 2000a; 2000b; Zabinski et al 1999). As we show below, the first three approaches, while appearing very different and, in existing practice, yielding quite different estimates, are actually closely related. The fourth approach is idiosyncratic in the particular structural model employed and difficult to compare to the other three and therefore we do not discuss it any further.

⁶ Blumberg and Nichols' simulation methods also incorporates a novel and rather sophisticated set of general equilibrium conditions, which are beyond the scope of our discussion here.

The discussion in the previous section assumed an observationally homogeneous population. In the matrix approach, however, observable covariates play an important role. Therefore, to illustrate the relationships between the three approaches, we need a formulation with observationally distinct groups. Consider a case with two groups, one, denoted E, is already eligible, while the other denoted N, is initially ineligible. Each group initially contains both insured and uninsured individuals. The take-up rate among the uninsured is distinct for each subgroup.⁷ The market price for insurance is initially p_0 and is p_1 in the subsidized program.

As we have emphasized in the previous section, elasticities vary across individuals and depend on the range of price changes. Each approach, therefore, will be applied to identical people and an identical price change. In this case, the price elasticity of demand for insurance is specific to the relevant population group, N, and calculated based on the discrete price change that corresponds to the policy proposal. Thus, the semi-elasticity for the N group is:

$$\epsilon_s^N = \frac{F_N(p_1) - F_N(p_0)}{p_1 - p_0} \cdot p_0.$$

Applying that semi-elasticity in an EA microsimulation to the entire

N population would result in a change in the number of insured of:

$$\Delta T(EA) = \sum_{i \in N} \epsilon_s^N \left(\frac{p_1 - p_0}{p_0} \right) = N_N (F(p_1) - F(p_0))$$

where N_N is the number of people in the N group.

The matrix approach uses take-up rates, rather than elasticities. A take-up rate is the fraction of the uninsured that participate in the new program. Calculating the take-up rate for the relevant population group, N, based on the same discrete price change of the policy proposal,

$$\text{results in } \tau_N = \frac{F(p_1) - F(p_0)}{1 - F(p_0)}.$$

Using the matrix approach, the take-up rate would be applied to

the uninsured members of the newly eligible group, resulting in:

$$\Delta T(MA) = \sum_{i \in U_{0,N}} \tau_N = N_N (F(p_1) - F(p_0))$$

The discrete choice approach starts by estimating a discrete choice model, generally a logit or probit. To predict enrollment, the fitted model is applied to each individual to predict his

⁷ For the matrix approach, there is also generally a different take-up rate for programs such as Medicaid among the already insured. Since we have only two states, insured and uninsured, that is not relevant here. Our already insured simply stay insured.

probability of participating at the new price. This is equivalent to estimating the cumulative distribution function for each individual and adding them up.

$$\Delta T(DCA) = \sum_{i \in UI_0, N} F(p_1) - F(p_0) = N_N(F(p_1) - F(p_0))$$

Clearly, if all parameters are estimated from the affected population facing the proposal to be simulated, then all approaches yield the same predictions. This begs the question of why empirical estimates derived using different methods generally produce inconsistent results.

One reason estimates differ is that they use parameters drawn from different sources. Occasionally, parameters are estimated from the same population to which a proposal will be applied. Sometimes, parameters are estimated from a larger sample that includes the small affected population (for example, a discrete choice model that uses the full population to estimate parameters more precisely and then applies these to an expansion targeted at those <250% FPL). Alternatively, parameters may be drawn from an entirely different population (for example, an elasticity computed from an independent natural experiment) to avoid problems of selection bias and endogeneity. Since the decision to acquire insurance is based on unobservable characteristics that may be correlated with price, any estimate of price responsiveness will falsely attribute all causality to the price change. Using estimates of consumer responsiveness obtained from a different population, however, may lead to errors because the elasticity can vary substantially across the population. The employment of these external sources must balance these bias and relevance concerns. Variations in predictions may also stem from differences in the choice of functional form and, implicitly, the treatment of unobservable heterogeneity. The relationship between functional form assumptions and predictions will be elucidated in the next section.

IV. An Empirical Illustration

Defining homogeneous sub-groups in a theoretical context is simple: one simply defines subscripts. In practice, this is considerably more difficult. In effect, the individual characteristics used as either regression model controls or a means of stratification, determine the populations effectively treated as homogeneous. In this section, we employ data to demonstrate the equivalence of the discrete choice approach (DCA), elasticity approach (EA) and matrix

approach (MA) in one particular case and the various conditions under which the approaches do not result in the same predictions.⁸ We also illustrate the impact of unobservable heterogeneity by deriving the empirically predicted variation in responsiveness among observationally identical individuals using the DCA.

The demonstration of the equivalence of modeling approaches proceeds as follows. We first fit a discrete choice model for having insurance using log price and other covariates. To implement the DCA, we use that model to predict the impact of a single health insurance expansion, a reduction in price for a certain sub-group. Next, we use the results of those predictions as if they were real world data from a policy change to calculate the implied discrete-price-change elasticity and take-up rate, as described in the theory section. The EA predictions are made by applying that elasticity to each of the income categories used in the estimation. The MA predictions are made by applying the derived take-up rate to each income category.

We illustrate how the methods diverge in their predictions in three different ways. First, when implementing the EA, we use the point elasticity estimated from the discrete choice model, rather than using the discrete-price-change elasticity, as described above. Second, we re-estimate the discrete choice model on a moderately high-income sub-group and use the implied discrete choice elasticity to predict the impacts of the policy change on a low-income sub-group. Third, we more closely mimic the MA as practiced, by using take-up rates estimated for different price-income cells and applying them to those newly eligible for lower prices.

IV.A. Empirical Methods

For all calculations we will use the March 2000 Current Population Survey and include only the non-elderly (under 65) and those without Medicaid. We treat all forms of insurance, ESI, individually-purchased insurance and so on, as equivalent. For all models, we will consider a reduction of price by 50% for a certain sub-group^{9,10}, specifically those below 200% of the federal poverty line. We estimate a weighted logistic regression model,

⁸ Note that we are only examining predictions of the total number of insured, not the distribution of *who* is insured. Clearly, policymakers and individuals are also concerned about how insurance is distributed among income groups and other categories.

⁹ The prices used here are the out-of-pocket prices. There is controversy about whether the employer share should also be included for ESI. While this issue is important, it does not matter for the points we wish to illustrate.

¹⁰ Prices are assigned to individuals in the CPS by matching them to actuarial data using covariates. Non-group prices are assigned based on age, sex, Census region and health status. Group prices are assigned based on firm size and Census region.

$$\Pr(Ins_i = 1) = f_i(p_i) = \frac{\exp\{\beta \log p_i + \gamma X_i\}}{1 + \exp\{\beta \log p_i + \gamma X_i\}},$$

where the subscript i denotes the i th individual, $\log p_i$ is the log out-of-pocket price of insurance, X_i is a vector of income categories: <50% FPL, 50-99% FPL, 100-149% FPL, and 150-199% FPL and the weights are the survey sample weights. The estimated coefficients can be used to predict the probability of having insurance for each person in the sample. The predicted number of insured for the population, Ψ , is calculated by multiplying the predicted probability of each person in the sample by the person's sample weight w_i : $\Psi = \sum_i w_i f_i(p_i)$.¹¹

In our empirical exercise, we treat these predictions as if they were “real” data from an actual policy change. We use these to calculate the implied discrete semi-elasticity, $\epsilon_s = \frac{\Delta\Psi/N}{\Delta p/P}$,

and implied take-up rate, $\tau = \Delta\Psi/U$ where N is the number of people affected by the policy change and U is the initial number of uninsured affected by the policy change.

To implement the EA, we calculate the final number of insured within an income cell as: $\Psi_{cell}^f = \Psi_{cell}^0 + \epsilon_s \left(\frac{\Delta p}{p} \right) N_{cell}$, where the subscript *cell* denotes the income cell, the superscript f denotes the final (post-intervention) state, 0 denotes the initial (pre-intervention) state, and N_{cell} denotes the total number of people in the cell affected by the policy. We first implement the EA approach using the discrete price change semi-elasticity implied by the discrete choice model's predictions for the policy change. We later implement the EA in a manner more like that used in actual EA simulations, by using the point semi-elasticity implied by the discrete choice model. The point semi-elasticity is calculated by taking the derivative of the probability with respect to log price for each individual in the sample and then taking the weighted average of those

derivatives: $\frac{1}{N_{tot}} \sum_i w_i [f_i(p_i)]^2 \beta \exp\{-\beta \log p_i - \gamma X_i\}$.

¹¹ The logistic model fitting procedure ensures that predictions of the number of insured for the entire sample will return the correct number. The same, however, will not be true for any sub-group. This discrepancy can be and often is overcome through model calibration. We repeated the analysis using a calibrated model and there were no qualitative differences in our results.

To implement the MA, we calculate the final number of insured within an income cell as: $\Psi_{cell}^f = \Psi_{cell}^0 + \tau(N_{cell} - \Psi_{cell}^0)$, applying the take-up rate τ to the uninsured group, as if they were eligible for a new program. We first use the take-up rate implied by the predictions of the discrete choice model for the policy change. We later implement the MA in a manner more like that used in actual MA simulations, dividing the sample into cells defined by the usual income categories and the price categories: \$0, \$1-\$499, \$500-\$999, \$1000-\$1499, \$1500-\$1999, \$2000-\$2499, \$2500-\$2999, \$3000-\$3499, and >\$3500. For example, one cell includes people with income 150-199% FPL who currently face prices of \$500-\$999 monthly. We calculate a cell-specific take-up rate by calculating the share of the cell that has insurance.¹² We then implement the program by reducing everyone’s price by 50% so that they move into new price-income cells. For example, to the people in the cell in the example above we now apply the take-up rate of those with income 150-199% FPL who face prices of \$1-\$499 monthly. We apply the appropriate take-up rate to all the uninsured *newly eligible* for the price. Those who are initially insured keep their prior insurance status. Those who do not move cells also keep their prior insurance status.

IV.B. Equivalence of Approaches

The change in the number of insured resulting from the 50% price reduction, as predicted by the discrete choice model, is 12.2 million (Table 1). The implied discrete-price-change semi-elasticity and take-up rates are used in the EA and MA, respectively, and yield essentially identical predictions (Table 1).

Table 1: Equivalence of Approaches and Point EA	
DCA: Predicted Change in Number Insured	12,222,288
Implied Discrete Change Price Semi-elasticity	.381
Implied Discrete Price Take-up Rate	.510
Implied Elasticity EA: Predicted Change in Number Insured	12,222,288
Implied Take-up Rate MA: Predicted Change in Number Insured	12,222,288
Point Semi-elasticity	.256

¹² This may seem like a contradiction with our definition of a take-up rate as the share of the uninsured who take-up a particular program. However, if one understands the existing data as being those otherwise uninsured people who respond to this price level by taking up insurance, the methods are consistent. In effect, this treats everyone in the cell facing a particular price as those eligible for a particular program (price level). The program we evaluate is then a new one, which changes the prices available to different individuals.

Point Elasticity EA: Predicted Change in Number Insured	88,869,780
Notes: Newly Eligible Population (those < 200% FPL without Medicaid):	69,183,929
Initial Number Insured in That Population:	41,413,953

This agreement is exactly that predicted by the theory section. Of course, this precise agreement requires the same data set and the same policy change. Moreover, the discrete-change-price elasticity is calculated from the very numbers of insured that resulted from the policy change. In our exercise, we create such numbers through the DCA, which serves as a benchmark for methodological comparisons. In effect, as our theory section showed, the results are the same, because the calculation is essentially the implementation of an identity.

IV.C. Discrete versus Point Semi-Elasticity

If instead, we make predictions using the point semi-elasticity-- a practice much more like that commonly employed by EA simulation-- we do not get the same results. As shown in Table 1, the predicted change in the number of insured is now 8.9 million. Both approaches use a semi-elasticity from the same model applied to the same data and apply it identically. The only difference between these methods is the size of the price change used to measure the semi-elasticity—the discrete vs. point semi-elasticity. Despite the tremendous similarities across methods, the results now differ by 27%.

The very different predictions that result from the point semi-elasticity vs. the discrete-price-change semi-elasticity belie a more general phenomenon. The elasticity varies throughout the population, as individuals vary in both observable and unobservable dimensions. The treatment of unobservable heterogeneity is embodied by functional form assumptions, which impose restrictions on the distribution of reservation prices, and, in turn, on the point elasticity variation with price.

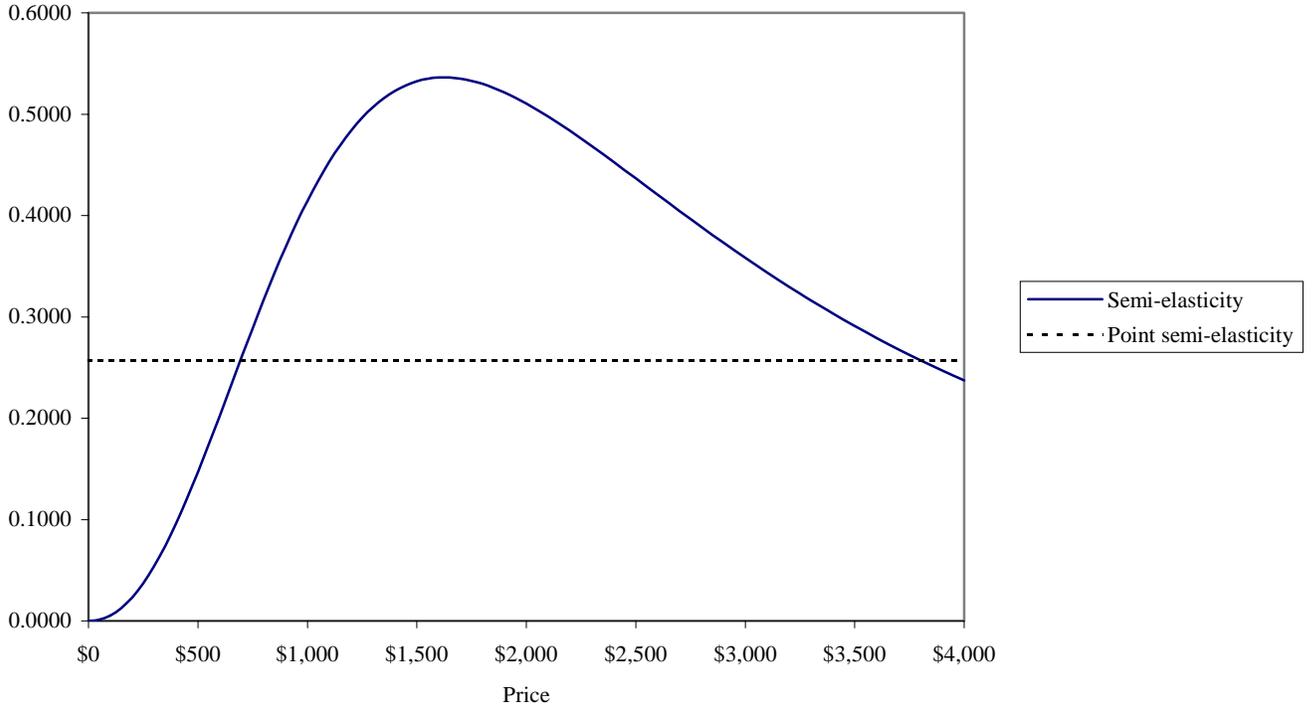
We illustrate the effects of different functional form assumptions using the logistic model. First, we calculate the point semi-elasticity, by calculating the effect of a marginal increase in price for each person in the sample, starting from that individual's current price. The weighted average of each of those effects is calculated to give a unique point semi-elasticity

based on current prices: $\frac{1}{N_{tot}} \sum_i w_i [f_i(p_i)]^2 \beta \exp\{-\beta \log p_i - \gamma X_i\}$.

Second, we calculate a continuum of point semi-elasticities at all possible prices. In this case, we calculate the effect of a marginal increase in price for each person in the sample, where each of those individuals is facing the same given price but otherwise maintains their original covariates. We then take the weighted average of the marginal effects over all individuals at that price, $\frac{1}{N_{tot}} \sum_i w_i [f_i(p)]^2 \beta \exp\{-\beta \log p - \gamma X_i\}$. We repeat this calculation throughout the distribution of prices.

Figure 5 presents graphically the results from these two calculations. It is the empirical counterpart to our earlier theoretical illustration in Figure 4. This illustrates the potentially enormous impact of unobservable heterogeneity. The constant semi-elasticity represents the simple assumption that semi-elasticity is invariant to starting prices. The variable semi-elasticity reflects the assumption that the coefficients of a logistic regression, including the coefficient of log price, are invariant to starting prices. Both approaches combine data with particular functional form assumptions, characterizing the unobservable heterogeneity. The correct characterization of unobservable heterogeneity is unknowable, but clearly different characterizations of that heterogeneity, through different functional form assumptions, will greatly influence predictions.

Figure 5: Variation of semi-elasticity with price. Dotted line is the constant semi-elasticity of the EA. Solid line is the semi-elasticity implicit in the DCA.



IV.D. Out-of-Sample Predictions

Table 2 shows the results from fitting the discrete choice model on the population $> 250\%$ FPL and using that model to predict the effects on those $< 200\%$ FPL. We do this for both the DCA and the EA. For the DCA we assume that the coefficients of the logistic model fitted on the high-income group apply to the low-income group. For the EA, we use the high-income logistic model to predict the effects of a 50% price reduction on the high-income group and use those predictions to calculate an implied discrete semi-elasticity. We then apply that semi-elasticity to the low-income group for the EA predictions. Clearly in this case the models do not yield identical results. The DCA predicts a change of 7.06 million while the EA predicts a change of 3.65 million. This difference arises despite the fact that the two methods differ only in the functional form with which parameters estimated are applied to predict changes for the lower income group.

It is interesting to note that the size of the predicted effect is dramatically smaller than that predicted when the model was fitted on the entire sample. This is not surprising because in our estimation high-income individuals are less responsive to price than are low-income individuals.

DCA: Predicted Change in Number Insured	7,060,728
Implied Discrete Change Price Semi-elasticity	.656
Implied Elasticity EA: Predicted Change in Number Insured	3,653,989
Notes: Newly Eligible Population (those < 200% FPL without Medicaid):	69,183,929
Initial Number Insured in That Population:	41,413,953

IV.E. Income-Price Cell Matrix Approach

The MA relies on tabulations of insured individuals disaggregated by a variety of covariates. An example of such a tabulation is shown in Table 3 with the initial number of insured, initial total population, and the take-up rates applied for each income-price cell. The pattern of variation by price is quite strange, reflecting the institutional realities of employer-sponsored health insurance, rather than any real behavior determined by price itself. A real-world matrix approach, of course, would make use of a variety of important other covariates and therefore would not imply such unrealistic behavior.

Table 3: INITIAL CONDITIONS

Price		Income relative to FPL				Total
		0%-50%	50%-100%	100%-150%	150%-200%	
\$0	<i>Take-up rate</i>	1.0000	1.0000	1.0000	1.0000	1.0000
	<i>Number insured</i>	3,559,070	947,331	1,569,684	1,928,792	8,004,877
	<i>Pop'n total</i>	3,559,070	947,331	1,569,684	1,928,792	8,004,877
\$0-\$499	<i>Take-up rate</i>	1.0000	1.0000	1.0000	1.0000	1.0000
	<i>Number insured</i>	475,008	1,841,223	4,072,107	5,805,723	12,194,061
	<i>Pop'n total</i>	475,008	1,841,223	4,072,107	5,805,723	12,194,061
\$500-\$999	<i>Take-up rate</i>	0.6302	0.6554	0.8027	0.8431	0.7532
	<i>Number insured</i>	1,436,140	1,690,627	2,905,743	3,309,970	9,342,480
	<i>Pop'n total</i>	2,278,830	2,579,480	3,620,046	3,925,951	12,404,307
\$1000-\$1499	<i>Take-up rate</i>	0.5253	0.6599	0.7595	0.8125	0.7078
	<i>Number insured</i>	427,336	618,707	1,026,751	958,625	3,031,419

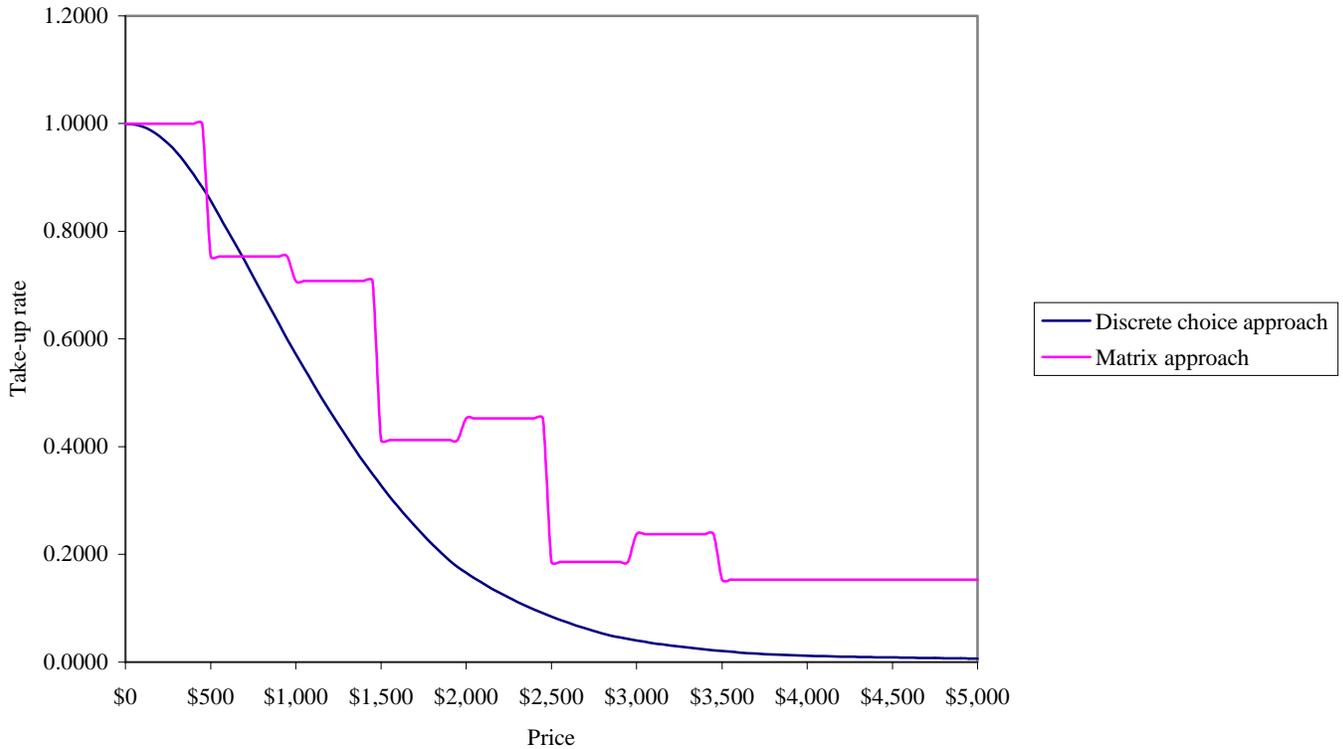
	<i>Pop'n total</i>	813,569	937,569	1,351,914	1,179,814	4,282,866
\$1500-\$1999	<i>Take-up rate</i>	0.4132	0.4069	0.3837	0.4512	0.4127
	<i>Number insured</i>	1,269,725	800,122	876,786	888,602	3,835,235
	<i>Pop'n total</i>	3,073,075	1,966,295	2,285,239	1,969,282	9,293,891
\$2000-\$2499	<i>Take-up rate</i>	0.5588	0.3552	0.3390	0.4066	0.4523
	<i>Number insured</i>	846,362	221,150	222,636	236,145	1,526,293
	<i>Pop'n total</i>	1,514,514	622,675	656,678	580,843	3,374,710
\$2500-\$2999	<i>Take-up rate</i>	0.1657	0.2074	0.2030	0.1765	0.1863
	<i>Number insured</i>	327,049	264,966	318,081	219,183	1,129,279
	<i>Pop'n total</i>	1,973,742	1,277,429	1,567,026	1,241,837	6,060,034
\$3000-\$3499	<i>Take-up rate</i>	0.2354	0.2172	0.2412	0.2586	0.2374
	<i>Number insured</i>	274,248	140,334	209,386	148,020	771,988
	<i>Pop'n total</i>	1,165,186	646,120	868,177	572,397	3,251,880
\$3500+	<i>Take-up rate</i>	0.1341	0.1479	0.1630	0.1974	0.1530
	<i>Number insured</i>	579,215	320,274	368,946	309,886	1,578,321
	<i>Pop'n total</i>	4,319,915	2,164,923	2,262,958	1,569,507	10,317,303
Total	<i>Take-up rate</i>	0.4795	0.5272	0.6338	0.7353	0.5986
	<i>Number insured</i>	9,194,153	6,844,734	11,570,120	13,804,946	41,413,953
	<i>Pop'n total</i>	19,172,909	12,983,045	18,253,829	18,774,146	69,183,929

We implement the MA for the proposed policy change by reducing by 50% the price of insurance for each person below 200% FPL. The proportion of the currently uninsured newly entering each cell that acquire insurance is equal to the old proportion of those in the cell who took-up insurance (the take-up rate marked). Table 4 shows the results of the MA with 16,806,756 million people predicted to be newly insured.

Table 4: Comparison of MA methods

DCA: Predicted Change in Number Insured	12,222,288
Implied Take-up Rate MA: Predicted Change in Number Insured	12,222,288
Price Cell Based MA: Predicted Change in Number Insured	16,806,756
Notes: Newly Eligible Population (those < 200% FPL without Medicaid):	69,183,929
Initial Number Insured in That Population:	41,413,953

Figure 6



The price-cell MA assumes that if anyone is given a particular price, his insurance take-up behavior will be the same as those who currently face that particular price. In contrast, the DCA predicts insurance take-up behavior using a model estimated from cross-sectional correlations in price and insurance status. Figure 6 illustrates the take-up rate as a function of price implied by both methods.¹³

In this section we have shown the empirical conditions under which the different approaches produce identical predictions. In addition, we have shown some commonly used practices that yield different predictions across the approaches-- using the point instead of the

¹³ The DCA curve is calculated as follows. To trace out the curve, we give each person in the sample the same new insurance price, starting with \$1 and repeating over the entire range of prices. The take-up rate is the share of the initially (predicted) uninsured that is predicted to become insured at the new price. However, if the new price is higher than the original price, we assume that they stay with the original price and that individual does not take-up at all.

discrete-price-change semi-elasticity, applying parameters estimated for one sub-group to another, and calculating take-up rates as existing average take-up rates. We have also demonstrated the potential of unobservable heterogeneity to render out-of-sample predictions enormously inaccurate.

V. Conclusions

Given the variety of people engaged in estimating the responses to policies to expand health insurance as well as the great variety of policy changes examined, there is a great onus on researchers to precisely and exhaustively explain all modeling assumptions (Glied, Remler and Graff Zivin). In order to allow a conversion of one simulation’s parameters into the terms of another simulation, it is critical to spell out exactly how an elasticity or take-up measure is actually employed.

Since all predictions inevitably involve out of sample predictions, it is important to assess how well parameters estimated in one situation generalize to the situation in which they will be employed. Obviously, there are limited sources for parameter estimation, particularly sources that do not suffer from selection bias, such as natural experiments. Therefore, there is an unavoidable tension between “clean” parameters and “relevant” parameters. While researchers may differ in how they handle this tension, they do policy makers a great service by clearly delineating the limitations of their choice.

People are different. It is not sensible to assume that individuals who heretofore have resisted acquiring insurance will respond to a subsidy in the same way as those who eagerly scooped up less generous subsidies. Hence, unobservable heterogeneity will always be important. All modeling approaches implicitly make assumptions about functional form that impose restrictions on unobservable heterogeneity. Such functional form restrictions can have a profound impact on the quantitative predictions of policy simulations. Researchers should explore the sensitivity of their predictions to specific functional form assumptions in order to disentangle the data-driven and model-driven elements of their predictions.

Formally, $\tau_{DCA} = \frac{\sum_i w_i (f_i(\min(p, p_i)) - f_i(p_i))}{\sum_i w_i (1 - f_i(p_i))}$.

While our focus has been on simulations to model health insurance expansions, many of the same concerns arise in simulations of other types of policy proposals (e.g., tax policy and non-health care benefits policies). The choice of functional form implies an assumption about unobservable heterogeneity that can have substantial effects on the estimates.

The fates of many policy proposals hinge on the actual numbers predicted for them: how many will participate and at what cost. It is often the case that evaluators, using different methods, will generate widely different numbers for the same proposal. Policy makers predisposed to reject the proposal will chose the less favorable numbers while those predisposed to support the proposal will chose the more favorable numbers. Without the ability to discuss the particular assumptions embedded in each estimate, there is little potential for constructive debate.

References

Baumgardner JR (1998): Providing health insurance to the short-term unemployed. *Inquiry* 35: 266-279.

Blumberg LJ, Nichols LM (2000). Decisions to buy private health insurance: Employers, employees, the self-employed, and non-working adults in the Urban Institute's Health Insurance Reform Simulation Model (HIRSM). Washington, DC: Department of Labor, Pension and Welfare Benefits Administration.

Feder J , Uccello C and O'Brien E (1999): The Difference Different Approaches Make: Comparing Proposals to Expand Health Insurance. The Henry J. Kaiser Family Foundation Project on Incremental Health Reform, October 1999.

Gruber J, Poterba J (1994): Tax incentives and the decision to purchase health insurance: Evidence from the self-employed. *Quarterly Journal of Economics* 109(3): 701-733.

Glied S, Remler D and Graff Zivin J (2001): "Inside the Sausage Factory: Understanding and Improving Estimates of the Effects of Health Insurance Expansion Proposals Using a Reference Case Approach" Columbia University Manuscript.

Gruber J (2000): Tax subsidies for health insurance: Evaluating the costs and benefits. (Prepared for the Kaiser Family Foundation)

Holahan J, Uccello C, Feder J (1999): Children's health insurance: The difference policy choices make. *Inquiry* 37(1): 7-22.

Pauly M, Herring B (2000a): Cutting taxes for insuring: Options and effects of tax credits for health insurance. Unpublished manuscript.

Pauly M, Herring B (2000b): Expanding insurance coverage through tax credits: Tradeoffs and outcomes. Unpublished manuscript.

Robert Wood Johnson Foundation (RWJF) (2000): Health coverage 2000: Cost and coverage analysis of eight proposals to expand health insurance coverage. (Prepared by The Lewin Group)

Zabinski D, Selden TM, Moeller JF, Banthin JS (1999): Medical savings accounts: Microsimulation results from a model with adverse selection. *Journal of Health Economics* 18:195-218.