Modeling Health Insurance Choices in "Competitive" Markets

Keane, Michael

2004

Modeling Health Insurance Choices in "Competitive" Markets

Michael P. Keane
Department of Economics
Yale University

First Draft, June 2004
Revised, September 2004

Prepared for presentation as a distinguished lecture in health economics at the Center for Health Economics Research and Evaluation (CHERE), at the University of Technology, Sydney, Australia – June 24, 2004. I wish to thank (without intending to implicate) Randall Ellis, Hanming Fang, Jane Hall, Katherine Harris, Laurel Hixon and Ahmed Khwaja for many helpful comments.

I. Introduction

All OECD countries have some form of government provided health insurance, although countries differ substantially in terms of whether this coverage is universal, or limited to certain groups, and whether or not government institutions combine the health care delivery and insurance functions. In many countries, including the U.S. and Australia, private firms can offer alternatives and/or supplements to government provided health insurance, and recently there has been considerable interest in whether such private competition is beneficial for consumers.

If several health insurance plan options are available in the market, we can use the tools of choice modeling to analyze how consumers choose among the competing plans, and how they value plan attributes. The consumer welfare implications of introducing new types of insurance plans (i.e., plans with different attributes) can also be analyzed using these methods. Thus, choice modeling can help us understand the extent to which giving consumers more health insurance options might increase consumer welfare.

This paper describes how state-of-the art methods of choice modeling can be used to analyze consumer choice behavior in "competitive" health insurance markets, using, as an example, the insurance choices of senior citizens in one region of the U.S.. I also discuss the more general issue of whether choice is beneficial for consumers in the health insurance context. In fact, consumers do not necessarily benefit if we allow more "choice" by letting private firms offer health insurance in competition with government. Rather, consumers may be better served if a single payer offers a menu of insurance options. If government were to use market research tools to design that menu, consumers would still have scope for welfare enhancing choices. Even if a single payer solution is politically infeasible, government clearly should use choice modeling techniques to help design the insurance option(s) that it offers in competitive markets.

To begin, it is useful to carefully define a "competitive health insurance market." I'll follow Van de Ven and Ellis (2000), who state: "By competitive, we mean markets in which individual consumers have a periodic choice of health plan and health plans may take actions, such as designing, pricing and marketing their products, to attract or repel enrollees."

By this definition, a "competitive" market could take the form of:

(i) a single payer who provides a menu of choices

(ii) a market with one or more private insurers and no government involvement, or

(iii) some hybrid of the two, in which private firms compete with government provided insurance.

For example, in the U.S., senior citizens have a choice between:

(i) the government provided Medicare fee-for-service plan ("Basic Medicare"),

(ii) Private HMOs that are subsidized by the government ("Medicare HMOs"), or

(iii) Basic Medicare plus supplemental <u>private</u> insurance to cover services or costs not covered by Medicare ("Medigap insurance").

Notably, the definition of "competitive" in Van de Ven and Ellis does not require private firms. A "competitive" market could involve a single payer (i.e., government) providing a menu of health insurance options. But, when policy makers discuss "competition" in health insurance, they typically mean letting private firms offer alternatives to government provided insurance.

The notion that allowing private firms to compete with government in the provision of health insurance is a good idea seems to rest on two assumptions:

(i) Choice is good. Government provided insurance is "one size fits all," while private firms can provide plans better tailored to individual preferences.[1]

[1] As Cutler and Zeckhauser (2000) note: "Health insurance choice is a natural way to meet differing individual preferences. Some people will prefer managed care insurance, which limits utilization but costs less, while others will opt for a more open ended indemnity-style policy."

(ii) Competition among alternative plans will promote market "efficiency," because plans will have to keep expenses down to survive in a competitive market place

But, the problem with the argument for "competition" in health insurance markets is that neither of the two assumptions on which it rests is obviously valid. The reasons are as follows:

1) Market research tools for designing products that appeal to consumer tastes are well known. Thus, while it is not the tendency of government to be responsive, there is no necessary reason that government could not use market research to design a menu of options that appeal to heterogeneous consumer tastes. Consumer choice is possible within a single payer system.

2) Private insurers have incentives to "cherry pick," which means trying to attract people who are good risks (i.e., people who will be profitable because they are unlikely to need services). In general, this raises average costs, and hence premiums, among those who stay with government insurance. The notion that "more choice is good" rests on the assumption that attributes of existing options stay fixed when new options are added. Then, adding a choice can't hurt anyone, and can help some people. But, if attributes of existing choices change when new choices are added, consumers can be made worse off. Thus, letting private firms offer insurance plans in competition with government will not necessarily benefit consumers.

3) Competition will not reduce costs if private firms seek profits primarily through cherry picking activity (e.g., marketing, advertising) rather than through more efficient service delivery.

In most of this paper I will focus on point (1). That is, I will explain how state-of-the-art methods of market research, also known as choice modeling, can be used to:

(i) Analyze consumer preferences for attributes of health insurance plans,

(ii) Predict demand for new health insurance products (with particular attributes),

(iii) Predict consumer welfare effects of adding new insurance products.

Since the knowledge of how to use these methods is not limited to private firms, I would argue that government should use such techniques to help design its insurance plan offerings, whether that be in the context of a designing a menu of health insurance options within a single payer system, or, alternatively, to design its own health plan offering(s) within a competitive system.[2]

After describing how choice modeling techniques work, I explain why a single payer system, where government applies these methods to generate a menu of options for consumers, may be preferable to a system where private insurers compete with government. The basic argument is as follows: Given concerns about equity and market failure, we can safely assume that there will continue to be substantial government intervention in the health insurance market. This will take the form of government provided insurance along with, perhaps, government subsidies to private insurance providers. As long as government provides insurance and/or subsidies, any private firms operating in the market will have an incentive to engage in cherry picking behavior. A single payer providing a menu of insurance options avoids the cherry-picking problem, while still providing a measure of consumer choice.

Furthermore, the computational problem that a government has to solve in order to design an appealing menu of insurance options is greatly complicated by the presence of private firms in the market. This is because, in order to predict the costs and welfare implications of any menu of

[2] Interestingly, President Clinton's Health Security plan did involve offering consumers a menu of insurance options. The plan required the U.S. States to create health care "alliances." These alliances would pool together employees of small to medium firms, government employees, the unemployed and self-employed, and negotiate (on their behalf) a menu of insurance plan options with private insurers. The members of the alliance could then choose from this menu in an annual open enrollment. (Large employers could continue to negotiate with insurance companies on their own). For employed members, the employer would pay 80% of the premium for the selected plan, while the employee would pay 20%. Unemployed and self-employed members would still pay 20%, with the remainder financed by the State. Thus, the idea of the plan was that it would make health insurance affordable for the unemployed and self employed, or for those whose employers did not offer insurance, by giving them the same 80% premium subsidy that is typical of large firms, and by making them part of a large alliance that could negotiate favorable rates (again, like employees of large firms). The law did specify some aspects of the menu of options to be offered by alliances. For instance, it had to include a traditional fee for service plan that did not restrict provider choice, along with other plans with certain specified features. But Clinton's health care task force apparently did not take seriously the possibility of using consumer research methods to help design the menu of options.

insurance options it might offer, the government has to predict the responses of the private firms, including which plans they would offer and their cherry picking behavior – an intractable problem. Given that government is going to intervene in the insurance market anyway, the computational problem that it must solve to do so in a sensible and efficient manner is greatly simplified by ruling out private competition.

II. An Application of Choice Modeling to the Health Insurance Market

II. A. The Data

To illustrate the application of market research techniques to the health insurance market, I'll describe some work I did with Katherine Harris a few years ago. In Harris and Keane (1999), we modeled how senior citizens living in a particular region of the U.S. choose among insurance options. In my (admittedly self-serving) view, this is the most sophisticated application of choice modeling to the health insurance market to date, both because of the rich pattern of consumer taste heterogeneity that is accommodated in the modeling, and because of the new methods we developed to estimate "unmeasured" attributes of alternatives. The data that we used come from the "Twin Cities" of Minneapolis and St. Paul, Minnesota, and were collected by HCFA in 1988. The sample size was N = 1274, and the mean age of the sample members was 74.

In order to understand the choice problem faced by consumers in these data, it is important to understand two things about this market. First, the basic Medicare "fee-for-service" program, which provides insurance coverage to those 65 and over, requires significant cost sharing (especially for hospital stays) and leaves a number of services, such as preventive care and, until recently, prescription drugs, uncovered.[3] Thus, many senior citizens buy supplemental insurance, known as "medigap" plans. These plans may cover Medicare deductibles and co-pays,

[3] The Medicare Modernization Act of 2004 introduced rather limited drug coverage. The new benefit does not take effect until 2006, and there are substantial cost sharing requirements.

as well as additional services and/or prescriptions. There were many such plans offered by private insurance companies in the Twin Cities in 1988, but we found they could be fairly accurately categorized into those that provided drug coverage and those that did not, with other plan features (like premiums) fairly comparable within each of those types.

Second, two basic types of "managed care" options were available in the data. Both are plans offered by private firms, known as health maintenance organizations (HMOs). These "Medicare HMOs" receive a per enrollee government subsidy (i.e., a "capitation" payment) that is somewhat less than the government's cost of insuring a typical Medicare enrollee. The basic idea is that, if the HMO can serve the person for less than the subsidy amount, then it makes a profit and the government saves money. It would seem that everyone is better off, but the situation is complicated if the HMO saves on costs not only through enhanced efficiency in service provision but also via cherry picking (i.e., attracting relatively healthy, low cost senior citizens to enroll). I'll return to that issue latter, but for now it is only necessary to understand that there are two basic types of HMOs. The first is called an independent practice association (IPA), while the second is called a group or network HMO.

In an IPA, the private insurance company contracts with a set of health care providers, and plan members can choose to obtain services from any of them. The idea here is that the IPA can obtain cost savings by negotiating favorable reimbursement rates with the providers who join. Ideally then, these providers have to contain costs in order to still make profits from serving the IPA patients, so the efficiency of health care provision is enhanced. In a group HMO, the private insurance company actually employs a staff of providers, thus combining the health care delivery and insurance functions. Then, it can attempt to enhance efficiency of service provision internally, via the incentives it creates for the employed doctors.

Thus, the consumer choice set contains five insurance options:

1) Basic Medicare
2) Medicare + a "medigap" insurance plan without drug coverage
3) Medicare + a "medigap" insurance plan with drug coverage
4) An HMO of the independent practice association (IPA) type
5) A Network or Group HMO

The key attributes of plans that we observe in the data are described in Table 1. These are: the premium, whether the plan covers drugs, covers preventive care, and allows provider choice, and whether an enrollee must submit claims for reimbursement after using medical services.

Table 1: Health Plan Attributes (Twin Cities 1988 Market)

	Basic Medicare	Medicare +medigap w/o drugs	Medicare +medigap w/drugs	IPA	HMO
Monthly premium	$28	$71 to $82 (based on age)	$95 to $109 (based on age)	$53	$40
Drug Coverage			Yes		Yes
Preventive Care				Yes	Yes
Provider Choice	Yes	Yes	Yes	Yes	
Must Submit Claims	Yes	Yes	Yes		

Crucially, two important attributes of health insurance plans are not measured in the data: quality of care and cost sharing requirements. This isn't a specific failure of these data, because these attributes are intrinsically difficult to measure. First, there is a large literature on quality measures in health care, and it doesn't come to a clear consensus on how such measurement should be done. Second, cost sharing rules of insurance plans are quite complex. There tend to be many different cost-sharing requirements for different types of services under different circumstances. Thus, it is very difficult to come up with any overall measure of "cost sharing."

The lack of quality and cost-sharing measures is an important problem for two reasons. First, a choice model that ignores these two attributes may give very misleading estimates of how consumers value the other attributes. Second, these two attributes are a critical aspect of any insurance plan, so, unless we know how consumers value them, we can't measure the welfare implications of adding new plans. However, a key aspect of the Twin Cities Medicare data is that it contained attitudinal data in which consumers were asked how much they valued various attributes of a health insurance plan. A key contribution of Harris and Keane (1999) was to show how this type of attitudinal data could be combined with consumers observed health plan choices to measure both: 1) how consumers value the unobserved attributes, and 2) the levels of the unobserved attributes possessed by each plan in the market (as perceived by consumers).

The attitudinal data were obtained from questions in which respondents were asked whether, in order to consider an insurance plan, it would "have to have" a certain attribute, or whether they would just "like to have" the attribute, or whether the attribute "doesn't matter" in deciding if a plan is considered. The questions and response frequencies are described in Table 2.

Economists typically eschew these type of data as measures of preferences, because they tell us nothing about a consumer's willingness to pay for various attributes. That is, there would appear to be no way to convert consumer responses to such questions into monetary measures of how consumers value attributes. However, in the approach developed by Harris and Keane (1999), the responses to such attitudinal questions are treated as "noisy" indicators of consumer preferences when estimating a model of consumer choice behavior. This enables one to construct estimates of consumer willingness to pay for the unobserved attributes (while also allowing one to construct more precise estimates of consumer willingness to pay for observed attributes). To describe how this approach works, it is necessary to lay out the choice model in some detail.

Table 2: Stated Attribute Importance Measures

("Tell me if you would …. to consider a plan")

	"Have to Have"	"Like to Have"	"Doesn't Matter"
Observed Attributes:			
Lowest Premium	23%	59%	18%
Drug Coverage	22%	60%	18%
Preventive Care	32%	55%	13%
Provider Choice:			
Choice of Physician	35%	55%	10%
Choice of Hospital	26%	60%	14%
Low Paperwork	38%	53%	9%
Unobserved Attributes:			
Low Cost Sharing	31%	60%	9%
Quality:			
Highest Quality	44%	52%	4%
Referral to Specialists	41%	54%	5%
Not Rushed from Hospital	33%	56%	11%

Notes: Each attitude scale was coded: 1="Doesn't Matter," 2="Like to have," 3="Have to Have."

The importance of quality measure was created by summing the three quality related questions and dividing by 3. The importance of provider choice measure was created by summing the two provider choice questions and dividing by 2.

II. B. The Choice Model

The insurance choice model in Harris and Keane (1999) is laid out as follows: Let X_j denote the vector of the observed attributes of insurance option j, where $j = 1,\ldots,5$ indexes the five options listed in Table 1. X_j includes:

(i) Premium (in $ per month)
(ii) Drug coverage (a 0/1 indicator)
(iii) Preventive Care (a 0/1 indicator)
(iv) Provider Choice (a 0/1 indicator)
(v) Must Submit Claims (a 0/1 indicator)

Next, let A_j denote the vector of un-observed attributes of insurance option j. A_j includes:

(i) Cost Sharing
(ii) Quality

Then, letting U_{ij} denote expected utility to person i if he/she chooses insurance option j, we have:

(1) $U_{ij} = X_j\beta_i + A_j W_i + \varepsilon_{ij}$

where:

β_i = the vector of weights that person i attaches to the observed attributes

W_i = the vector of weights that person i attaches to the un-observed attributes

ε_{ij} = an idiosyncratic component of preferences, specific to how person i evaluates alternative j.

Of course, we cannot observe the person specific attribute importance weights β_i and W_i. Rather, we seek to learn about these parameters by observing choice behavior. This is standard in choice modeling.[4] The innovation in Harris and Keane (1999) is to show that the stated attribute importance measures described in Table 2 can give us important additional information about how different people value the attributes, and enable us to develop better choice models.

Harris and Keane use the attitudinal questions to obtain information about the attribute importance weights as follows: First, we code the responses to the attribute importance questions as 1 for "doesn't matter," 2 for "like to have" and 3 for "have to have." Then, letting:

S_{ik} = the importance (1, 2 or 3) that person i <u>says</u> he/she assigns to attribute k,

β_{ik} = the weight that person i <u>truly</u> attaches to observed attribute k,

we assume that:

(2) $\beta_{ik} = \beta_{0k} + \beta_{1k} S_{ik} + \mu_{ik}$

where β_{0k} and β_{1k} map the 1, 2, 3 scale into utility units, and μ_{ik} is "measurement error." Thus, we are allowing for the possibility that respondents who <u>say</u> they value an attribute more actually <u>act</u>

[4] We assume the unobserved idiosyncratic preference terms ε_{ij} are independent type I extreme value distributed. Then, if we were to ignore the unobserved attributes A_j, and assume that preference weights β_i on observed attributes are equal for all respondents, we would obtain the multinomial logit model. Most applied choice modeling still uses this simple model, which assumes homogenous consumer tastes for observed product attributes. By allowing for preference weights to differ across consumers, and/or for unobserved common attributes, we obtain the "heterogeneous logit model."

as if they value the attribute more. If that is true, then we should obtain $\beta_{1k} > 0$ if an attribute is "good," and $\beta_{1k} < 0$ if the attribute is "bad."

For example, we have that:

$k = 1$ corresponds to the Premium (X_{j1}).

β_{i1} = the weight person i puts on premiums (presumably this is negative).

S_{i1} = the stated importance of low premiums(on a scale of 1 to 3).

A person who responds that a plan would "have to have" the lowest premium has S_{i1}=3. A person who responds that the premium "doesn't matter" has S_{i1}=1. If the stated attribute importance measures are indicative of actual preferences, then a person who says he/she would "have to have" the lowest premium (S_{i1}=3) will probably put a bigger (negative) weight on premiums in his/her utility function than one who says the premium "doesn't matter" (S_{i1}= 1). This means that in the equation:

(2') $\beta_{i1} = \beta_{01} + \beta_{11} S_{i1} + \mu_{i1}$

we expect the slope parameter β_{11} to be negative (i.e., the bigger the stated importance of premiums S_{i1}, the bigger will be the negative coefficient on premiums, β_{i1}).

The "measurement error" term μ_{ik} captures the fact that:

(i) People may not respond carefully to the questions (e.g., someone who says the premium "doesn't matter" might actually care quite a bit about premiums).

(ii) Different people may mean different things by the same answer (e.g., If two people say they would "Like to Have" low premiums, one may actually care quite a bit more about premiums than the other).

Problems like these are part of why economists have traditionally eschewed attitudinal data. It is important to stress, however, that the approach in Harris and Keane (1999) does not assume *a*

priori that the stated attribute importance data is a good predictor of individual level preferences. Rather, we let the choice data to tell us whether the attitudinal data is informative.

Intuitively, if people who say they care a lot about a particular attribute tend to choose alternatives with a high level of that attribute, then our estimates will indicate that the slope coefficients in equation (2) are significant.[5] In other words, if the stated attribute importance data helps to predict individual level choices, then our estimates will imply that it helps to predict individual level preferences. On the other hand, if the stated preference data is not useful for predicting behavior, then the measurement error terms will be "big," and the estimates of the slope parameters in (2) will tend to be insignificant and close to zero.

If it turns out that the attitudinal data are uninformative, so that the slopes in (2) are zero, then the intercept terms in (2) would tell us the average importance that people place on each attribute. This can be inferred from observed choice behavior alone, as in any simple choice model. Clearly, we can't learn more than the average preference weights (across all consumers in the population) if the individual level stated importance measures are uninformative.

As the final component of the model, we specify that the preference weights on the unobserved attributes are given by the equation:

(3) $W_{ip} = W_{1p}\, S_{ip}^{*} + \upsilon_{ip}$ $\quad p$=1 (cost share), 2 (quality).

This is like equation (2), except that S_{ip}^{*} denotes the person's stated importance for *un*-observed attribute p, the slope coefficient that maps the stated attribute importance into true attribute importance is now denoted W_{1p}, and the measurement error term is now denoted υ_{ip}.[6]

[5] Interestingly, the stated attribute importance data could also predict behavior because people who say they care a lot about an attribute tend to choose alternatives with low levels of that attribute. That is, the slope coefficients in (2) could be significant but with the wrong sign. This would mean that people care about the attribute, and that the attitudinal data helps measure how much they care about the attribute, but that their perceptions are inaccurate. That is, they think the health plans with high levels of the attribute actually have low levels of the attribute.

[6] We assume that the measurement error terms μ_{ik} in (2) and υ_{ip} in (3) have normal distributions. The variances of these distributions are additional parameters that must be estimated as part of the model.

The key difference between (2) and (3) is that (3) has no intercept term. Harris and Keane (1999) explain in detail why (3) does not contain an intercept, but the basic intuition is as follows: Consider an unobserved attribute like quality. Our method will infer that an alternative has high quality (as perceived by consumers) if people who say they care a lot quality tend to pick that alternative. Thus, if the stated importance of quality were not predictive of behavior, it would be impossible to estimate the perceived quality levels of each alternative. It would be impossible to estimate even the mean weight that people place on quality (let alone the distribution of taste heterogeneity).[7] Since we would have no information at all about how people value quality if the slope coefficients in (3) were zero, setting the intercepts to zero guarantees that we would not be trying to do the impossible (i.e., estimate the intercepts) in this case.

It is simple to estimate the model given by (1)-(3) using simulated maximum likelihood (SML). If the attribute importance weights β_i and W_i were known, the choice probability for a person would have a simple multinomial logit form. Since β_i and W_i are unobserved (we are estimating the parameters of their distribution), the simulated probability that person i chooses plan j is just the average over draws for β_i and W_i of multinomial logit choice probabilities:

$$P(j \mid \theta, S_i, S_i^*) = D^{-1} \sum_{d=1}^{D} \left[\exp(X_j \beta_i^d + A_j W_i^d) / \sum_{k=1}^{5} \exp(X_k \beta_i^d + A_k W_i^d) \right]$$

Here θ is the vector of all model parameters and S_i and S_i^* are attitudinal measures for person i.

II. C. The Parameter Estimates

Table 3 presents estimates of equation (2), which describes how people value the observed attributes of the insurance plan options. The estimates imply that the stated attribute importance data is highly predictive of individual level preferences, so that using such data does

[7] Recall that, in (2), if the attitudinal measures provide no information about preferences, then the slope coefficients will be zero, and the intercepts tell us the mean importance that people place on each attribute, as inferred from observed choice behavior alone.

indeed enable us to get a better predictive model. For each of the five observed attributes included in the choice model, the slope coefficient that maps the stated attribute importance measures into true attribute importance weights is significant and has the expected sign.

Table 3: Parameter Estimates, Observed Attributes

Observed Attribute:	Intercept	Slope
Premium	.014 (.011)	-.007** (.003)
Drug Coverage	.057 (.912)	.384** (.145)
Preventive Care and No Claims	1.887** (.498)	.766** (.202)
Provider Choice	-.395 (1.081)	1.430** (.489)
Must Submit Claims	Collinear with Preventive Care (Plans with preventive care do not have claims)	-.274** (.130)

Note: The "slope" coefficient must be multiplied by the stated importance weight S_i = 1, 2, or 3, and the result then added to the intercept to obtain the predicted importance weight for person *i*. Standard errors are in parenthesis below the estimates. A "**" indicates significance at the 5% level.

For example, Table 4 details how the model's prediction of the importance weight that a person puts on drug coverage differs, depending on whether the person says this is an attribute that he/she would "have to have," or would "like to have," or that "doesn't matter." Notice that the utility weight ranges from a low value of 0.441 if the person says the attribute "doesn't matter," to a high value of 1.209 if the person says it is an attribute that he/she would "have to have." Thus, consumers who say they "have to have" drug coverage act as if they place nearly 3 times as much value on that attribute as the consumers who say this attribute "doesn't matter." But does a coefficient estimate of 1.209 mean that these consumers care a lot about drug coverage? In a choice model, the best way to interpret the magnitudes of the coefficient estimates

it to look at what they imply about how changes in plan attributes would affect market shares, an exercise I'll turn to in section II. D.

Table 4: Predicted Utility Weight on "Drug Coverage" for Different Levels of Stated Importance

S=1 "Doesn't Matter"	S=2 "Like to Have"	S=3 "Have to Have"
.057+(1)·(.384) = .441	.057+(2)·(.384) = .825	.057+(2)·(.384) = 1.209

It is interesting that even consumers who say drug coverage is an attribute that "doesn't matter" act as if they place a significant positive value on drug coverage (according to our model estimates). This might seem inconsistent, but it is important to remember exactly how the stated attribute importance questions are phrased. Consumers were asked whether a plan had to have a particular attribute in order for them to consider the plan. It is perfectly consistent to answer that an attribute "doesn't matter" when deciding which plans to consider, but that the attribute would matter for which option one actually chooses.

Pursuant to this point, one might observe that the attitudinal questions in the Twin Cities data are actually phrased rather oddly if they are intended to measure preference weights. One might also question why we choose to code the responses as 1, 2 and 3. Is there any reason to think that the preference weight for a person who responds they "have to have" an attribute exceeds that of a person who responds "like to have" by exactly the same amount that the preference weight for a person who responds "like to have" exceeds that of a person who responds "doesn't matter"?

But, despite these problems, it turns out that responses to these rather imperfectly phrased attitudinal questions, coded in our admittedly rather coarse way,[8] are very predictive of actual choice behavior. In fact, the improvement in the log-likelihood function when we included the stated attribute importance measures in the model was over 100 points (from –1956 to –1834), a very dramatic improvement.[9] This was beyond our wildest expectations of how useful such data might be in predicting behavior. It is possible that more refined questions, or a more refined coding of responses, might yield a predictive model that is better yet. But they key point is that our exercise revealed the predictive power of even rather crude attitudinal measures.

Finally, Table 5 presents our estimates of equation (3) and of the unobserved attribute levels (A_j) for each insurance plan. Let's first consider the second unobserved attribute, quality of care. It is worth noting that we can only measure the quality of each plan relative to some base or reference alternative, since only differences in quality affect choices in our model. In Table 5, we set the quality of Basic Medicare to zero (i.e., it is the base alternative) and then estimate the quality of the other plans relative to Basic Medicare.[10] Thus, the positive estimates of A_2 for options 2 and 3 imply that consumers perceive these plans as providing higher quality than Basic

[8] It is worth noting that we are not really committing the sin of coding ordinal variables as cardinal variables, because we are not interested in using the model to predict how changes in consumers' stated attribute importance levels would affect choice probabilities. We are only interested in how changes in the attributes of the insurance plans affect market shares for each plan. As far as the stated importance weight measures are concerned, the only issue is whether our coding generates a variable that is a good predictor of individual importance weights (or whether some other coding might have provided a better predictor), not whether our coding is consistent with the scale of the attitudinal data (which would seem to be a rather amorphous concept anyway).

[9] One does not need to estimate a complicated heterogeneous coefficients model like the one we laid out in equations (1) through (3) to see the predictive power of the attitudinal data. If one estimates a simple multinomial logit model with the five observed attributes in Table 1 as predictors of behavior, and then compare this to a simple multinomial logit model that also includes interactions between the observed attributes and the stated attribute importance measures (thus letting the logit coefficients on each observed attribute differ depending on the stated attribute importance weight) the improvement in the log likelihood function is again roughly 100 points.

[10] Another technical point, explained at some length in Harris and Keane (1999), is that it is difficult to estimate both the scale of W_{lp} in equation (3) and the scale of the unobserved attribute levels A for each plan. To deal with this problem, Harris and Keane restricted W_{lp} to equal the inverse of the estimated standard deviation of the measurement error in equation (3), which, in turn, was restricted to be the same as the standard deviation of the measurement error in equation (2). Intuitively, these restrictions imply that the stated attribute importance measures are just as good at predicting peoples' preference weights on the unobserved attributes as they are at predicting peoples' preference weights on observed attributes.

Medicare. This is as we would expect, since options 2 and 3 are Basic Medicare plus medigap insurance that covers additional services. Thus, care under these options should be at least as good as under Basic Medicare alone.

Table 5: Parameter Estimates, Unobserved Attributes

Un-Observed Attribute Importance:
Estimates of Equation (3):

$W_{ip} = 2.688 \cdot S_{ip}^{*} + \upsilon_{ip}$ $\quad p$=1 (cost share), 2 (quality)

Estimates of the Un-observed attribute levels for each insurance plan

Un-observed (or "Latent") attribute 1:
Cost Sharing Relative to Basic Medicare

Basic Medicare	A_{11}	0
Medigap without Drug Coverage	A_{21}	-.270
Medigap with Drug Coverage	A_{31}	-.355
IPA type HMO	A_{41}	-.414
Group HMO	A_{51}	-.271

Un-observed (or "Latent") attribute 2:
Quality of Care Relative to Basic Medicare

Basic Medicare	A_{12}	0
Medigap without Drug Coverage	A_{22}	.269
Medigap with Drug Coverage	A_{32}	.261
IPA type HMO	A_{42}	-.081
Group HMO	A_{52}	.161

Note: The unobserved attribute levels for Basic Medicare are normalized to 0 since it is the base alternative. Attribute levels for the other plans are measured relative to Basic Medicare. In equation (3), S_{ip}^{*} is the weight (from 1 to 3) that person i says he/she puts on attribute p, and υ_{ip} is "measurement error."

The estimates of the perceived quality levels for the HMO plans are quite interesting. The negative value of A_2 for the IPA plan implies that consumers perceived the care provided under this plan as being low quality. In contrast, consumers felt that the care provided under the group

HMO plan was higher than under Basic Medicare. Still, the quality of care under the group HMO was perceived as lower than under Basic Medicare plus either medigap plan.

The results for the first unobserved attribute, cost sharing requirements, are rather surprising. As we see in Table 5, the estimates of A_{21} through A_{51} are all negative. Since the preference weight that multiplies this attribute is a preference for "low cost sharing," a negative attribute level means that the plan requires more cost sharing than the base alternative (Basic Medicare). Thus, these estimates imply that the survey respondents perceive every alternative health insurance plan as having greater cost share requirements than Basic Medicare. In fact, Basic Medicare has the highest cost share requirements of any option.

At this point, it's worth recalling the intuition for how we can estimate the levels of plan attributes that are not observed in the data, such as quality and cost sharing. Basically, if people who say they care a lot about quality tend (*ceteris paribus*) to choose a particular plan, it implies the plan is perceived as high quality. Similarly, if people who say they care a lot about low cost sharing tend to choose a plan, it implies the plan is perceived as having low cost sharing. Since the people who say they care most about low co-pays are also the most likely to choose Basic Medicare, our estimates imply that people perceive Basic Medicare as having low co-pays.

We can't readily judge if respondents' quality perceptions are accurate, because quality is so difficult to measure. However, as we noted earlier, its seems that quality must improve if you add a medigap plan on top of Basic Medicare, and our estimates are consistent with that. In contrast, while it is difficult to form an overall quantitative measure of co-pay requirements, we do know qualitatively that Basic Medicare has the highest co-pays of any plan. Thus, we can tell that respondents have rather fundamental mis-perceptions about cost sharing, even though we can't easily form an objective ranking of all five plans on the cost-sharing dimension.

There is a literature suggesting that senior citizens have mis-perceptions about Medicare and the supplemental insurance market. Examples are Cafferata (1984), McCall et al. (1986) and Davidson (1992). This is also a literature showing that consumers have difficulty understanding health insurance plans more generally. See, for instance, Cunningham et al. (2001), Gibbs et al. (1996), Isaacs (1996) and Tumlinson et al. (1997). Given this, it does not seem surprising that we find that senior citizens have mis-perceptions about cost sharing requirements.

Interestingly, however, our estimates do not imply consumer misperceptions about the five observed plan attributes in our model. That is, consumers who say they care a lot about premiums do act as if they place a relatively high weight on low premiums (in the sense that they tend to choose plans with low premiums), consumers who say they care a lot about drug coverage do act as if they place a high weight on drug coverage (in the sense that they tend to choose plans with drug coverage), etc.. Why should mis-perceptions be more important for cost-sharing requirements than for these other attributes?[11]

My hypothesis is that cost-sharing requirements are very hard for consumers to understand for the same reason they are hard for a researcher to measure/quantify. Health plans tend to specify a wide range of different co-pays that differ across treatments and the circumstances under which those treatments are obtained. Patients' out-of-pocket costs will also vary depending on how physician billing for a procedure compares to the reimbursement rate

[11] It is worth emphasizing that our method could have also implied consumer misperceptions about observed attributes. I discussed this in footnote 5. For example, if consumers thought the plans that allow provider choice actually did not allow choice (and vice-versa), then consumers who said they care a lot about provider choice would act as if they placed relatively small utility weights on provider choice. On the other hand, our results should not be taken as implying that consumer perceptions of the observed attributes (premiums, drug coverage, etc.) are completely accurate. They simply mean that perceptions of these attributes are sufficiently accurate to generate the correlation that those who say they care more about an attribute are also more likely to choose a plan that has that attribute. This is consistent with some inaccuracy of information. For example, even if consumers did not know the premiums for each plan exactly, but only knew the ranking of plans by premium, one would get the pattern that consumers who care more about premiums tend to choose plans with lower premiums. Perceived attributes would have to be negatively correlated with objective attributes to completely flip the sign of the slope coefficients in (2).

under Medicare or under the other plans, and according to whether particular procedures are covered at all. Given uncertainty about what services one will require, how one will be billed, and what any insurance plan will cover, it is very difficult for a trained statistician, let alone a typical consumer, to predict future out of pocket costs conditional on enrollment in a particular health care plan. Thus, cost-sharing requirements may be harder to understand than other plan features, since they vary a lot by procedure, and may not be experienced until an illness occurs. In contrast, a plan attribute like provider choice is more evident "up front," since, for example, one either chooses a doctor or not when one joins a plan.[12]

This finding has important implications for the design of "competitive" health insurance markets. As Hall (2004) notes: "to choose rationally across insurers [consumers] must be well informed about … the plans offered. … It is worth noting that many consumers … have not had substantial experience in obtaining health care until they face … illness." Thus, our finding that consumers have important misperceptions about their insurance options undermines a key tenet of the standard "choice is good argument."

II. D. Simulations of the Model

Given an estimated choice model, one can use it to simulate the impact of a change plan attributes (like premiums or drug coverage) on the market shares of the various plans. One can also use the model to predict whether there would be substantial demand for new plans with particular attributes. Some examples of these type of simulations are provided in Table 6.

[12] An alternative hypothesis is that people with low incomes place a great weight on low co-pays, but that they simply cannot afford supplemental insurance or the extra cost of joining an HMO. We find this story implausible for two reasons. First, we dropped respondents who used Medicaid, the medical insurance program for the poor, or who has SSI benefits (which are disability benefits), or who couldn't pay the Medicare Part B premium of $28 per month. Thus, the poorest respondents are not represented in the data. Second, the HMO options only cost a little more than Basic Medicare, so it seems implausible that liquidity constraints would preclude those options.

Table 6: Some Illustrative Experiments Using the Model

	Basic Medicare	Medigap w/o Drugs	Medigap w/ Drugs	IPA	Group HMO	"New Plan"
Baseline	9.1%	9.4%	12.4%	25.6%	43.6%	
Medicare adds Drug Coverage	17.7%	8.2%	10.9%	22.2%	41.2%	
IPA adds Drug Coverage	6.7%	7.1%	9.1%	41.7%	35.5%	
IPA plan removes Provider Choice	11.4%	12.1%	16.3%	2.3%	57.7%	
Add "New Plan"	6.8%	7.4%	9.9%	19.6%	30.6%	25.8%

The first row of Table 6 reports a "baseline" simulation of the model, which simply gives the model's predictions regarding the market shares of the various plans. These predictions line up reasonably closely with the actual market shares observed in the data, although the model somewhat overstates enrollment in the IPA plan (25.6% predicted vs. only 21.7% in the data) and in the group HMO (43.6% predicted vs. only 36.4% in the data) and correspondingly under-predicts actual enrollment in the Medicare and medigap options.[13] A notable aspect of the Twin Cities health insurance market is the very high penetration rate of the Medicare HMOs. Nationwide, participation in such plans is quite a bit lower.

The second row of Table 6 reports our model's predictions of what would happen to the market shares of the five plans if Basic Medicare were to add prescription drug coverage. The model predicts that the market share of Basic Medicare would increase substantially, from 9.1% to 17.7%. This suggests that many consumers find prescription drug coverage to be a very attractive feature of a health plan. This impression is reinforced in the third row of Table 6,

[13] Our choice model could be made to fit the overall market shares of the five plans just about perfectly if we were to include plan specific intercepts. The problem with including intercepts is that it makes it impossible to predict the what market share would be for a new plan with a particular set of attributes, because we wouldn't know how to set its intercept. As Elrod and Keane (1994) discuss, an intercept captures average consumer tastes for the unique attributes of an alternative.

which shows the model's prediction of what would happen if the IPA plan were to introduce drug coverage. The model predicts that its market share would increase substantially, from 22.2% to 41.7%.

Similarly, the fourth row of Table 6 presents the model's prediction of what would happen if the IPA plan were to remove provider choice. The model predicts that its market share would dwindle to almost zero (2.3%). This is not surprising, as in this case the IPA plan would be completely dominated by the Group HMO. That is, it would have a slightly higher premium, it would not cover drugs while the group HMO does, and it would have worse perceived quality and higher perceived cost-sharing (see Table 5). Other simulations (not reported here) implied that shares of the medigap plans would drop substantially if they were to restrict provider choice.

In other simulations reported in Harris and Keane (1999) we found that moderate changes in premiums (i.e., $20 per month increases) would have very small effects on plan enrollments. Thus, our estimates imply that consumers care quite a lot about provider choice and prescription drug coverage, but that they aren't very sensitive to premiums (at least not within the rather limited range of premiums exhibited in these data).

In the bottom row of Table 6, we use the model to predict what would happen if a new health insurance plan were introduced. The "New Plan" is designed to fill a gap that existed in the Twin Cities insurance market. Consider a segment of consumers who place a high value on provider choice and preventive care, but little value on prescription drug coverage. Given the structure of the Twin Cites market in 1988, the plan best tailored to these tastes was the IPA plan. However, the IPA plan was perceived as being of very low quality (as well as having very high cost sharing), thus leaving these consumers without a very appealing option. The fact that so many people choose the IPA plan anyway (21.7%) suggests that this configuration of

preferences is rather common. The "New Plan" was designed to be like the IPA on observed attributes, but to have the same perceived quality as the group HMO (A_{62}=.161) and to have less perceived cost sharing (A_{61}=-.150).

Our model predicts that the "New Plan" would be very popular, with a market share of 25.8%. This implies a substantial welfare improvement from its introduction (holding other plan attributes fixed), since every consumer who chooses the "New Plan" is better off than they were before, while consumers who stay with the existing plans are made no worse off. Note that the "New Plan" differs from the group HMO primarily in that it allows provider choice but doesn't cover drugs. Our estimates imply that a substantial segment of the population likes that option, provided it is also of reasonably high quality.

One could use the model to formally calculate the increase in consumer surplus that arises from introducing the "New Plan," holding existing plan features fixed. But, Harris and Keane (1999) did not do that, so I can't report the calculation. Thus, I'll stick with the informal statement that the welfare gain is "large" since the new plan would be quite popular.[14]

II. E. The Importance of Controlling for Unobserved Attributes

A key finding in Harris and Keane (1999) was that failure to control for the unobserved attribute dimensions of cost-sharing and quality leads to severe bias in estimates of consumer preferences for the observed attributes of insurance plans. Most notably, when we estimated

[14] Consumer surplus is the sum over all consumers who buy the new plan of the difference between what they would be willing to pay for it and what they actually have to pay (i.e., the premium). The calculation is actually rather trivial if one takes the choice model literally. However, such welfare calculations can be rather sensitive to the shape of the demand curve implied by the model at very high price levels. The logit model, because of the extreme value error assumption, implies that some small number of people would want to buy any new product even at a very high price. The model therefore predicts huge welfare gains for this small group when a new product is introduced. Since we wouldn't really take the model's predictions of demand extrapolated to very high prices literally, it may, in practice, be better to stick with the informal statement that welfare gains are large if the new product is popular, or, in doing the formal calculation of consumer surplus, to truncate consumer willingness to pay at some maximum value. Surplus calculations will always be somewhat arbitrary since we can never directly observe willingness to pay, only demand.

models that ignored the unobserved attributes,[15] the estimates implied the completely implausible outcome that consumers dislike provider choice.

The reason for this odd outcome is as follows: Only the Group HMO restricts provider choice, but this plan has a very high market share. Thus, a model that ignores quality as a determinant of insurance plan choice has to assume that consumers don't care about provider choice in order to explain the high market share of the Group HMO. In contrast, our model estimates imply that the Group HMO has high perceived quality, which we infer because consumers who say they care a lot about quality are very likely to choose the Group HMO. Because of this, our model can explain the high market share of the Group HMO on the basis of perceived quality, rather than by assuming consumers don't care about provider choice.

For the econometric sophisticate, let me point out that this argument can also be stated as follows: Observed insurance plan attributes are "endogenous" in the statistical sense that they are correlated with the error terms (i.e., unobserved plan attributes). The use of stated preference or attitudinal data to control for unobserved plan attributes and obtain consistent estimates of preference parameters is an alternative to the conventional econometric approach of using "instrumental variables." But, unlike instrumental variables, this approach works in non-linear models, like the heterogeneous logit model considered here. This observation is a key part of the methodological contribution in Harris and Keane (1999).

II. F. Summary of Main Findings and Subsequent Work

The main findings of the empirical choice modeling exercise in Harris and Keane (1999) can be summarized as follows:

[15] These models included simple logit models that use the observed plan attributes to predict choices, as well as heterogeneous logit models that allow for consumer heterogeneity in tastes for observed plan attributes but that do not estimate unobserved plan attribute levels. The latter can be obtained just but setting all the A parameters equal to zero in the model described in section II.B.

1) Consumers are not very sensitive to premiums when they choose health insurance plans (at least not within the rather limited range of premiums exhibited in the Twin Cities Medicare data).

2) Many people care a lot about Drug Coverage and Provider Choice when choosing a health insurance plan (i.e., plans' market shares are quite sensitive to these attributes, and people are willing to pay quite a lot for them).

3) Senior citizens have important misperceptions about the cost sharing requirements of Basic Medicare vs. medigap and HMO options. People who say they care most about low cost-sharing are the most likely to choose Basic Medicare (which has the highest cost-share requirements of any option).

Subsequently, Harris, Feldman and Schultz (2002) – henceforth HFS – used a similar methodology to analyze insurance plan choices of employed workers who were under 65, and hence not yet eligible for Medicare.[16] HFS used data from the Buyers Health Care Action Group (BHCAG). The BHCAG is a coalition of two-dozen employers in the Twin Cities area that contracts directly with health care providers (rather than negotiating plans with insurance companies). Employees of BHCAG member companies have a choice among several alternative health insurance plan options. Employees were surveyed about their plan choices in 1998, and they were also asked a series of questions about how much they valued various plan options.

Similar to Harris and Keane, HFS used questions about how much consumers valued various aspects of quality, along with choice data, to infer perceived quality levels of the various plans. The HFS study differed from Harris and Keane in several ways: (1) they attempted to uncover different dimensions of perceived quality, (2) the plans offered by the BHCAG coalition

[16] Parente, Feldman and Christianson (2004) use this approach to study health plan choices of University of Minnesota employees.

had identical cost sharing requirements, so HFS made no attempt to estimate the effect of perceived cost-sharing on choices, and (3) HFS pretended they did not observe premiums, in order to ascertain if the Harris and Keane methodology could successfully uncover the premium differences across plans by using data on survey respondents' stated importance of premiums.

Like Harris and Keane, HFS found that the use of stated attribute importance data led to dramatic improvements in model fit, and also led to more sensible coefficient estimates for observed attributes. They found that premium differences across plans were accurately uncovered by the methodology. Their estimates imply that perceived quality differs greatly across plans. When quality is decomposed into different components, what appears to have the biggest impact on choice is service quality (i.e., access to specialists, convenience of clinic locations, wait time for specialist appointments) rather than provider quality. This result is consistent with a literature suggesting that consumers tend to pay relatively little attention to various measures of provider quality.

III. Choice Modeling and the Design of "Competitive" Health Insurance Markets

Having described how choice models can be used to (i) analyze consumer preferences for attributes of health insurance plans, (ii) predict demand for new health insurance products (with particular attributes), and (iii) predict consumer welfare effects of adding new insurance products, I would now like to discuss the potential role of choice modeling techniques in the design of "competitive" health insurance markets.

Let me start by noting the main limitation of the empirical application that I described in section II. Unfortunately, the Twin Cities Medicare data does not contain information on health status and retrospective service use that would be critical for forecasting medical expenses of each respondent. In most marketing applications of choice modeling, all one cares about is

predicting market share of products (including new products or existing products given a change in price or other attributes). The cost of supplying a product doesn't typically depend on the identity of who buys it, but only on the quantity supplied (i.e., on production costs). But a key problem in health insurance markets is that the cost of providing the service clearly does depend on the identity of the consumers who buy it.[17]

In forecasting the cost of a particular insurance plan, we cannot assume that the consumers who choose each plan would be typical of the population. We need to consider adverse selection (i.e., those who choose more comprehensive insurance are likely to be those who will require more services), as well as possible moral hazard and elastic demand (i.e., a person in a more comprehensive plan may use more services than would the same person in a less comprehensive plan, either because he/she takes worse care of himself when protected by insurance, or simply because demand for services is elastic with respect to price). Thus, one cares not only about predicting market share of each insurance plan, but also about predicting characteristics of purchasers – especially their utilization of services.

Why is predicting utilization important? For a private health insurance plan to be sustainable, the premiums plus cost sharing (plus any government subsidies) must be sufficient to cover the expense of providing services to the participants. Thus, we need to predict utilization to determine if a plan provided by a private firm would be sustainable.

On the other hand, if there is a single payer (government) providing an array of insurance options, the differential between revenue and cost of each plan determines the pattern of cross-subsidization needed to sustain the menu of plans. For any proposed menu of insurance options that the single payer might wish to provide, one wants to look at the cross-subsidy pattern, along

[17] Of course, there are other markets besides health insurance where this is true, e.g., credit cards, phone service, and auto insurance. In each of these cases, a marketer would also care about forecasting the identity of buyers.

with premiums and co-pays in each plan, to analyze equity and sustainability. Here, a plan may be sustainable even if it loses money, provided the cross-subsidy needed to sustain the plan is politically and financially feasible.[18]

There are some fundamental similarities between private vs. public insurance provision:

1) In either case, there may be "adverse selection," whereby "sicker" people choose more "generous" plans. To cover costs, the more generous plans must charge higher premiums than if their participants had average health. This works against equity.

2) With a single payer, the government can use cross-subsidies to enhance equity. Similarly, with private insurance, the government can provide bigger subsidies to plans that take on "sicker" populations.

A system of government subsidies to private insurance plans, based on the expected utilization of services of those enrolled in the plans, is called "risk adjustment." Van de Ven and Ellis (2000) provide an excellent discussion of risk adjustment methodology and the challenges one faces in implementing it.

However, there is one key difference between private vs. public insurance provision: Given any system of risk adjusted subsidies that the government puts in place, the private insurers will have an incentive to "cherry pick." That is, private firms will design plans so as to attract people who are predictably profitable (i.e., their expected utilization costs are low relative

[18] Note that there is a basic tension here between the tripartite goals of providing choice, achieving equity and achieving political sustainability. To give a simple example, suppose the government offers a limited insurance plan at a low premium, and a very comprehensive plan at a higher premium. The premiums are set so the limited plan is preferred by the healthy while the comprehensive plan is preferred by the unhealthy. If there is to be choice, the premium difference must be large enough to maintain this "separation" (i.e., if the difference were too small the healthy would no longer prefer the limited plan, and only the comprehensive plan would continue to exist, and vice-versa). But equity concerns suggest that the premium difference should not be made too great. On the other hand, if the premium difference is set too small, then the cross-subsidy from the healthy to the sick may become too great to be politically sustainable.

to premiums plus co-pays plus subsidies). But, in contrast, this cherry picking problem does not arise in a single payer public system where the government offers a menu of options.

The cherry picking problem that would arise under any system where private firms offer insurance plans in competition with government undermines the "choice is good" argument. If new plans are designed with the goal of cherry picking, consumer welfare will not generally increase, since the assumptions underlying "Choice is Good" and "Competition is Good" do not hold. For instance, if the low cost people are drawn away from Basic Medicare into private Medicare HMOs, known as "Medicare+Choice" plans, then Basic Medicare premiums (or else taxes) must rise to cover increased average costs, and those left in Medicare are worse off.[19] Under this scenario, a key assumption of the "choice is good" argument is violated, since attributes of the existing plan do not stay fixed when the private insurance option is introduced.[20]

This phenomenon is of more than academic interest. According to GAO (2000), "… we estimate that aggregate payments to Medicare+Choice plans in 1998 were about $5.2 billion (21percent) … more than if the plans' enrollees had received care in the traditional FFS program." In general, there appears to be a wide consensus in the literature that Medicare HMOs in the U.S. have achieved at least a substantial part of their cost reductions via cherry picking. For example, see Glied (2000), Greenwald, Levy and Ingber (2000), Brown et al (1993).

Of course, the cherry picking problem can be avoided if government can design its system of risk adjusted subsidies in such a clever way that it is impossible for private firms to

[19] Currently, there is a direct link between costs of the Medicare fee-for-service program and premiums. To see why, one needs to understand that Basic Medicare consists of Part A, which basically covers hospitalization, skilled nursing facilities, hospice and home health care, and Part B, which basically covers outpatient hospital, physician office visits, and laboratory services. Medicare enrollees get Part A for "free" (i.e., it is financed from payroll taxes), but pay a premium for Part B. Under the Balanced Budget Act (BBA) of 1997, the part B premium is set equal to 25% of expected Part B spending, so any increase in Part B costs is directly reflected in premiums. Although the link is not direct, higher Part A costs are likely to be reflected in higher Part A deductibles and co-pays.

[20] An example from the economics of education that is analogous is the introduction of private competition to public schools. If the private schools attract the "best" students, then the public school students may be made worse off.

locate and attract predictably profitable people. So, a key question is this: Could the government design a system of risk adjustment that is good enough so that private insurance firms couldn't "beat the system" and find ways to attract predictably profitable (i.e., low cost) clients? I strongly suspect that the answer is "No," especially since risk adjustment must be based on observed attributes like age, sex, region, ambulatory cost group, etc., and there is substantial heterogeneity in cost even with such groups (see Shen and Ellis (2002)).[21, 22]

It is worth emphasizing that the cherry picking problem arises because "asymmetric information" (i.e., consumers know more about their health state than do insurers) leads to adverse selection (i.e., more comprehensive insurance plans will tend to attract unhealthy, and hence high cost, consumers). In an important series of papers, Rothschild and Stiglitz (1976), Wilson (1977) and Spence (1978) studied the nature of competitive equilibrium in markets with adverse selection. Basically, these papers show that one tends to get segregation of consumers, with the unhealthy buying comprehensive insurance at high premiums, and the healthy buying limited insurance at low premiums. This creates both equity and efficiency problems. Obviously, the unhealthy end up paying high premiums. More subtly, the equilibrium is inefficient because the healthy are led to underinsure, since that is the only way they can get low premiums. If the inexpensive health plans aimed at the healthy were to cover too much, then at some point the unhealthy would find them attractive, and they couldn't remain inexpensive.

[21] Furthermore, in the real world, insurance company lobbyists would probably play a major role in designing the risk adjusted subsidy system. The Medicare Modernization Act of 2003, which raised capitation payments to help bail out the failing Medicare HMO sector, appears to be a good illustration.

[22] Note that Clinton's health care plan, described in footnote 2, involved a risk adjustment system. To understand this, note that the Clinton plan did not alter the private insurance system. Rather, it simply added another layer to the system (i.e., the alliances). The alliances were supposed to negotiate with private insurance companies to create menus of health insurance options. They would then collect premiums from the alliance members and their firms, and use risk adjustment techniques to determine payments to the private insurance companies. The hope was that the alliances could do this in a sophisticated enough way so as to avoid cherry picking behavior by the private insurers.

But, as Wilson (1977) and Spence (1978) pointed out, equity and efficiency gains are often possible in such a market if the government can engineer a premium subsidy from the healthy to the unhealthy. If the plans that appeal to the healthy cross-subsidize the plans that appeal to the unhealthy, it becomes possible for the healthy to get more comprehensive insurance. Since the subsidy lowers the premium in the comprehensive plan, the unhealthy are better off. Furthermore, the limited plan aimed at the healthy can expand its coverage without attracting the unhealthy. As long as the subsidy that the healthy must pay to the unhealthy is less than their willingness to pay for this expanded coverage, they are made better off too.

The problem with this idea is that private insurance firms aren't going to offer profit making policies to the healthy and use them to cross-subsidize loss making policies aimed at the unhealthy. Government regulation or intervention is necessary. Wilson (1977) pointed out that one way to implement a cross-subsidy is for the government to require all consumers to purchase a "Basic" insurance policy, and to allow private insurers to offer supplemental policies. Since all consumers get partial coverage from the Basic plan at the same premium, a subsidy from the healthy to the unhealthy is implemented. But the unhealthy can buy supplemental coverage, and do so without causing the Basic plan risk pool to split up (since everyone is required to belong).

This was exactly the situation in the U.S. with Basic Medicare and medigap insurance, prior to the advent of Medicare HMOs. The latter don't fit into this framework, because they allow people to opt out of the Basic Medicare risk pool. Indeed, it is uncontroversial that HMOs have drawn relatively healthy people out of the Basic Medicare risk pool.[23]

[23] Prior to the Balanced Budget Act (BBA) of 1997, Medicare HMOs received a per participant subsidy ("capitation payment") that was 95% of average spending for people in the same age/gender/region in Medicare fee-for-service. As noted earlier, there is general agreement that this primitive risk adjustment system was inadequate to eliminate profits from cherry picking. The BBA complicated the formula for paying Medicare HMOs, but did not resolve the basic problem. It is worth emphasizing, however, that problems created by cherry picking can remain severe even if risk adjustment is done so well that potential profits from cherry picking are rendered quite small. Even then, the market is contorted by the fact that private insurers design insurance plans with cherry picking objectives in mind.

Wilson (1977) and Spence (1978) pointed out that an equivalent way to implement a cross-subsidy is to have the government offer two insurance options: a comprehensive policy aimed at the unhealthy, and a more limited policy with a lower premium aimed at the healthy. Unlike private insurers, the government is willing to use the former plan to subsidize the later.[24]

For these reasons, I would argue that a very different approach to offering consumers choice in health insurance plans should be seriously considered. The alternative is for a single payer (i.e., government) to offer consumers a menu of health insurance plan options, with the techniques of choice modeling being used to design those options in a way that would appeal to heterogeneous consumer tastes.[25] This approach has two key virtues. First, since the government designs the menu of insurance plans, they can, at least in principle, be designed with the goal of appealing to consumer tastes – as opposed to having their design being contorted by cherry picking objectives. Second, the government can attempt to enhance consumer welfare by generating cross-subsidies from the plans aimed at the healthy to those aimed at the unhealthy.

Now, all this is fine in theory, but, as Spence (1978) pointed out, actually implementation of a menu of insurance options that would increase both equity and efficiency requires that we know a great deal about consumer taste heterogeneity. It also requires that we be able to predict

---

Such plans will not in general be optimized to appeal to consumer tastes, as there will be over (under) supply of attributes that appeal to healthy (sick) consumers (see Frank, Glazer and McGuire (2000), Miller and Luft (1997)).

[24] Glazer and McGuire (2000) discuss how risk adjustment can also be designed, in principle, to implement a cross-subsidy. The idea is design the capitation payments so that it is profitable for private health plans to take on relatively unhealthy populations. For example, the government could pay HMOs more than expected costs for older individuals, and less than expected costs for younger individuals. But, given the tremendous heterogeneity in health status within any age or other demographic category, I am skeptical of whether private firms couldn't "beat" any such system the government devises by using clever enough marketing (e.g., I assume that Medicare HMOs could use market research to find attributes desired by healthier senior citizens within given age ranges).

[25] Wilson (1977) shows that, in a world where insurance plans are differentiated on only two dimensions (premium and a uni-dimensional measure of coverage), the single payer menu-of-options approach and the approach of allowing private firms to offer supplemental policies in addition to a required basic plan are actually equivalent, since one can optimize the subsidy from the healthy to the unhealthy by appropriate choice of the basic plan. But, if there are several plan features over which consumer have heterogeneous preferences, the menu of options approach would appear to be more flexible.

the average costs (i.e., the health services utilization) of the type of consumers who would select each plan offered in a particular menu. This poses a formidable econometric challenge.

The main point I would like to make is that, in my view, the necessary econometric techniques to pursue such a strategy may now be available. Using state-of-the-art choice modeling methods such as those developed in Harris and Keane (1999), we can estimate the distribution of consumer tastes for various health plan features. Then, given any hypothetical menu of insurance options that a single payer might offer to consumers, a choice model could be used to predict the market shares of each plan, as well as the composition of people who choose each plan. Next, we can also develop models of health care service utilization, and predict the cost of offering each plan as a function of the type of consumers who would select into it. Of course, for this to be possible, we need the data to include good predictors of utilization, like health status, prior health care utilization etc..

The approach I am describing would enable us to simulate the cross-subsidy pattern under any hypothetical menu of insurance options. One could then determine if the pattern of cross-subsidies needed to sustain that menu appeared socially and politically acceptable (i.e., do the premium differences appear "fair"?). We could also calculate consumer welfare under alternative menus that might be offered, subject to the constraint that the menu as a whole must break even (i.e., the plans that make losses must be subsidized by plans that make profits). And we could analyze choice by asking whether all important consumer segments are "covered," in the sense that a plan appealing to their tastes is offered.[26]

[26] Subsequent to my CHERE lecture, Randall Ellis pointed out to me that the single-payer menu of options approach I am advocating here is quite similar in spirit to the health insurance reform plan proposed by Diamond (1992). He advocated that the government of the U.S. divide the whole population into regional groups. Private insurance companies would then bid for the right to offer a menu of health insurance options to each group. His reason for wanting a single payer to offer the whole menu of options was the same as mine: "With joint bids on the entire menu, competition among … options … would be based primarily on … individual preferences over different ways of managing care, not the attempt to attract the best risks." The key difference in Diamond's plan is his proposal that

One idea for achieving equity is to design a menu of options, and then price each one so as to cover its predicted expenses conditional on its enrollees being drawn from a random sample of the population. For instance, the premium for a plan that offered drug coverage could be set at an increment above that of an otherwise equivalent plan that did not, with the increment set equal to the expected cost of drugs for a random member of the population.[27] In this way, peoples' premiums would be based only on the services they choose, and not on the risk pool to which they belong. This scheme would involve cross-subsidies from the healthy to the unhealthy, since the premiums paid by the healthy would exceed their expected utilization costs.

It appears that the computational problem faced by the government in pursuing this single-payer menu-of-options strategy would be much simpler than the problem it would face in designing a risk adjustment system such that, given the pattern of private insurance firms' response to the system, the resultant market outcome is desirable. Either approach involves "central planning," but the planning under the single payer system would be much simpler to implement. The reason is that under the private competition scenario, the government actually

---

private insurance companies offer the menu of options, rather than the government. As I see it, he adopts this approach for two reasons: 1) a view that eliminating private insurance companies is not a politically feasible option in the U.S., 2) the notion that the design of a menu of insurance plan options is something that private insurance firms would do better than government (as long as the cherry picking incentive is removed by requiring a single firm to offer the entire menu). Diamond also stresses the innovative role of the private market: "Having private insurance companies offering different ways of organizing medical services may well be a useful part of exploring alternatives and of adapting to the ever changing technology, both of medical services and of administration." In principle, my approach could also be implemented by having the government design a menu of health insurance options, and letting private insurers bid on the right to offer the whole menu. However, a potential problem with a Diamond like approach is that private insurance companies, even if required to offer a menu of options, might be able, through their marketing efforts, to manipulate which consumers enroll in each plan in such a way as to shift people into plans where they would be most profitable. Under the scheme that I have proposed, the government would need to accept bids from provider organizations for provision of the plans. The provider organizations would have incentives to hold down costs, and the government would need to regulate quality. An advocate of Diamond's approach might argue that private insurers would be better at negotiating with providers than government, but the experience of Medicare since 1996 contradicts such a view. During that period, Medicare used its substantial monopsony power to negotiate steep discounts from hospitals and physicians, leading to excellent performance in cost containment. The Medicare Modernization Act of 2003 interfered with this process, leading to substantial recent cost increases.

[27] Under such a pricing scheme, consumers with expected utilization of prescription drugs below some threshold would prefer the plan without drug coverage, while those expected utilization above some threshold would prefer the plan with drug coverage. The threshold would equal the expected level of utilization for the population as a whole, minus some increment since consumers are risk averse.

has to predict the behavior of the private firms in order to predict the market outcome, and this problem is an order of magnitude more difficult than that of just predicting market shares and utilization rates of participants if the government offers a menu of options itself.[28]

Of course, a key problem with the single-payer menu-of-options strategy that I have proposed is that the necessary modeling to implement the strategy requires very rich data. Using choice modeling techniques to predict market shares from any given menu of options would be rather straightforward. In fact, even if market data like the Twin Cities Medicare data were not available, there are well-established methods for applying choice modeling to so called "stated preference" (SP) choice experiments that could be used instead. In these approaches, consumers are asked what choices they would make among hypothetical menus of insurance plan choices with which they are presented. Harris (2002) is a good example of this approach.[29]

However, finding the data necessary to also predict the health plan utilization is more challenging. In particular, I know of no single data set that contains all the data necessary to forecast both insurance choice and utilization of the respondents who choose each plan. In order to model choice, one needs to know the insurance plan options that each person in a data set faced. In order to predict utilization, one needs information on personal demographics, health status and prior utilization. One also needs data on the characteristics of the insurance plan in which a person is actually enrolled (since a person with given characteristics would generally have different utilization of services under plans with different coverage).

---

[28] In fact, Rothschild and Stiglitz (1976) and Wilson (1977) point out that, in insurance markets with adverse selection, our predictions regarding the equilibrium behavior of firms in terms of what policies they will offer is extremely sensitive to the particular equilibrium concept one adopts.

[29] Harris (2002) uses SP choice experiments to analyze how giving consumers more information about health plan quality would affect their choices. She finds that the availability of quality information (either in the form of expert or consumer assessments) causes the impact of HMO network features on choice to fall substantially. This suggests that consumers use features like a large network or the ability to self-refer to specialists as a signal of high quality, or perhaps as insurance against low average physician quality. This type of question would be very difficult to examine using market choice data, given the difficulty in finding the right variability in information regimes.

Unfortunately, data sets like the Twin Cities Medicare data, that enable us to model choice, don't have the additional information needed to model utilization. And data sets like the household component of the U.S. National Medical Expenditure Survey of 1987 (NMES), or its successor, the household component of the Medical Expenditure Panel Survey (MEPS) begun in 1996, the that enable us to model utilization, don't have information on consumers' insurance choice sets. They only have information on the plan in which a person was actually enrolled. So such data sets cannot be used to model choice.

The NMES and MEPS also contain establishment surveys in which employers are asked about the insurance options they offer to employees. Cardon and Hendel (2001), Blumberg, Nichols and Banthin (2001) and Vistnes and Banthin (1997) have linked the household and employer components of these data sets in order to model insurance choice. As Blumberg, Nichols and Banthin discuss, the success rate in linking is only about 30%, so there is a serious issue of whether the linked sample is representative. Of greater concern, in my view, is that the linked samples only contain about one to two thousand people. This is far too small a sample size to reliably model utilization, given that a small fraction of people account for most medical costs.

One possible strategy is to use different data sets to estimate different parts of the model. For instance, one could think about using a data set like the Twin Cities data to model choice, and then using the NMES or MEPS household surveys to model utilization. In this strategy, one would develop utilization models based on the NMES or MEPS that predict utilization based only on characteristics of respondents that were also collected in the Twin Cities data (i.e., age, gender, income). Then, given predictions from our choice model of the demographics of respondents who would choose a particular plan, we could predict utilization based on those same demographics using the NMES or MEPS data.

A problem with this multiple data set strategy is that the demographic information in the data sets available for choice modeling is not sufficiently rich to construct good predictive models of utilization.[30] A promising alternative strategy is to collect new insurance choice data from SP choice experiments, and, in this new data collection effort, to also obtain from the respondents rich information about health status and medical history. One could then use the SP choice model to predict the characteristics of responds who chose each insurance option based not only on simple demographics like age, gender and income, but also in terms of health and medical history. All these variables could then be used to predict utilization, based on NMES type data. In my view, this is a critical avenue for future research.

In the choice modeling part of this exercise, market choice and SP choice data could be combined to create a better predictive model. Specifically, choice models based on market and SP data should predict similar market shares for insurance plans, both unconditionally (i.e., for the populaton as a whole) and conditional on the demographic information that is common to the data sets used to estimate each model. See Hall, Viney, Haas and Louviere (2004) for discussion of health applications of SP choice modeling, and Hensher, Louviere and Swait (1999) or Louviere, Hensher and Swait (2000) for discussions of merging market and SP choice data.

Obviously, in any effort to develop a single-payer menu-of-options system, the forecasts of market shares and utilization used in the design stage would only serve as preliminary guides. Actually implementation of such a plan would, of course, reveal discrepancies with the forecasts, presumably leading to subsequent modifications of premiums and/or plan attributes to achieve more precisely the desired choice, equity, and cross-subsidy patterns. Econometric forecasting

[30] More subtly, there is also a selection problem that arises because the average expenditures among people who chose a particular plan will in general differ depending on the original choice set – since the segment of the population that selects into a particular plan depends on the choice set. For this reason, modeling of utilization should be done jointly with modeling of insurance plan choice.

can be a useful guide to designing policies and predicting their impacts, but forecast errors and subsequent revisions have to be expected.

Finally, while I have focused on a "maximal program" of designing a single payer system, let me conclude this section by arguing that choice modeling should play a role even in more modest reform proposals. Indeed, more modest reforms, like adding benefits to Basic Medicare while continuing to allow private competition, or regulating private insurers to require that they offer certain types of coverage, are certainly more politically feasible. It strikes me as extraordinary that Clinton's Health Security Plan, which required alliances to offer plans with certain features, was not based on any attempt to ascertain how consumers value those features. Similarly, the MMA of 2004 added a prescription drug benefit to Medicare without any attempt to use market research to ascertain the distribution of consumer willingness to pay for such a benefit. Such seat-of-the-pants policy making seems very unwise.

IV. Challenges in Designing a Single-Payer Menu-of-Options System

In the preceding sections I have argued that the "cherry picking" problem can be avoided if a single payer offers a menu of insurance options. However, designing an appealing menu of options to be offered is a challenging problem. In this section, I will discuss the problems that adverse selection and moral hazard created for design of a menu of options, relate this to some recent work on the practical relevance of adverse selection and moral hazard problems, and describe how choice modeling techniques might be used to shed further light on these issues.

IV. A. Adverse Selection and Felicitous Selection

I've already commented at some length about the problems that adverse selection creates for the design of the menu of options to be offered in a "competitive" health insurance market. Basically, we would expect that "sicker" people will select into more generous insurance plans.

This makes it hard to predict the impact of offering any particular menu of insurance options, either in terms of total cost or the resultant cross-subsidy pattern: To predict the cost of each plan we have to also predict the health characteristics of the people who will select into each plan.

However, while we often take it for granted that adverse selection is a serious problem, some recent work by De Meza and Webb (2001) suggests its importance may be exaggerated. The reason is that adverse selection may be offset by what they call "advantageous selection," or what I will call "felicitous selection." The idea is that it is not only sicker people who want to buy more generous health insurance. More risk averse people want to buy more generous insurance too. If more risk averse people also tend to be healthier, then healthy people may demand just as much insurance as unhealthy people. In fact, it seems quite plausible that more risk averse people would be healthier, because their risk aversion may also induce them to take better care of their health.

In my view, there is some interesting qualitative evidence suggesting that felicitous selection is important. This comes from comparing the market for health insurance with some other insurance markets. As is well known, adverse selection can cause a market for insurance to fail to exist, or, at least to make the market quite small. The reason is that adverse selection may lead to a vicious cycle of rising premiums that reduces demand to low levels or even to zero.[31]

[31] The "death spiral" can arise as follows (see Akerlof (1970)): Suppose that all consumers have the same degree of risk aversion, and that a health insurance plan is offered at a certain premium – lets call it P_0. Then, the only consumers who will be willing to buy the plan are those with expected health care costs above some threshold (where the threshold is some increment below the premium since consumers are risk averse). Now, suppose that the expected cost to the insurance company of covering these consumers exceeds P_0. Then, the insurance company will lose money at the original premium, and it will have to raise the premium to $P_1 > P_0$. When the premium rises, the marginal consumers with relatively low expected utilization drop out, and a riskier group of consumers is left in the plan. Now suppose the expected cost of covering this smaller group of consumers exceeds P_1. The insurance company will again have to raise the premium to cover costs. Depending on the distribution of expected health care costs among consumers in the population, this vicious cycle can continue until this insurance plan (or for that matter, any possible plan) ceases to exist. However, if there is heterogeneity in risk aversion, and risk averse people are healthier, it counteracts the vicious cycle, by making healthier people willing to pay higher premiums.

Now, suppose we look across markets for different types of insurance, and ask which markets are small and which markets are large. Well-known examples of very small insurance markets (i.e., almost no one has coverage) are the market for long term care (LTC) insurance (i.e., insurance to cover the costs of nursing home stays) and the market for annuities (which is essentially insurance against living too long so that one runs out of savings before death). In these markets, "adverse selection" arises if healthier people buy. That is, if you are an insurance company, you want to sell annuities and LTC insurance to sick people who are likely to die fast, because such people won't make much use of the product and will therefore be profitable. Thus, the "bad risks" in these markets are healthy, long lived people.

Now, if risk averse people indeed tend to be healthier, it will reinforce the adverse selection problem in LTC and annuity markets. That is, risk aversion generates another reason that the "bad risks" will want to buy insurance. Plausibly, this tendency of the positive risk aversion-health correlation to reinforce, rather than counteract, the adverse selection on health mechanism explains why these markets are so small.

In contrast, consider the market for auto insurance. This market is like the market for health insurance, in that (i) the "bad risks" are the poor and/or reckless drivers, who are analogous to unhealthy people, and (ii) with population heterogeneity in risk aversion, the good risks, who in this case are the good/safe drivers, will tend to buy more insurance, thus counteracting adverse selection. The same argument applies to life and homeowners insurance.

Thus, it appears that the felicitous selection story may be able to explain why markets for auto, health, life and homeowners insurance are big (i.e., many or most people buy them), while markets for LTC insurance and annuities are very small. This hypothesis is in principle testable, because economists have developed methods to measure risk aversion (e.g., by asking people

whether they would accept particular gambles). However, the felicitous selection hypothesis cannot be tested using existing data, because data sets containing the information we need to model insurance choice do not contain risk aversion measures. The sort of SP choice experiments I described in section III could be used to address this problem, because it would not be difficult to append questions about risk aversion to survey instruments used to collect SP data.

If felicitous selection is important, it has important implications for modeling insurance choices. As an example, suppose the population of consumers looks like the following:

Risk Aversion / Health	Low	Medium	High
Low	20% (L)		
Medium		20% (M)	20% (H)
High		20% (L)	20% (M)

In this situation, there is a 20% segment of the population that is rather indifferent to risk and that also has poor health (because they do not take good care of their health). This segment will demand low levels of insurance (L), and will also have high utilization of services. Among the other 80% of the population, risk aversion and health status are independent. That is, a medium risk aversion person is just as likely to be in good health as a high risk aversion person. Thus, among the 80% of the population with medium to high risk aversion, adverse selection will be operative: people with medium health will demand more insurance and have higher utilization, on average, than people with good health. On the other hand, if we look at the population as a whole, those with low demand for insurance (L) are, on average, less healthy than those with medium demand (M).

If insurance companies cannot identify members of the risk-indifferent segment, this leads to important changes in the sorts of adverse selection equilibrium modeled by Rothschild and Stiglitz (1976), Wilson (1977) and Spence (1978). As De Meza and Webb (2001) point out, felicitous selection can lead to an equilibrium where the marginal insurance buyer has expected utilization costs that exceed his/her premium, and where the highest expected utilization consumers are uninsured. Thus, there is a positive correlation between insurance coverage and health status, in contrast to the negative correlation that arises in adverse selection models.[32]

Chiappori and Salanie (2000) have argued that, if consumers have information about their expected utilization that insurers don't have, there should be a positive correlation between insurance coverage and utilization. Thus, they take their finding that uninsured drivers have more accidents than insured drivers as evidence against asymmetric information in the auto insurance market. But, as my simple example illustrates, and as De Meza and Webb (2001) point out more generally, such a positive correlation need not hold in a market with felicitous selection.[33]

---

[32] De Meza and Webb show that, in a felicitous selection equilibrium, a welfare improvement can be achieved by taxing insurance purchases, and making lump sum distributions to the population. This gets the risk indifferent off insurance (which they were buying too much of), lowers premiums, and allows the risk averse to buy insurance at a lower net cost. This is very different from the welfare enhancing policies discussed by Wilson and Spence for the adverse selection case, which involve the healthy subsidizing insurance for the unhealthy. De Meza and Webb also show that felicitous selection can induce another important change in the nature of equilibrium. Rothschild and Stiglitz showed that a Nash equilibrium with adverse selection, if it exists, is always a "separating" equilibrium: the unhealthy have comprehensive coverage at a high premium, and the unhealthy have limited coverage at a much lower premium. Any "pooling" equilibrium where the same policy is offered to both groups would be broken, because one can always find a policy with lower premiums and less coverage that would attract the healthy consumers away from the population risk pool. With felicitous selection, however, the healthy don't want less coverage than they would have in a pooling equilibrium, so such an equilibrium may be sustainable. For completeness, it is worth noting that Wilson (1977) showed that a pooling equilibrium is sustainable in an adverse selection model if the equilibrium concept is modified from Nash to "Wilson foresight," a concept which rules out any deviation that would subsequently be rendered unprofitable by other firms' reactions.

[33] Chiappori, Jullien, Salanié and Salanié (2002) argue that in a competitive market, if two insurance plans coexist, then the more comprehensive plan must have riskier (i.e., higher utilization) enrollees. The reason is that, if enrollees of both plans had equal rates of utilization, then, given the zero profit condition, the higher premium of the more comprehensive plan would exactly cover its higher payouts. It's premium couldn't be higher than that, or its profit would be positive, inducing entry. If risk averse agents can choose between two insurance plans, and choosing the more expensive plan only involves an actuarially fair increase in premium, they will always choose the more comprehensive plan (since, by doing so, they hold their expected consumption fixed while lowering its variance). However, if the enrollees of the more comprehensive plan are sufficiently riskier, its premium can rise to the point

However, using choice modeling techniques, combined with data on risk aversion and health care utilization, it would be possible to estimate the size of market segments like those in my example, and examine directly the importance of adverse and felicitous selection.

IV. B. Moral Hazard and the Mantle Effect

A traditional fear about generous provision of subsidized public insurance is that it may generate moral hazard. That is, a better insured person may take less care of their health, since they know they will get free or inexpensive treatment should they get sick. Traditionally, the moral hazard problem has been analyzed in a static choice framework. But recent work by Khwaja (2003) suggests that moral hazard may be less of a problem in a dynamic model of investment in health over the life-cycle (see Grossman (1972)).

The key point is that, in a dynamic model, more generous insurance can increase your life expectancy, since a better insured person can afford to obtain better treatment should they become sick. Increased life expectancy, in turn, enhances ones incentive to invest in health. Technically, the reason is that, in any dynamic model, a longer planning horizon (i.e., in this case, life-span) increases returns to investment. More intuitively, if one expects to live longer, it gives one an incentive to make investments in health that will enhance quality of life in old age. This dynamic effect counteracts the static moral hazard effect on investment in health, pretty much completely canceling it according to Khwaja's estimates, obtained using the U.S. Health and Retirement Survey (HRS).[34]

---

where the lower risk consumers do not find it attractive. Hence, Chiappori et al argue that a positive correlation between risk and coverage is a very "general" implication of models with asymmetric information (i.e., adverse selection or moral hazard). However, De Meza and Webb note that this argument only works if there are no fixed costs of providing insurance. If the more comprehensive policy has a higher premium in part because there is a fixed cost of providing more comprehensive coverage (e.g., one has to devote more resources to underwriting), then the premium increment in going from the limited to the more comprehensive plan need not be actuarially fair, even under a zero profit condition. Also, the "positive correlation" argument clearly fails if there is imperfect competition.

[34] Khwaja operationalizes investment in health as occurring through exercise, the avoidance of risky behaviors (like excessive drinking), and through medical treatment for chronic conditions and/or preventive care. He then estimates

I like to call the positive effect of life expectancy on investment in health the "Mickey Mantle effect," after the great Yankee centerfielder of the 1950s and 60s. Mantle (1931-1995) was one of a handful of sports stars whose celebrity transcended sports to make him a cultural icon. He was legendary for having perhaps the greatest natural talent of any baseball player. But a string of serious injuries, combined with heavy drinking and lack of regard for his health, ended his career prematurely. After many years of alcoholism, leading ultimately to liver failure, Mantle finally died of cancer at the age of 63. He later explained his reckless behavior with the observation that he never expected to live past his early 40s, because many males in his father's line had died young due to Hodgkin's disease.[35] He was surprised to have lived into his 60s, and observed: "If I knew I was going to live this long, I'd have taken better care of myself."[36]

To test the hypothesis that life expectancy has a positive effect on investment in health would require data on investment in health (such as exercise, avoidance of risky behaviors and preventive care), and measures of life expectancy. The HRS does contain measures of life expectancy (Manksi (2004) discusses measurement of expectations using survey instruments), along with measures of investment, so the Mantle hypothesis could perhaps be tested using these data. However, a fundamental problem that must be addressed is reverse causality. Clearly, the rate of investment in health will feed back and affect life expectancy.

---

a dynamic model of investment using the HRS data. Simulations of the model imply that subsidized insurance has almost no effect on levels of exercise or risky behaviors. Note that this does not mean that more generous insurance doesn't raise use of health services. Khwaja's estimates do imply that more generous insurance raises utilization (i.e., demand curves do slope down). Khwaja's paper is an extremely ambitious exercise, being the only attempt to model health insurance choices, investment in health and health care utilization jointly, and all within a dynamic life-cycle framework. A limitation of his work, however, is that consumer's health insurance choice sets, which differ by location and/or employer, are not observed in the HRS. Khwaja assumes that all consumers have access to a standard set of insurance options.

[35] Mantle's father died of Hodgkin's disease at age 39, while his two uncles died at 32 and 41. Of Mantle's four sons, one died of Hodgkin's disease a year before Mantle (at age 36), while the other died of cancer in 2000. Mantle's two surviving sons are active in the Mickey Mantle Foundation, which promotes organ donations.

[36] One could view this as ex-post rationalization, but Mantle made related statements in his youth. For instance, as a rookie he told the Yankee player representative: "You don't have to talk to me about pensions. I won't be around long enough to collect one."

To clarify the nature of the problem, suppose that investments in health, life expectancy, one's general level of optimism, one's level of health and one's insurance choice are determined by the following system of equations:

(A) Investment in health = f(current health, price of investment in health, income, taste for health, optimism, life expectancy)

(B) Life expectancy = g(current health, price of health care, investment in health, environmental risk factors, genetics, optimism)

(C) Optimism = k(current health, investment in health, genetics)

(D) Current Health = h(lagged health, lagged investment in health, environmental risk factors, genetics, exogenous shocks to health)

(E) Insurance coverage = I(current health, insurance plan options, income, risk aversion, taste for health, tastes for insurance plan options)

Our interest is in estimating equation (A), in order to test if the "Mantle effect," that is, the effect of life expectancy on investment in health, is quantitatively important. I'll assume the error term in equation (A) arises because some part of the "taste for health" variable is unobserved. Then, a problem arises since, according to equation (B), investments in health affect life expectancy. Thus, life expectancy is endogenous in equation (A), because a person with a high unobserved taste for health will tend to have both a high investment in health and high life expectancy.

In this framework, one approach to consistently estimating equation (A) is to find an "instrumental variable" for life expectancy. That is, a variable that affects investment in health only through its affect on life expectancy (and not through any other channel). In the system of equations I've written here, one's genetic health endowment can play this role. Thus, as suggested by the Mantle story, one possible way to generate an instrument is to obtain data on congenital family diseases.[37] According to the above framework, measures of environmental risk

[37] One problem is that, while a family history of congenital disease would certainly reduce life expectancy (independent of any affect on investment in health), having parents with such health problems might also affect ones

factors are also potential instruments. Living in a risky environment may reduce life expectancy, but conditional on life expectancy it should not affect investment in health directly.[38]

It is also interesting to consider the role of a person's general level of optimism in the model. If optimism were a fixed characteristic of a person, it might quickly jump to mind as a potential instrument for life expectancy (i.e., a more optimistic person will have a higher life expectancy, *ceteris paribus*). However, I would not buy such an instrument because, as I write in equation (C), I suspect that optimism is affected by investment in health (e.g., exercise may have physiological effects that enhance ones general sense of well being, perhaps making one feel more optimistic). And, if optimism is reflective of ones general sense of well being (i.e., mental health), it might plausibly affect investment in health directly, as I write in (A). Also, there may simply be a positive correlation between unobserved tastes for health and optimism. As I've written the model, genetic factors could also be used as instruments for optimism in (A).

Next, note that the set of variables "price of investments in health" that appear in equation (A), and the set of variables "price of health care" that appear in equation (B), would both include aspects of a person's insurance coverage. Additionally, the price of investments in health would also include such things as prices of alcohol and tobacco, proximity to and cost of healthy food, proximity to athletic facilities, etc.. This leads to the following two observations:

First, it is clear that aspects of a person's insurance coverage would tend to be correlated with unobserved tastes for health as well. We will tend to have a selection bias whereby people with greater taste for health also have more comprehensive insurance (and hence, a lower cost of investment in health). This means that a proper estimation of equation (A) requires us to deal

---

investment in health through other channels (e.g., having fewer financial resources in youth if parents were ill). However, it can be plausibly argued that this problem is resolved by conditioning on current health status in the first equation, since this would control for prior investments in health that might have arisen due to family background.

[38] This assumes that risky environment is not endogenous in the sense that people with low tastes for health will also choose to live in a risky environment.

with this selection bias. That in turn, means that we should estimate equation (A) jointly with the choice model for insurance coverage, which I have written as equation (E).

Second, Diamond (1992) has suggested that, if the healthy are to subsidize insurance for the unhealthy, as efficiency in the adverse selection model suggests, then one could try to avoid any resultant moral hazard problems by simultaneously imposing "sin" taxes. Thus, another reason for interest in equation (A) is that we need estimates of the price effects in (A) in order to determine the level of sin taxes necessary to undo any effects of moral hazard that are induced by more generous insurance (and not counteracted by the Mantle effect).

In summary, it is interesting that the empirical literature has come to very weak and conflicting conclusions about whether asymmetric information in the form of either adverse selection or moral hazard is important in health insurance markets. This literature has attempted to infer the importance of adverse selection and moral hazard almost entirely by asking whether those with larger policies tend to have larger claims. Given such limited information, it is not surprising that attempts to determine whether adverse selection or moral hazard are important, to distinguish one from the other, or to determine if heterogeneity in risk aversion is correlated with risk type, all rely on very strong assumptions about market structure. Identification would be much easier if we had measures of risk aversion, perceived risk (i.e., self rated health status), and measures of risky behaviors. The HRS contains such measures, so it might be usable for this purpose. There is no reason one could not collect stated preference (SP) insurance choice data, and in the same instrument also collect measures of risk aversion, health, expected utilization of services, and risky behaviors. The collection of such additional data might allow one to form strong tests of the importance of adverse selection, felicitous selection, moral hazard and the Mantle effect. This is obviously an important avenue for future research.

V. Summary

The main points that I have made in this paper can be summarized as follows:

- A standard argument for "Competition" in insurance markets is that giving consumers more "Choice" will enhance welfare.

- Indeed, there is clear empirical evidence, such as Harris and Keane (1999), of substantially heterogeneity in consumer tastes for attributes of insurance plans, so there is scope to enhance welfare by giving consumers more choices.

- If private firms design the menu of insurance options, the goal will be "cherry picking," which means more "Choice" will not necessarily increase consumer welfare.

- Furthermore, consumers seem to have important misperceptions about insurance options. This also undermines the "choice is good" argument.

- Thus, the assumptions underlying the standard argument for private competition do not hold.

- Given the problems with private competition, the best way to satisfy heterogeneity in consumer tastes for health care plans is for a single payer (i.e., the government) to design a menu of health insurance plan options, and to contract with providers so it can offer the whole menu to consumers.

- The choice modeling techniques used to analyze consumer preferences and design products that appeal to various consumer segments are well known, and could be applied by government just as well as by private firms.[39] The government could use these methods to design an appealing menu of insurance options that would meet certain equity and efficiency goals.

- To implement this agenda, we need to model insurance plan choices and health care utilization conditional on insurance plan choice jointly. Unfortunately, existing data sets permit one to model either insurance plan choice or health care utilization, but do not contain enough information to do both. Thus, we need to collect more data.

[39] It does not appear to me that private insurance firms are doing a great deal of innovating in terms of coming with genuinely new insurance product attributes. By this time, the attributes of insurance plans seem to be fairly standard, and designing new insurance options largely comes down to mixing and matching those attributes. For example, the latest "hot idea" in health insurance seems to be the "consumer driven health plan." But this basically just amounts to combining catastrophic insurance with a large deductible with a health care spending account. All these individual features are well known. In contrast, the superiority of private firms at product innovation is the key feature of differentiated consumer goods markets that would invalidate any argument for government design of choice options in that context.

A number of people have asked me why I would advocate substantial government intervention in the market for health insurance, when I am generally skeptical of government intervention in other markets. For instance, why would I argue that the government should provide health insurance, but not, say, auto accident insurance? I am going to argue that health care is "different," but it is important to be careful about how. The health economics literature can be confusing on this point, because it is filled with discussions of why the market for health insurance is "different" because problems of asymmetric information, adverse selection, agency and moral hazard lead to "market failure" that prevents an efficient competitive equilibrium outcome from being achieved. But many other markets (especially other insurance markets) suffer from similar problems. The really key point is that in health insurance, unlike, say, auto insurance, even if all these potential sources of market failure could be magically made to vanish, and a fully efficient competitive equilibrium could be achieved, just about everyone would agree that the efficient outcome is <u>not</u> desirable.

In an efficiently functioning competitive market without informational problems, each individual would pay premiums equal to their expected cost of health care utilization. But, on equity grounds, we don't think that intrinsically less healthy people should have to pay more for health insurance just because they are unlucky enough to have poor health. This is in contrast to auto insurance markets, where most people would say it is desirable for less safe drivers to pay higher premiums.

Because of equity considerations, there obviously <u>is</u> going to be government intervention in the health insurance market, whether it takes the form of a single payer system (with or without any scope for consumer choice), or a system that involves private insurance firms offering competing plans, combined with risk adjusted government subsidies. In light of this, I

feel that someone like myself, who distrusts the ability of government to do sophisticated central planning, should gravitate toward a system of intervention that involves the simplest computational problem for the government. I have argued that this is a single payer system, since introduction of private firms along with risk adjusted subsidies substantially increases the computations that must be undertaken by government.

While I have heavily emphasized the value of single payer system in overcoming the adverse selection problem, it should be noted that there are other benefits to such a system. These are well illustrated by the U.S. Medicare program. The large size of the program enables it to have very low administrative overhead relative to private insurers, and during the second half of the 1990s, when Congress permitted, it was very successful at using its monopsony power to negotiate low payment rates with hospitals and physicians (through its prospective payment system and physician fee schedule - see Berenson (2001) and Foster (2000) for discussions). Indeed, Medicare fee-for-service cost increases we so slow in the late 1990s that many private Medicare HMOs were driven out of business, despite the fact that their capitation payments remained, by most estimates, above what their enrollees would have cost under Medicare fee-for-service. This experience can't be encouraging for the idea that Medicare HMOs can achieve cost savings on Basic Medicare. Yet, under the Medicare Modernization Act of 2003, the government started pumping more money into Medicare HMOs to keep them afloat.

The original idea behind HMOs was that they could deliver health care more efficiently by organizing providers into competing groups, thus driving down provider prices. As discussed by Nichols et al (2004), this idea has floundered because consumers are so attached to provider choice. The strength of this preference was revealed by the estimates I reported in Section II. Providers have been able to exploit this to gain market power. Instead of HMOs threatening

providers with loss of patients if they are unwilling to accept discounted fees, we have provider groups able to "dictate terms to health plans on the premise that their absence from a network would make [it] unattractive to consumers."[40] In contrast, recent history shows that a large single payer like Medicare does have the countervailing power to dictate terms to providers.

Recently, Enthoven (2004) has argued that to make the managed care idea work we need to use antitrust laws to break up provider monopolies and come up with accurate methods of risk adjustment. Whether breaking up provider networks would enhance welfare seems unclear, given the clear consumer preference for large networks. And to admit that managed care relies on accurate risk adjustment seems tantamount to admitting it is impossible (see Newhouse (1998) for further critical comments on risk adjustment technology).

I have emphasized the problems with Medicare HMOs, but it is worth stressing that allowing private firms to offer medigap plans creates problems as well. Most notably, by covering the deductibles and co-pays that exist under Basic Medicare, medigap plans increase the use of services covered by Basic Medicare, thus undermining efforts to control utilization by having consumers bear some fraction of costs - see Christensen and Shinogle (1997). Again, the government's problem is simpler under a single payer system, because it can implement cost sharing without worrying about how private insurers' actions might undo its efforts.[41]

---

[40] On this point, it is interesting to look back at the classic article by Stockman (1983). He said "The fourth premise on which any kind of plan ought to be centered is the notion of healthy provider competition and marketing of health care plans on a retail basis …. Once we establish a retail market among the consumers, *we will automatically and perforce get fierce competition among various provider units*." (emphasis added).

[41] Recently, a number of analysts have suggested reforming Medicare to make it look more like the U.S. Federal Employees Health Benefits Plan (FEHBP). For instance, see Feldman, Dowd and Coulam (1999). The basic idea is to run Medicare like employer subsidized health insurance: the government pays some share of premiums, and negotiates with a set of private insurance plans that agree to offer their products to senior citizens during an annual open enrollment period, using community rated premiums. The basic problem with this idea is that it does nothing to alleviate the adverse selection problem, unless plans are somehow given risk adjusted contributions. But this again raises the question of whether it is possible to do risk adjustment well. Feldman, Dowd and Maciejewski (2001) point out that larger employer subsidies towards the premiums of more comprehensive (higher cost) plans "can accomplish some of the benefits of formal risk adjustment," but I am skeptical that cherry picking problems would not remain severe. See Moon and Davis (1995) and Merlis (1999) for critical analysis of the FEHBP model idea.

My advocacy of a single-payer menu-of-options approach to organizing the health insurance market is certainly not new. Indeed, Wilson (1977), Spence (1978) and Diamond (1992) all noted the appeal of such an idea long ago.[42] But Spence also noted that the key challenge in implementing this idea is how to design the menu, since this would require a great deal of information about consumer tastes, the share of each type of consumer in the population, etc.. The key point of this lecture was to describe new techniques in choice modeling that might actually make this exercise feasible.

Specifically, the state-of-the art choice modeling techniques that I have described here provide a technology that enables us, at least in principle, to design menu of insurance options that would appeal to consumer tastes and, at the same time, enhance equity and efficiency. As I've pointed out, the main obstacle to pursuing this important agenda is the available of data that would enable us to model health plan choices and service utilization simultaneously. To do this, we will need to combine data from a number of sources, including data on consumer insurance choice in actual markets, stated preference and attitudinal data, and data of health status and medical service utilization. Harris and Keane (1999), as well as some subsequent papers adopting similar approaches, point the way toward how this could be done.

In his conclusion, written over 25 years ago, Spence (1978) called for exactly the sort of research of agenda that I am proposing. He said "Publicly provided insurance can improve on the private market. … Neither goal, improving efficiency, or redistributing benefits, is inconsistent with maintaining a reasonable array of consumer options. It might be objected that the informational problems make it difficult to calculate exactly what the second best menu would

[42] As an applied econometrician who is not an expert on the theory of insurance markets, I was gratified to discover that such an eminent group of theorists had reached similar conclusions to my own. It is worth noting that Spence (1978) suggested that the menu-of-options idea could be implemented either by government regulation of insurance companies (i.e., requiring insurance companies to offer a well designed menu of options, and requiring any company to offer the entire menu), or by having the government offer the same menu itself.

look like. That is certainly true. But that hardly seems a reason to ignore the problem… by pretending that individuals … are … sufficiently similar to make a differentiated menu unnecessary. That judgment should be empirically based. Perhaps the easiest way to make it is to offer a portfolio of options and observe the choices that are made." Well, we have already learned that there is substantially heterogeneity in consumer tastes for health insurance plan features, so design of a differentiated menu should be the policy goal. But we have a lot of work left to do before we can determine what the menu should look like.

Finally, while I have advocated a single-payer menu-of-options approach, and the potential usefulness of choice modeling in designing such a system, it is also true that choice modeling techniques can and should play an important role even in more modest reform proposals. As I've noted, it seems extraordinary that no market research was done to help design the basic options that were to be offered under the Clinton health care plan, or to try to ascertain the population distribution of willingness to pay for prescription drug benefits before passage of the Medicare Modernization Act of 2003. Given that the technology to learn about consumer preferences is available, it seems inexcusable that government fails to exploit this technology to help design better public health insurance plans.

References

Akerlof, George (1970). The market for lemons: Qualitative uncertainty and market mechanism. Quarterly Journal of Economics, Vol. 84, No. 3 (August), p. 488-500.

Berenson, Robert A. (2001). Medicare+Choice: Doubling or disappearing? Health Affairs, Web exclusive, November 28.

Blumberg, Linda J., Len M. Nichols and Jessica S. Banthin (2002). Worker decisions to purchase health insurance. International Journal of Health Care Finance and Economics, Vol. 1, No. 3-4 (Sept.-Dec.), p. 305-325.

Brown, Randall S., Dolores G. Clement, Jerrold W. Hill, Sheldon M. Retchin and Jeanette W. Bergeron (1993). Do Health maintenance Organizations Work for Medicare? Health Care Financing Review, 15:1 (Fall), p. 7-23.

Cafferata, Gail L. (1984). Knowledge of their health insurance coverage by the elderly. Medical Care, Vol. 22, No. 9 (September), p. 835-847.

Cardon, James H. and Igal Hendel (2001). Asymmetric information in health insurance: evidence from the National Medical Expenditure Survey. Rand Journal of Economics, Vol. 32, No. 3 (Fall), p. 408-427.

Chiappori, Pierre-André, and Berhard Salanié (2000). Testing for asymmetric information in insurance markets. Journal of Political Economy, Vol. 108, No. 1 (February), p. 56-78.

Chiappori, Pierre-André, Bruno Jullien, Berhard Salanié and Franćois Salanié (2002). Asymmeric Information: General Testable Implications. Working paper, University of Chicago.

Christensen, Sandra and Judy Shinogle (1997). Effects of supplemental coverage on use of services by Medicare enrollees. Health Care Financing Review, 19:1 (Fall), p. 5-17.

Cunningham, Peter J., Charles Denk and Michael Sinclair (2001). Do consumers know how their health plan works? Health Affairs, Vol. 20, No. 2 (March/April), p. 159-166.

Cutler, David M. and Richard J. Zeckhauser (2000). The Anatomy of Health Insurance. In A.J. Cuyler and J.P. Newhouse (eds.), Handbook of Health Economics, Vol. 1A, Elsevier Science: Amsterdam, p. 563-643.

Davidson, Bruce, Shoshanna Sofaer and Paul Gertler (1992). Consumer demand and biased selection in the demand for coverage supplementing Medicare. Social Science and Medicine, Vol. 34 No. 9 (May), p. 1023-1034

De Meza, David and David C. Webb (2001). Advantageous selection in insurance markets. The RAND Journal of Economics, Vol. 32, No. 2 (Summer), p. 249-262.

Diamond, Peter (1992). Organizing the Health Insurance Market. Econometrica, Vol. 60, No. 6 (November), p. 1233-1254.

Elrod, Terry, and Michael P. Keane (1995). A factor analytic probit model for representing the market structure in panel data. Journal of Marketing Research, Vol. 32, No. 1 (February) p. 1-16.

Enthoven, Alain C. (2004). Market forces and efficient health care systems. Health Affairs, Vol. 23, No. 2 (March/April), p. 25-27.

Feldman, Roger, Bryan E. Dowd and Matthew Maciejewski (2001). A demand side view of risk adjustment. Inquiry, Vol. 38, No. 3 (Fall), p. 280-289.

Feldman, Roger, Bryan E. Dowd and Robert Coulam (1999). The Federal Employees Health Benefits Plan: Implications for Medicare reform. Inquiry, Vol. 36, No. 2 (Summer), p. 188-199.

Foster, Richard S. (2000). Trends in Medicare expenditures and financial status, 1966-2000. Health Care Financing Review, Vol. 22, No. 1 (Fall), p. 35-49.

Frank, Richard G., Jacob Glazer and Thomas G. McGuire (2000). Measuring adverse selection in managed health care, Journal of Health Economics, Vol. 19, No. ** (****), p. 829-854.

Gibbs, Deborah A., Judith A. Sangl and Barri Burrus (1996). Consumer perspectives on information needs for health plan choice. Health Care Financing Review, Vol. 18, No. 1 (Fall), p. 55-73.

Glazer, Jacob and Thomas G. McGuire (2000). Optimal risk adjustment in markets with adverse selection: An application to managed care. American Economic Review, Vol. 90, No. 4 (September), p. 1055-1071.

Glied, Sherry (2000). Managed Care. In A.J. Cuyler and J.P. Newhouse (eds.), Handbook of Health Economics, Vol. 1A, Elsevier Science: Amsterdam, p. 707-753.

Greenwald, Leslie M., Jesse M. Levy and Melvin J. Ingber (2000). Favorable selection in the Medicare+Choice program: New evidence, Health Care Financing Review, Vol. 21, No. 3 (Spring), p. 127-134.

Grossman, Michael (1972). On the concept of health capital and the demand for health. Journal of Political Economy, Vol 80, No.2 (March/April), p 223-255.

Hall, Jane (2004). Can we design a market for competitive health insurance? CHERE discussion paper #53 (February).

Harris, Katherine M. (2002). Can high quality overcome consumer resistance to restricted provider access? Evidence from a health plan choice experiment. Health Services Research, Vol. 37, No. 3 (June), p. 551-71.

Harris, Katherine M. and Michael P. Keane (1999). A model of health plan choice: Inferring preferences and perceptions from a combination of revealed preference and attitudinal data. Journal of Econometrics, Vol. 89, No. 1-2, p. 131-157.

Harris, Katherine M., Roger Feldman and Jennifer Schultz (2002). Measuring consumer perceptions of quality differences among competing health benefit plans. Journal of Health Economics, Vol. 21, No. 1, p. 1-18.

Hensher, David A., Jordan L. Louviere and Joffre Swait (1999). Combining sources of preference data. Journal of Econometrics, Vol. 89, No. 1-2, p. 197-221.

Isaacs, Stephen L. (1996). Consumer's information needs: results of national survey. Health Affairs, Vol. 15, No. 4 (Winter), p. 31-41.

Khwaja, Ahmed (2003). Health insurance, habits and health outcomes: moral hazard in a stochastic life cycle model of investment in health. Working paper, Fuqua School of Business, Duke University.

Hall, Jane, Rosalie Viney, Marion Haas and Jordan Louviere (2004). Using stated preference discrete choice modeling to evaluate healthcare programs. Journal of Business Research, forthcoming.

Louviere, Jordan J, David A. Hensher and Joffre D. Swait (2000). Stated Choice Methods: Analysis and Application. Cambridge University Press.

Manski, Charles F. (2004). Measuring expectations. Econometrica, forthcoming.

McCall, Nelda, Thomas Rice and Judith Sangl (1986). Consumer knowledge of Medicare and supplemental health insurance benefits. Health Services Research, Vol. 20, No. 6 (February), p. 633-657.

Merlis, Mark (1999). Medicare restructuring: The FEHBP model. Report prepared for the Kaiser Family Foundation, www.kff.org/medicare/upload/14690_1.pdf.

Miller, R. H. and H. S. Luft (1997). Does managed care lead to better or worse quality of care?, Health Affairs, Vol. 16, No. 5 (Sept./Oct.), p. 7-25.

Moon, Marilyn and Karen Davis (1995). Preserving and strengthening Medicare. Health Affairs, Vol. 14, No. 4 (****), p. 31-46.

Newhouse, Joseph P. (1998). Risk adjustment: Where are we now? Inquiry, Vol. 35, No. 2 (Summer), p. 122-131.

Nichols, Len M., Paul B. Ginsburg, Robert A. Berenson, Jon Christianson and Robert E. Hurley (2004). Are market forces strong enough to deliver efficient health care systems? Confidence is waning. Health Affairs, Vol. 23, No. 2 (March/April), p. 8-21.

Parente, Stephen T., Roger Feldman and Jon B. Christianson (2004). Employee choice of consumer-driven health insurance in a multiplan, multiproduct setting. Health Services Research, Vol. 39, No. 4, Part II (August), p. 1091-1111.

Rothschild, Michael and Joseph Stiglitz (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information, Quarterly Journal of Economics, Vol. 90, No. 4 (November), p. 629-649.

Shen, Yujing and Randall P. Ellis (2002). How profitable is risk selection? A comparison of four risk adjustment models. Health Economics, Vol. 11, No. 2 (March), p. 165-174.

Spence, Michael (1978). Product differentiation and performance in insurance markets. Journal of Public Economics, Vol. 10, No. 3 (December), p. 427-447.

Stockman, David A. (1981). Premises for a medical marketplace: A neoconservative's vision of how to transform the health system. Health Affairs, Vol. 1, No. 1 (Winter), p. 5-18.

Tumlinson, Anne, Hannah Bottigheimer, Peter Mahoney, Elliot Stone and Ann Hendricks (1997). Choosing a health plan: What information will consumers use? Health Affairs, Vol. 16, No. 3 (May/June), p. 229-238.

U.S. General Accounting Office (2000). Medicare+Choice: Payments Exceed Cost of Fee-for-Service Benfits, Adding Billions to Spending. GAO/HEHS-00-161, U.S. Government Printing Office, Washington, D.C..

Van de Ven, Wynand P.M.M and Randall P. Ellis (2000). Risk adjustment in competitive health plan markets. In A.J. Cuyler and J.P. Newhouse (eds.), Handbook of Health Economics, Vol. 1A, Elsevier Science: Amsterdam, p. 707-753.

Vistnes, Jessica P. and Jessica S. Banthin (1997). The demand for Medicare supplemental insurance benefits: The role of attitudes toward medical care and risk. Inquiry, Vol. 34, No. 4 (Winter), p. 311-324.

Wilson, Charles (1977). A model of insurance markets with incomplete information. Journal of Economic Theory, Vol. 16, No. 2 (December), p. 167-207.