An Agenda for Simulating Plan Offer Choice by Establishments and Offer Price Setting by Plans*

Martin R. Holmer†
HR&A Policy Simulation Group

April 25, 1996

This note outlines a proposed development agenda that would add and test two major capabilities to the state health insurance market simulation model. The model is being developed to provide policy analysis capabilities for states considering reforms that seek to make health insurance more affordable and accessible, especially to employees of small businesses. The current version of the model implements the general design of the market simulation model developed by Holmer, Long, and Marquis in an earlier project planning paper [2] as well as specific policy simulation capabilities described by the same authors in a recent policy analysis paper [3]. The model utilizes establishment data from the Robert Wood Johnson Foundation Employer Health Insurance Survey, which was conducted in ten states during late 1993 and early 1994.

The first section of the note describes the goals of the next stage of model development. The second section outlines the proposed approach to adding the capability for establishments to choose one or none out of several health insurance plan offers. The third section describes the proposed approach to adding the capability for plans to set the price of each health insurance offer based on the attributes of the establishment soliciting the offer, the nature of the rating regulatory regime, and the plan's simulated underwriting experience in the prior year. The ideas in all three sections have benefited from discussion with Roald Euller.

*This research is being supported by a technical assistance grant from the Robert Wood Johnson Foundation to The RAND Corporation and a RAND subcontract to HR&A. The grant is part of the Foundation's State Initiatives in Health Care Finance Reform program.

†Principal, HR&A Policy Simulation Group, 1156 15th Street, NW, Suite 1202, Washington, DC 20005 (voice: 202-986-0050, fax: 202-986-0052, email: holmer@digex.net).

Goals for this Stage of Model Development

The two capabilities that will be added in this stage will permit the model to simulate market dynamics under both experience rating and pure community rating regulatory regimes. Modified community rating regulatory regimes will be added in a subsequent stage of model development. One of the new model features will add the capability for establishments to choose one or none out of several health insurance plan offers. The approach to modeling offer choice will allow an establishment to consider a variable number of offers and each offer's attractiveness. The other new model feature will add the capability for plans to set the price of each offer based on prior experience. The approach to modeling price setting will permit a plan to consider its prior year's simulated underwriting profits and market share, the nature of the rating regulatory regime in place, and the attributes of the establishment soliciting the offer. The simulation of more sophisticated strategic and competitive considerations in offer price setting will be considered for addition in a subsequent stage of model development.

Behavioral Modeling Approach. Two principles guided the development of the approach proposed in this note.

The first principle is that the addition of new behavior to the model should be accomplished in a way that avoids the need to replace or change the model's existing behavioral capabilities. This is a reasonable goal given that the original model design effort tried to anticipate the interaction of a broad range of behavior by numerous market actors.

The second principle is that the new behavioral capabilities should be based on models of behavior that have been used in other areas of economic research. This will allow the characterization of behavior in terms that are familiar to other researchers and permit this project to draw on the insights and findings of prior work with these behavioral models.

Behavioral Parameter Calibration. Both of the new capabilities being added to the simulation model require behavioral models whose parameters are not easily estimated with currently available data. In lieu of empirically estimated behavioral parameters, a series of model sensitivity runs will be conducted to determine which combination of parameter values generates simulation results that are consistent with the stylized facts of a small group health insurance market under an experience rating regulatory regime. This calibration technique has been used in the natural sciences in situations where experimental data are either not possible to generate or not yet available. See Wilford [8] for a recent example in astronomy where sensitivity testing of a computer simulation model of the expansion of the universe has led to a major revision in beliefs about a key parameter value. This kind of sensitivity testing can be conducted over a subset of all possible parameter value combinations in cases where other research with the behavioral model suggests a likely range of values for some of the parameters.

The stylized facts of an experience-rated small group health insurance market include: (1) a substantial fraction of establishments do not have health insurance; (2) premiums for the same kind of coverage vary widely among establishments; (3) some persistence in establishment health insurance decisions (as revealed in retrospective questions and in the longitudinal subsample of the next wave of the establishment survey); and (4) adverse selection spirals (in which a plan's enrollees become increasingly high cost and its relative premium rises sharply over a number of years) happen occasionally. Also, Holmer, Long, and Marquis [3, page 3] refer to state reforms that "are also aimed at stabilizing premiums in the small group market and thereby reducing churning over time." It would be desirable to quantify the degree of premium volatility and churning so that this stylized fact can be added to the list.

The goal of the calibration work will be to determine which set of behavioral parameters produce simulation results that correspond to these styled facts and which do not.

Behavioral Parameter Estimation. If the calibration work indicates that realistic market simulation results are sensitive to the exact value of one or more behavioral parameters, the case for empirical estimation of these parameters will be quite strong. This empirical work may require the collection of data that can be used to estimate the behavioral model that includes the parameter. If this turned out to be the case, it would demonstrate the ability of simulation modeling to play a leading (rather than a lagging) role in a research program, providing an informal benefit-cost analysis of alternative data collection and behavioral estimation activities, rather than simply combining existing data and estimated behavioral models.

Modeling Plan Offer Choice by Establishments

It seems clear that the modeling approach in this area involves some kind of generalization of the simple discrete choice model now used in the simulation model. The model currently incorporates a binary probit offer model to estimate the probability that an establishment with a given set of characteristics offers its employees health insurance or offers them no health insurance. This binary probit offer model includes the cost of the offered plan, but not its coverage attributes.

The following discussion of plan choice by establishments begins with a list of features that an ideal approach would possess. The discussion continues with an outline of an interim approach that adds a first-phase select-an-offer model based on a conditional logit discrete choice framework to a second-stage offer-or-not model based on the current version of the binary probit offer model. The discussion concludes with an outline of a subsequent modeling approach that integrates the two phases of the interim approach into a single nested multinomial logit model of plan choice, which represents offering no health insurance as selecting the null plan that is available to all establishments.

Ideal Approach. The ideal approach to modeling plan choice by establishments would faithfully represent four key features of the health insurance choice situation.

First, the choice algorithm used in the simulation model should be able to handle situations in which the number of plan offers available to an establishment in the market varies. There may be just two plans assumed to be operating in an establishment's area in one model run, while five could be operating in the same area in another model run. In a multiple-state run, there is no reason to expect that all establishments will be facing the same number of plans making offers in the local market.

Second, the choice algorithm should be able to handle situations in which there are differences across plan offers not only in cost (*i.e.*, premium), but also in coverage (*i.e.*, cost sharing, provider choice, and quality of care). A typical market situation that will need to be simulated involves choice between HMO and FFS plans that are being offered with different premiums.

Third, the choice algorithm should be able to recognize that establishments with different work-force characteristics evaluate plan cost and coverage attributes in different ways, leading to differences in plan choice probability distributions across establishments facing the same choices.

And fourth, the choice algorithm should be able to recognize that the price sensitivity of selecting among health insurance plan offers may be greater than the price sensitivity of deciding whether or not to offer a plan at all. The results from the binary probit model of offering any kind of health insurance indicate a relatively low price elasticity, which is consistent with findings from a number of other studies. But this finding is not logically inconsistent with establishments being very price sensitive when selecting among alternative plan offers.

Interim Approach. Since there are no data currently available to estimate a complete model of offer choice, the interim approach calls for the calibration and sensitivity testing of a conditional logit model of offer selection. The conditional logit model, which was originally developed by McFadden [5], will be used to generate the establishment's offer selection under the assumption that it is going to be offering some kind of health insurance to its employees. The second-phase offering choice will be simulated using the binary probit model that has already been estimated with the survey data.

This approach views an establishment's choice process as sequential: the first phase involves selecting a plan offer from those available in the market and the second phase involves deciding whether or not to offer health insurance to employees or not. The establishment is viewed in the second phase as comparing the situation of offering employees no health insurance versus offering them the health insurance plan offer selected in the first phase of decision-making. Adopting this sequential choice process assumption is equivalent to assuming that the random factors that influence the two sequential decisions are independent (Maddala [4, pages 49–51]). The nested multinomial logit model proposed below as part of the subsequent approach can be used with appropriate data to estimate an integrated model of decision-making within which this independence assumption can be tested.

Under the assumption that an establishment offers its employees some kind of health insurance, the conditional logit model represents the random utility function as

$$v_{ij} = \beta x_{ij} + \epsilon_{ij}$$

where x_{ij} denotes a vector of attributes of offer j as perceived by establishment i. The functional form of the offer selection probabilities is derived from the assumptions of random utility maximization and of the unobserved additive utility factors ϵ_{ij} being independently and identically distributed with a type I extreme value distribution. Under these assumptions the probability of establishment i selecting offer j, which is denoted by p_{ij}, takes the form

$$p_{ij} = e^{\beta x_{ij}} / \sum_{k=1}^{n_i} e^{\beta x_{ik}}$$

where n_i denotes the number of offers available to establishment i. The β vector of utility coefficients can be interpreted as representing the hedonic prices of offer attributes. The higher the β value for an attribute, the more sensitive is an establishment's decision to differences across the offers in the value of that attribute. See Maddala [4, Chapter 3] for an introductory discussion of this and other probabilistic choice models in econometrics.

A simple specification of the offer attribute vector would call for just two attributes: coverage (an overall measure of the cost sharing, provider choice, and quality of care features of the offered plan as perceived by the establishment) and cost (the plan premium offered to the establishment). Plan coverage values would be normalized so that the one plan implicitly assumed in the current binary probit offer model would have a coverage value of exactly one. A plan that is perceived to offer less coverage would have a coverage value less than one. The cost of a plan offer would be measured by the premium for a self-only plan just as is now done in the binary probit offer model.

Given this simple two-attribute specification of the conditional logit selection model, the calibration task is not too complex. The utility coefficients of coverage and cost can be specified with a range of alternative values in the sensitivity testing. Appropriate values of the cost coefficient should equal or exceed in absolute value the premium coefficient in the binary probit offer model after the standard probit-to-logit coefficient translation is done. The selected offer's values of coverage and cost would be used to construct a coverage-adjusted cost of the conditionally selected plan offer for use in the second-phase binary probit offer model. Using this coverage-adjusted premium, the probit model would be used to determine if the establishment decides to offer any kind of health insurance to its employees. If it does, the offer selected in the first decision phase would be offered to employees.

In both this first and second phase of the sequential decision-making, the probabilistic choice models are estimated or calibrated with cross-sectional data for a single year. When using these models in a dynamic simulation, consideration must be given to the possibility of serial correlation across time in the unobserved random factors ϵ_{ij} in the utility function for a given establishment. It seems plausible to assume that there is some persistence to the unobserved factors that influence an establishment's health insurance decisions from year to year. The proposed method for handling this in the interim approach is to induce serial correlation in the uniformly distribution random numbers generated to realize an establishment's fitted selection or offer probability distribution. Of course, a serial correlation coefficient of zero would be included in the range of values used in the sensitivity testing.

If u_t denotes the uncorrelated uniform variate generated for year t, then the correlated uniform variate w_t for year $t = 0$ is

$$w_0 = u_0$$

and for years $t = 1, \ldots, T$ is

$$w_t = \rho w_{t-1} + (1 - \rho) u_t$$

where ρ has a value between zero (no serial correlation) and one. It seems appropriate to experiment with serial correlation in both the first-phase conditional logit selection model and the second-phase binary probit offer model. In addition to persistence in the unobserved characteristics of the establishments and plans, serial correlation would be induced by the existence of various types of financial and employee-morale costs associated with changing the selected plan as well as dropping or adding health insurance coverage.

This interim approach is consistent with the characteristics of the ideal approach outlined above, and moves towards the kinds of statistical models that could be used to estimate the plan offer choice behavior of establishments if more complete data on offers were to become available. The main virtues of the interim approach are that it builds on current capabilities in the market simulation model and that it provides the additional capabilities required to begin modeling health insurance market dynamics with competing plans setting offer prices and establishments responding to the plan offers.

Subsequent Approach. After a thorough review of the considerable number of methods and examples presented in the book by Maddala [4], it appears as if the most promising approach to an integrated model of plan offer choice is the nested multinomial logit model in which the option of offering employees no health insurance is considered offering the null plan, which has no coverage and a zero premium. The nested model is commonly used to handle situations in which the utility function coefficients differ between subgroups of choice alternatives, which seems likely to be a feature of the plan choice situation. The assumed nesting can be characterized as a tree decision structure in which there is a binary choice between the null plan and the rest of the plans and a selection of one of the non-null plan offers. This second "limb" could even be divided into three "branches" representing FFS, managed-care, and HMO plans.

The nested multinomial logit model for plan choice is the generalization of the conditional logit model of random utility maximization in which the additive random utility factors are distributed with a generalized extreme value distribution (see Maddala [4, pages 67–73]). This is the same generalization that is involved in moving from a binary probit model to a multinomial probit model, but the resulting logit model is much easier to estimate. The estimation process is sequential with the use of inclusion values, which are constructed from the results of the first stage of estimation, in the second stage of estimation (see Maddala [4, page 73]).

The nested multinomial logit model of plan choice is similar in spirit to the two-phase model described in the interim approach section above. The main virtue of the nested model is that it can actually be used to estimate the parameters of the plan choice model (instead of just calibrating the parameters of the offer selection decision branch as in the proposed interim approach). But, of course, this estimation would require data on the offers that a number of establishments rejected as well as the offers they accepted.

One of the main objectives of the interim approach is to determine if the character of the overall market simulation results depend critically on the values of the parameters of the offer selection model. If they do, then a case has been made for the policy research benefits of gathering the data necessary to estimate a nested multinomial logit model of plan choice. The cost of that data collection activity could be reduced by drawing a relatively small sample of establishments for which complete plan offer sets are constructed.

In addition to performing this informal benefit-cost study on additional data collection, it would be prudent to explore the econometric issues involved in estimating this kind of nested multinomial logit model by conducting a trial estimation study that uses one or more hypothetical data sets. This exploration would be facilitated by use of the nested multinomial logit model TSP package (developed in December 1994 by Paul Ruud of UC-Berkeley) that is available free to users of TSP. HR&A has a registered version of TSP as well as a copy of Ruud's package.

Modeling Offer Price Setting by Plans

The previous section outlined an approach to providing establishments in the health insurance market simulation model with the capability of choosing among several plans whose offered premiums will be changing from year to year. This section switches the focus from the demand side of the market to the supply side. Given that each simulation model run assumes that a constant number of plans operate every year and that each plan offers the same coverage every year, the sole supply-side behavior of plans is setting the price quoted to each establishment that solicits an offer. Currently, the simulation model represents the offer price as a vector of premiums, currently one for self-only coverage and one for family coverage.

It seems clear that the modeling approach in this area involves implementing a behavioral algorithm that adds the capability for plans to set the price of each health insurance offer based on the attributes of the establishment soliciting the offer, the nature of the rating regulatory regime, and the plan's simulated underwriting profits and market share in the prior year. The proposed approach to adding this modeling capability combines this expanded plan behavior with insights drawn from recent research on market trading institutions as described by Friedman and Rust [1].

Market Institution. The findings from this market trading literature are remarkably consistent whether they are from experimental studies with human subjects acting as market traders or simulation studies with different kinds of computer programs acting as traders. See, for example, the review of prior experimental findings and the new simulation results reported by Rust, Miller, and Palmer [7].

In addition to being influenced by the initial endowments of the market traders, a market's allocation is determined by the rules that determine how allowable messages from the traders are translated into transactions. Most market trading research has focused on the double auction market institution in which supply-side market participants make offers and demand-side participants make bids. In this kind of market institution, which is also called an open-outcry market and is typified by the commodity trading pit at the Chicago Board of Trade, transactions can occur only between the seller with the lowest ask and the buyer with the highest bid.

Most of the experiments and simulations that have been conducted using this double auction market institution are structured so that sellers have a known and fixed unit costs against which to compare their selling revenue and buyers have known and fixed unit values with which to compare their buying costs. In this kind of situation, the completely uncoordinated behavior of the traders making bid and ask offers leads quickly to a competitive equilibrium allocation. This convergence to the competitive equilibrium is accomplished without the "Walrasian auctioneer" or *tâtonnement* process so common in traditional economic theory. And furthermore, quick convergence to a competitive equilibrium is accomplished when there are as few as four traders on each side of the market. This result is in stark contrast to the conventional theoretical notion that many traders are required on both sides of a market for a competitive equilibrium allocation to be generated.

While these research results will provide some guidance, they are not likely to be replicated for the health insurance market. The market for health insurance differs in at least two important ways.

First, health insurance is not bought and sold in a double auction market. The simulation model is structured so that each establishment solicits an offer from every plan operating in its geographical area and then decides which one of the offers to select. Using this provisionally selected plan offer, the establishment decides whether or not to offer health insurance to its employees (*i.e.*, chooses between the provisional plan offer and the null plan). This is similar to the posted-offer market institution that is used to sell Treasury bonds. It will be important to determine whether these differences in the market institutions of the health insurance market and the double auction market lead to substantially different kinds of market allocations. It should be possible to specify a model simulation run that is simple enough to calculated its competitive equilibrium and then compare the simulated market results to the calculated competitive equilibrium. This will identify how much of the difference in overall market simulation results is related to differences in market institutions and how much is related to the second major difference in the markets. (The assumption that establishments solicit offers from all available plans may be somewhat unrealistic except for the fact that insurance brokers, who play a major role in the small group market, provide establishments with information about many plans and their likely offers. It is too early to know whether this assumption needs to be relaxed in order for the simulation model to generate the stylized facts of the small group market.)

The second important difference between the double auction and health insurance markets is that the cost of providing health insurance to an establishment is not known by the plans. The cost of providing health insurance depends on the uncertain volume and size of claims filed by the establishment's covered employees. The fact that establishments may know more about likely claims costs than do plans has been shown by Rothschild and Stiglitz [6] to create a situation in which no competitive equilibrium exists. It will be important to use the health insurance market simulation model to determine under what behavioral parameter combinations this adverse selection occurs and what sort of market dynamics are generated when a static competitive equilibrium does not exist.

More details on the research agenda will be provided after plan price setting behavior is discussed in the next section.

Plan Behavior. In addition to the influence of market institutions and initial endowments, the kind of pricing behavior exhibited by plans would seem to be the major determinant of the nature of the market allocation. The most surprising finding from the double auction market research is that this is not the case. In both laboratory experiments and computer simulations, the market allocation has moved relatively quickly to the competitive equilibrium regardless of the bid and ask strategies of the traders. In fact, Rust, Miller, and Palmer [7] find that the most successful computerized trader in their market simulation tournament uses a very simple decision-making algorithm. This simple rule-of-thumb algorithm consistently out performs other algorithms that use sophisticated optimization and learning techniques.

One implication of these findings is that this line of health insurance market research should not hesitate to start with a relatively simple behavioral algorithm for plan price setting behavior. Before outlining the proposed algorithm, it will be necessary to discuss how a plan will know the values for its simulated underwriting profits and market share.

Plan revenues and market share are relatively easy to add to the current version of the health insurance market simulation model. Once an establishment chooses a plan and determines the number of plan enrollees among its employees, the plan's offered price can be used to calculate total premiums for both self-only and family coverage. The current version of the model is structured so that establishments submit premiums to plans after plan selection and plans keep a tally of aggregate premiums and enrollees. It should, therefore, be easy for each plan to determine its underwriting revenue each year. It seems reasonable to assume that each plan is aware of the market share of all the other plans operating in its area. This means that each plan can calculate its relative market share.

Generating a plan's underwriting costs during a year is more difficult. In the real world, the annual utilization experience of plan enrollees generates plan claims. The cost sharing features of the plan interact with utilization to determine the enrollee's out-of-pocket cost and the plan's reimbursable claims and administrative costs. The current version of the health insurance market simulation model has not developed the behavioral details of its family and individual modules even though the model actually processes both a sample of establishments and a mock sample of families.

In lieu of a developed family module that would generate utilization and claims, it seems as if the premium-quote equation, which is estimated with data from an experience rating regulatory regime, could be pressed into service as a loaded-claims equation. The idea is to use this equation for two different purposes. First, as currently done in the model, it will be used to generate premium quotes under experience rating for establishments with no underwriting experience with the plan. In this use, the equation can be thought of as a plan's actuarial rate book. Second, the equation will be used to generate aggregate loaded claims for the plan enrollees of an establishment. The claims generate by the equation will be fully loaded for administrative costs as well as normal profits given the nature of the data used to estimate the equation. This approach allows the simulation model to be developed without undertaking the substantial effort required to prepare family survey data, refine the algorithm that matches employed individuals with establishment employee groups, and develop behavioral algorithms for the family module.

It will be important to consider the possibility of serial correlation in the normal variates used to generate the loaded-claims amounts. If there are unobserved attributes of an establishment's enrollee group that are somewhat persistent, then some degree of positive correlation in the random terms would be appropriate.

Given that the model is enhanced along these lines, at the end of each simulation year a plan will know its underwriting revenues, costs, and market share. Given that information, how will the plan change its pricing strategy for the next year? The proposed approach here is to specify a reaction function that produces a proportional change in the general level of the plan's price quotes. This proportional change and the current level would be used to produce the plan's relative pricing index for the subsequent year. This pricing index, which is plan specific and has a value of one in the initial year of the simulation, is used as a multiplicative adjustment factor by the plan to generate premium quotes for individual offers. Under an experience rating regulatory regime, the pricing index would be multiplied by the fitted values generated from the premium-quote equation for establishments the plan did not cover in the prior year. Premiums would be adjusted to prior year's loaded claims for establishments the plan did cover. Under a pure community rating regulatory regime, the pricing index would be multiplied by the plan's initial year community-rating premiums.

The sensitivity testing of the enhanced model is likely to generate new ideas about how best to specify this reaction function as well as provide an understanding of how different numerical values of the function's parameters influence the market simulation results. But to get the model development and testing started, an initial reaction function is needed. Let P_t denote the plan's pricing index in year t. For year $t = 0$ the pricing index is by definition

$$P_0 = 1$$

and for year $t = 1, \ldots, T$ the value of the pricing index is

$$P_t = R_{t-1} P_{t-1}$$

where R_t denotes the value of the plan's reaction function in year t. Let E_t denote the plan's normalized excess profit in year t (*i.e.*, underwriting revenues minus underwriting costs all divided by underwriting costs) and S_t denote the plan's market share measured relative to the average market share of all plans operating in its area. Since both P_t and R_t must be positive, it is convenient to specify the reaction function as returning a value of $log(R_t)$ given the values of E_t and S_t.

For a reaction function to be plausible it should have the following characteristics. First, $log(R_t)$ should be positive (causing an increase in the pricing index) when normalized excess profits are negative and relative market share is greater than one. This represents a situation in which the plan offered insurance at too low a price leading to a large enrollment and underwriting losses. Second, $log(R_t)$ should be zero (causing no change in the pricing index) when normalized excess profits are positive and relative market share is greater than one. This represents the desirable situation in which the plan earned excess profits on a relatively large enrollment. And finally, $log(R_t)$ should be negative (causing an decrease in the pricing index) or zero (causing no change in the pricing index) when relative market share is less than one. This represents a situation in which the plan offered insurance at too high a price leading to a small enrollment.

One specification of the reaction function that has these three characteristics is

$$log(R_t) = \begin{cases} max[0, \alpha E_t(S_t - 1)] & \text{for } S_t > 1 \\ min[0, \beta E_t(S_t - 1)] & \text{for } S_t \leq 1 \end{cases}$$

where the behavioral parameters α and β are both negative. The reaction function has been specified in terms of profits and market share rather than in term of quoted premiums relative to other plan's premiums because the heterogeneity of plan coverage makes direct price comparisons meaningless. In addition, under an experience rating regulatory regime there are many prices for the same plan depending on which establishment solicited the offer.

Research Agenda. After the model has been enhanced with these plan price setting capabilities, there will be a need to calibrate the behavioral parameters since there are no data available to estimate a plan pricing reaction function. As discussed above, the objective will be to conduct a broad range of sensitivity tests to determine which set of parameter values produces market simulation results that most closely resemble the stylized facts of the small group market operating under an experience rating regulatory regime.

Also, of considerable interest will be the market simulation results generated under the assumption that all establishments generate the same non-stochastic claims each year. In a market with several plans offering identical coverage, this constant claims assumption eliminates the possibility of adverse selection and permits study of whether the market institutions assumed for the health insurance market are sufficient to generate a competitive equilibrium. Without this assumption, the model market simulation model should be able to be used to study the nature of adverse-selection dynamics.

Once, these tests have been conducted and behavioral parameters have been calibration, attention will shift to generating market simulation results under a pure community rating regulatory regime. This market setting will provide an opportunity to explore the effectiveness of alternative risk-adjustment policies.

References

[1] Friedman, Daniel, and John Rust (editors), *The Double Auction Market: Institutions, Theories, and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume XIV, Addison-Wesley, 1993.

[2] Holmer, Martin R., Stephen H. Long, and M. Susan Marquis, "Work Plan for Analyzing the Implications of Health Insurance Alliance Design Choices," working paper, March 15, 1994.

[3] Holmer, Martin R., Stephen H. Long, and M. Susan Marquis, "The Effects of Small Group Reform on Employers' Decisions to Offer Insurance: Some Preliminary Results," paper presented at the American Economics Association meetings, Washington, DC, January 7, 1995.

[4] Maddala, G. S., *Limited-Dependent and Qualitative Variables in Econometrics*, Cambridge University Press, 1983.

[5] McFadden, Daniel, "Conditional Logit Analysis of Qualitative Choice Behavior," in P. Zarembka (editor), *Frontiers in Econometrics*, Academic Press, 1973.

[6] Rothschild, Michael, and Joseph Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, November 1976.

[7] Rust, John, John Miller, and Richard Palmer, "Behavior of Trading Automata in a Computerized Double Auction Market," in Daniel Friedman and John Rust (editors), *The Double Auction Market: Institutions, Theories, and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume XIV, Addison-Wesley, 1993.

[8] Wilford, John Noble, "Computer Enlists Einstein's Enigmatic Force and Little Mass to Explain the Universe," *New York Times*, April 9, 1996, pages C1 and C7.