International Journal of Technology Assessment in Health Care, **17:1** (2001), 27–37.

ESTIMATING THE BAYESIAN LOSS FUNCTION

A Conjoint Analysis Approach

Mohan V. Bala
Centocor, Inc.

Josephine Mauskopf
Research Triangle Institute

Abstract

Current health economic literature does not provide clear guidelines on how uncertainty around cost-effectiveness estimates should be incorporated into economic decision models. Bayesian analysis is a promising alternative to classical statistics for incorporating uncertainty in economic analysis. Estimating a loss function that relates outcomes to societal welfare is a key component of Bayesian decision analysis. Health economists commonly compute the loss function based on the quality-adjusted life-years associated with each outcome. However, if welfare economics is adopted as the theoretical foundation of the analysis, a loss function based in cost-benefit analysis (CBA) may be more appropriate. CBA has not found wide use in health economics due to practical issues associated with estimating such a loss function. In this paper, we present a method based in conjoint analysis for estimating the CBA loss function that can be applied in practice. We illustrate the use of the methodology using data from a pilot study.

Keywords: Bayesian analysis, Loss function, Conjoint analysis

The objective of health economic analysis is to allocate societal resources to health interventions in order to maximize societal well-being. Cost-effectiveness analysis (CEA) is the most common analytical framework used to compare treatment options for any given disease and to choose the one which will achieve the stated goal of maximizing societal welfare. Even though alternate analytical frameworks such as cost-benefit analysis (CBA) exist, they have not found wide use in health economics.

Economic analyses of health interventions need estimates of both the efficacy and the cost of these interventions in the chosen disease populations. These estimates are commonly obtained from studies conducted in a sample of patients with the given disease. Hence, there is uncertainty surrounding the cost and efficacy estimates derived from these studies. One important question in health economic analysis is how this uncertainty can be appropriately incorporated in generating resource allocation decisions.

Initial efforts to account for uncertainty in CEA have been based in classical statistics, including generation of confidence intervals, and hypothesis testing. However, these analyses cannot be easily incorporated into a decision model for optimal allocation of resources. Bayesian analysis is a promising alternative to the classical statistics approach

for incorporating uncertainty into economic decision models. Even though few published economic analyses have used the Bayesian methodology, this is a topic that is gaining substantial interest in health economics (4;6;23).

One of the key components of the Bayesian decision model is the loss function. Historically, in decision analysis the term *loss function* has been used to refer to the loss associated with any combination of action and outcome, and this loss is usually expressed in monetary terms (16). In such analyses a gain is represented as a negative loss. From an economic analysis perspective, the loss function can be used to represent the loss in societal welfare for any given combination of intervention and health outcome. Traditionally, health economic analyses have used the CEA framework, where health outcomes have been quantified in terms of quality-adjusted life-years (QALYs). CEA implicitly uses a loss function where the contribution of the intervention to the societal welfare is expressed as the difference between the incremental QALY gain associated with intervention (converted to monetary terms using the societal value per QALY) and the incremental cost of the outcome (20). However, the CEA approach has several drawbacks from a theoretical standpoint, and a CBA approach may be preferable if welfare economics is adopted as the theoretical foundation for the analysis. The objective of this paper is to examine the potential use of conjoint analysis to estimate the loss function used in CBA.

In this paper, we will first provide a brief overview of the use of classical statistics to account for uncertainty in health economic analysis. We will follow this with a brief overview of the Bayesian approach to incorporating uncertainty into economic analysis. We will then introduce conjoint analysis and describe its theoretical foundation and practical use. This will be followed by a comparison of the QALY-based and the conjoint-based approaches to estimating the loss function. We will then present a pilot study where we used conjoint analysis to conduct a health economic analysis of a new therapy for acute myocardial infarction (AMI). We will conclude with a brief summary of the key findings and discuss how conjoint analysis can be useful even if one rejects the welfare economic foundation of health economic analysis.

ANALYSIS OF UNCERTAINTY IN HEALTH ECONOMICS: THE CLASSICAL STATISTICS APPROACH

As stated earlier, the most common analytical approach used in health economics is CEA. CEA attempts to maximize societal welfare by ensuring that society pays no more than $V for every unit of additional health benefit. The health benefit is usually measured in terms of QALYs. The value V (usually $40,000–$50,000) is derived based on what society is willing to pay or has traditionally paid for a QALY. In practical application of CEA, the incremental cost-effectiveness ratio (ICER), defined as the ratio of the incremental cost to the incremental QALYs, is used to compare alternative therapies. ICER is compared to V to determine whether the gain in effectiveness is worth the additional expense. The welfare maximizing intervention in the CEA approach is the one with the highest ICER (compared to the next best alternative) below V, after the dominated interventions have been eliminated (18).

Most economic analyses present a point estimate of the ICER based on the mean values of cost and effectiveness. However, the uncertainty surrounding the point estimates of cost and effectiveness results in uncertainty around the point estimate of the ICER. Several different options based in classical statistics have been suggested to incorporate this uncertainty into health economic analysis.

One approach for incorporating uncertainty that has been proposed is to derive the confidence interval around the ICER. Bala and Mauskopf (2) present a review of the methods based in classical statistics that can be used to compute the confidence interval. Even though the panel on cost-effectiveness in health and medicine recommends that health economic

studies report confidence intervals around the ICER (22), there has been little investigation of how the confidence interval should be used in the optimal allocation of resources.

Another approach for incorporating uncertainty that has been proposed is to test hypotheses regarding the sign and magnitude of the ICER. In this case, the decision to adopt or reject a new therapy would be made based on whether the *p* value achieves a certain level of statistical significance, such as .05 (19). However, from the perspective of an economic decision model whose objective is to maximize societal welfare, it is unclear why .05 level of significance should be preferred over some other value (such as .1 or .025). Thus, the approaches based in classical statistics do not clearly indicate how to incorporate uncertainty into economic decision models.

ANALYSIS OF UNCERTAINTY IN HEALTH ECONOMICS: THE BAYESIAN APPROACH

Bayesian analysis provides an alternative framework to classical statistics in analyzing uncertainty. As opposed to the confidence intervals generated by classical statistics, Bayesian analysis generates probability distributions around the estimates of cost and efficacy. The Bayesian approach requires the specification of a prior probability distribution for the parameters of interest. This is the best estimate based on available data before the trial of interest has been conducted. The prior distribution is then combined with the trial data to generate a posterior distribution for the parameter. The methodology for computing the posterior distribution for different prior distributions is described in detail in the literature (10). In most cases large simulations are required to compute these distributions; however, the availability of computing power has made this task feasible.

As stated before, the objective of health economic analysis is to choose an intervention a from a possible set of health interventions $\mathbf{A}(a \in \mathbf{A})$ for a given disease. The consequence of the choice will depend on the "state of nature" $\theta(\theta \in \Theta)$. In the healthcare context the state of nature corresponds to how the patient responds to the chosen intervention. For each combination of a and θ, we define a loss function L(a, θ). From a health economics perspective, the loss function represents the change in societal welfare associated with (a, θ). Implicitly L(a, θ) is expressed relative to a baseline (a_0, θ_0) where L(a_0, θ_0) is set equal to zero.

In the absence of any information regarding the relative likelihoods of the different states of nature θ, criteria such as Minimax, Maximin, or Laplace criteria can be applied to choose the optimal intervention (16). However, if a prior probability distribution of θ is available, this in combination with experimental data can be used to compute the posterior distribution of θ. The posterior distribution can in turn be used to compute an expected loss, E[L(θ, a)], as well as an acceptance curve (4;23). Claxton argues that:

> If the objective is to maximize health gain for a given budget, then programmes should be selected based on the (posterior) mean net benefit irrespective of whether any differences are regarded as statistically significant or fall outside a Bayesian range of equivalence (6).

He illustrates how rejecting an intervention based on the expected net benefit not achieving statistical significance can impose unnecessary costs on society. It is the approach proposed by Claxton that we adopt in this paper. Hence, if L(θ, a) is the net benefit associated with intervention a and outcome θ, the optimal intervention is the one that minimizes E[L(θ, a)]. Thus, the focus of this paper will be on computing the expected loss.

The QALY loss associated with outcome θ can be used to compute the loss L(a, θ). However, this approach has some drawbacks, and we propose conjoint analysis as an alternative to the QALY approach to estimate the loss function. In the next section we

https://doi.org/10.1017/S0266462301104046 Published online by Cambridge University Press

present a brief introduction to conjoint analysis. We will follow this with a discussion of the relative merits of the QALY-based and conjoint-based approaches to estimating the loss function.

CONJOINT ANALYSIS: AN OVERVIEW

Conjoint analysis is a technique that can be used to relate the characteristics of a product or service to an individual's preference for that product. This technique has been studied extensively in the marketing science literature to quantify preferences (5;12;13;14). It has also been used in the design of new products (27).

In conjoint analysis, the product (in our case the health intervention) that is being evaluated is assumed to have several characteristics that could potentially affect an individual's preference for that product. In conjoint analysis terminology these characteristics are referred to as attributes. For a health intervention, potential attributes could be efficacy, safety (note that efficacy and safety could be themselves split into several different attributes), cost, dosing frequency, duration of therapy, mode of administration, etc. Each attribute can take on a range of values. In conjoint analysis, the attributes are usually required to take on one of a discrete set of values. For instance, the per diem cost of the intervention could be $1, $5, or $10. The potential values that an attribute could take on are referred to as levels of the attribute.

Conjoint analysis is based on principles of multi-attribute utility assessment, and the goal of the analysis is to estimate individual utility as a function of the different attribute levels. Let Y_{ij} be a binary variable where Y_{ij} is 1 if the given product has level j of attribute i, and zero otherwise. For a given product Y_{ij} satisfies the condition

$$\Sigma_j Y_{ij} = 1, \quad \forall i \tag{1}$$

Equation 1 indicates that each attribute takes on one, and only one, level in a given product. The objective of conjoint analysis is to estimate the utility of a product as a function of the attribute levels (25;27). This function is often expressed as:

$$U = \Sigma_i \Sigma_j \lambda_{ij} \times Y_{ij} \tag{2}$$

where λ_{ij} quantifies the contribution to product utility of level j of attribute i.

The steps in conducting a conjoint analysis have been described in detail in the literature (25). However, we provide a brief summary of these steps here. The process of conducting a conjoint analysis study can be divided into five discrete steps: a) define product attributes and levels; b) design conjoint analysis survey instrument; c) draw sample and administer survey; d) estimate utility/loss function; and e) interpret and use results.

Defining product attributes and levels is probably the most important step in conducting a conjoint analysis study. As stated earlier, the product attributes are used to describe the different characteristics of the health benefit that we are interested in valuing. One of the challenging tasks is to select a narrow set of characteristics that play a critical role in the decision-making process. This is usually done through in-depth interviews and focus groups. As mentioned before, the levels of each attribute are values that the attribute can take on in the conjoint survey. The levels are chosen to cover the range of products that we were interested in examining.

The conjoint survey is based on experimental design. The number of all possible product profiles that can be generated (referred to as a full factorial design) is large even for a small number of attributes and levels. For instance, the full factorial design for a study with six attributes at three levels each consists of 3^6 or 729 different product profiles. It is often

difficult to evaluate all these product profiles in a conjoint study. Hence, a carefully chosen subset of these profiles (a fractional factorial design) is usually used in conjoint studies. A commonly used fractional factorial design is one where the attribute levels are uncorrelated (or orthogonal) (25). This allows the estimation of the effect of each coefficient independent of others. However, such a design assumes independence of attributes. Independence of attributes in a health intervention study would, for instance, imply that dosing frequency (e.g., once vs. twice a day) affects product utility to the same extent whether the mode of administration is oral or by injection, which may not necessarily be true in practice. If specific attribute interactions need to be estimated, the fractional factorial design can be modified to achieve this.

Various types of questions, such as full profile questions, paired comparison questions, or choice questions can be used to elicit responses in a conjoint survey. In the full profile approach, the profile of a product with each attribute fixed at one if its prespecified levels is shown to the respondent. The respondent is asked to then either rank the profiles in order of preference or to rate the profile on a preference scale. In the paired comparison approach, two full or partial (where only a subset of product attributes are shown) profiles are presented side by side. The respondent is then asked to indicate which product they prefer as well as the strength of preference. In the choice-based approach, the respondent is presented with a set of (usually three or four) product profiles. The respondent is then asked to choose the product that he or she prefers most. A review of the relative merits of these techniques is beyond the scope of this paper.

Once the survey design is complete, the next step is to draw a sample and administer the survey. At this juncture we need to decide whose utility function we would like to elicit. If the analysis is based on welfare economics, a representative sample of the general population should be surveyed (11). However, if an extra-welfarist perspective is adopted, the survey could be administered to decision makers (7;28). We will examine the distinction between welfarist and nonwelfarist approaches in more detail in the Discussion section.

Once the survey data are available, an ordinary least-squares (OLS) regression model is usually used to estimate the utility function given by equation 2. The utility function can be estimated at the aggregate level by combining all the responses, or the responses can be used to estimate a utility function for each respondent. If there is substantial heterogeneity in preferences, the aggregate utility function can mask these differences, leading to erroneous interpretation of the results. We will return to this issue in the next two sections.

Instead of the OLS model, random utility models can be used to relate individual utilities to choices or ratings (15). In the random utility model, an individual's utility function is assumed to have a deterministic component and a random component. The deterministic component of the utility can be estimated based on the responses to the conjoint survey. However, quantifying the deterministic component of the utility of two product profiles for an individual will not allow us to predict his or her choice with certainty due to the presence of the random component. The random utility model only allows us to predict the probability that the individual will choose one product over the other. Based in random utility theory, discrete choice models can be used to model choices, and ordered probit model can be used to model ratings data. However, not much compelling evidence exists to dictate the choice of one estimation model over others (25).

Once the utility function is estimated, the results need to be interpreted and used to answer the questions posed at the start of the study. These data are commonly used to design new products or predict adoption of a new market entry. However, in our case the objective is to use the data within the framework of an economic model to inform decisions regarding allocation of healthcare resources. We discuss how this can be achieved in the next section.

ESTIMATING THE LOSS FUNCTION: THE QALY APPROACH VERSUS THE CONJOINT APPROACH

We want the loss function L(a, θ) to have the property that the decision with the lowest expected loss maximizes societal welfare. How a social welfare function should be derived from individual preferences for alternative outcomes has been a topic of extensive research. Interpersonal comparison of utilities is required if the welfare function is to capture not just individual rank ordering of alternatives but also individual strength of preferences. Even though several approaches to perform such comparisons have been proposed, a consensus has not been achieved among economists regarding the best approach (21).

Welfare economics is often adopted as the theoretical foundation for health economic analysis. Among the different analytical frameworks, CBA is most closely founded on welfare economics (9). However, due to some practical considerations, such as ease of analysis, CEA has been the analytical framework that has been more widely used in health economic analysis. In CEA the outcome θ is valued in terms of QALYs. Let the QALYs associated with θ be QALY(θ). Then the loss function used in CEA can be expressed as:

$$L(a, \theta) = C(a) - C(a_0) - \{[QALY(\theta) - QALY(\theta_0)] \times V\} \quad \textbf{(3)}$$

where V is the societal monetary value of a QALY, and C(a) is the per person cost associated with action a (6;20;26). QALY(θ) is usually estimated as the average of QALYs of a representative sample of the population (11).

There are several theoretical as well as practical problems associated with the QALY-based loss function in equation 3. First, QALYs may not accurately capture individual utility for health outcomes. It has been shown that QALYs can be interpreted as a utility measure only if restrictive assumptions such as constant proportional trade-off and utility independence are satisfied (24), and there is some evidence that these assumptions do not hold in practice (3;8). Second, CEA provides the same resource allocation as CBA only if every individual in society has the same monetary value for a QALY (17). Again, this assumption is not likely to hold, implying that CEA could possibly diverge preferred CBA methodology in terms of optimal resource allocation decisions.

From a practical perspective, it is difficult to value characteristics of a drug intervention, such as dosing frequency or mode of administration, in a QALY-based approach, since individuals may be unwilling to trade off years of life for these characteristics. However, when these characteristics of an intervention affect individual utility, and hence societal welfare, it is important to capture the effect of differences in these characteristics on the loss function. One advantage of the QALY approach is that it is extremely easy to use. Once the utility of the health states that are part of the health outcomes are estimated, the QALY of any health outcome (combination of health states and time periods) can be computed quite easily.

If welfare economics is adopted as the theoretical foundation for the economic analysis, as recommended by the panel on cost-effectiveness, it may be preferable to use the CBA framework (9). In CBA the loss function corresponding to societal net benefit is given by:

$$L(a, \theta) = C(a) - C(a_0) - W[(a, \theta), (a_0, \theta_0)] \quad \textbf{(4)}$$

where W[(a, θ), (a_0, θ_0)] is the average societal willingness to pay (WTP) for (a, θ) instead of (a_0, θ_0). One practical problem associated with CBA is that there could be a large number of potential outcomes (θ), and estimating the WTP for each of them may not be possible in practice.

https://doi.org/10.1017/S0266462301104046 Published online by Cambridge University Press

Conjoint analysis provides a practical means of estimating equation 4. As in equation 2 we can use conjoint analysis to derive the function:

$$U(a, \theta) = \Sigma_i \Sigma_j \Sigma_{ij} \times Y_{ij} + v \times C(a) \quad \textbf{(5)}$$

where $U(a, \theta)$ is the utility associated with intervention a and outcome θ. The coefficients λ_{ij} capture all relevant characteristics of (a, θ) other than cost, while v captures the effect on preference of the cost associated with intervention a. As stated earlier, equation 5 can be estimated at the aggregate level based on all the responses, or it could be estimated at the individual respondent level. The CBA loss function can be written as:

$$L(a, \theta) = [U(a, \theta) - U(a_0, \theta_0)]/v \quad \textbf{(6)}$$

based on the aggregate utility function. In equation 6 dividing by v converts the loss to monetary units as in the CBA loss function in equation 4. If we estimate utility functions at the individual level, the loss function is given by:

$$L(a, \theta) = \Sigma_k L_k(a, \theta) \quad \textbf{(7)}$$

where

$$L_k(a, \theta) = [U_k(a, \theta) - U_k(a_0, \theta_0)]/v_k. \quad \textbf{(8)}$$

The subscript k gives the utility value and the coefficient for individual k. If conjoint analysis is administered to a sample of the population, sampling weights can be used to aggregate the net benefits in equation 7. Equations 6 and 7 can lead to different allocation decisions if the preferences are not homogeneous. The methodology in equation 7 is preferred, since it more closely mirrors the CBA methodology.

Thus, conjoint analysis provides us a practical means of deriving the loss function used in CBA. Given the closer relationship of CBA to welfare economics, from a theoretical perspective this approach may be preferable to a loss function based in the QALY approach.

USING CONJOINT METHODOLOGY IN HEALTH ECONOMIC ANALYSIS: A PILOT STUDY

We will use a hypothetical new drug for treating AMI to illustrate the use of conjoint analysis in conducting economic analysis. We derived the attributes in Table 1 as key

Table 1. Attributes and Levels for AMI Conjoint Analysis Survey

Attribute	Levels
30-day mortality	6.3%
	5.3%
Patency at 60 minutes	84%
	74%
Rate of reinfarction	3%
	4%
Rate of stroke	1.55%
	2.05%
Drug cost	$2,300
	$2,700
Annual cost	$26,690
	$27,190

https://doi.org/10.1017/S0266462301104046 Published online by Cambridge University Press

Table 2. Description of Potential Outcomes of Drug X

Attribute	θ_1	θ_2
30-day mortality	5.3%	6.3%
Patency at 60 minutes	74%	84%
Rate of reinfarction	4%	4%
Rate of stroke	2.05%	1.55%
Drug cost	$2,700	$2,700
Annual cost	$26,690	$26,690

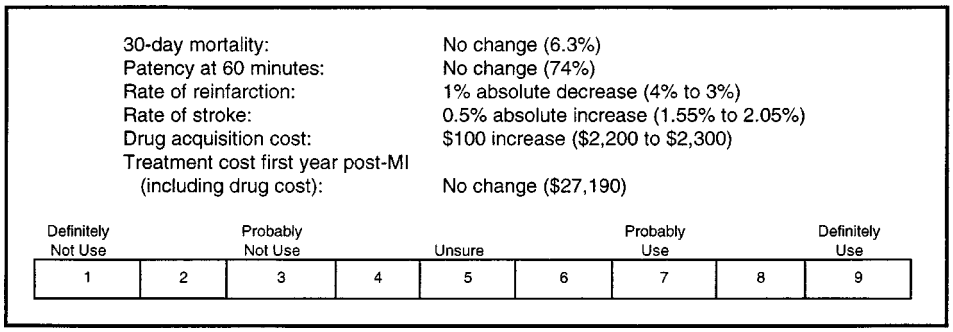

Figure 1. Example of conjoint question for AMI.

drivers of decision making regarding therapies for AMI after in-depth surveys with two decision makers. The levels were chosen to cover the range of characteristics that we can reasonably expect to see in a new therapy.

Let us assume that the objective of the analysis is to decide whether to adopt a new drug X, which will add $500 dollars to the cost of therapy. For the sake of simplicity, let as assume that there are two potential outcomes of therapy, θ_1 and θ_2 (shown in Table 2), each with a Bayesian posterior probability of .5. We now describe how the results of a pilot conjoint study that we conducted can be used to determine whether to adopt this new product.

In our pilot survey we used a full profile approach, where respondents were asked to evaluate different profiles for a new AMI drug. An example of a survey question is shown in Figure 1. For each drug profile the respondents were asked to state their likelihood of using the drug on a 1 to 9 scale. The survey was administered to 11 decision makers—six clinical cardiologists, two emergency department physicians, and three pharmacy directors.

The responses to the survey where used to estimate the coefficients in equation 5 for each of the respondents based on an OLS model. The product characteristics were used as independent variables, and the rating provided by the respondent minus five was used as the dependent variable. Thus, the dependent variable for a product that the respondent would definitely use (rating = 9) was 4, while the dependent variable for a product that respondent would definitely not use (rating = 1) was −4. The results of the estimation are shown in Table 3, which also shows the loss functions L(X, θ_1) and L(X, θ_2) computed using equation 8. Now, the expected loss is given by:

$$E[L(X, \theta)] = 0.5 \times L(X, \theta_1) + 0.5 \times L(X, \theta_2) \quad \textbf{(9)}$$

The expected loss is −$299 for drug X, and hence it should be adopted.

https://doi.org/10.1017/S0266462301104046 Published online by Cambridge University Press

Table 3. Results of OLS Model for AMI

	Regression coefficients							
Respondent	30-day mortality	Patency	Reinfarction	Stroke	Drug cost	Annual cost	$L_k(X, \theta_1)$	$L_k(X, \theta_2)$
1	0.5439	−0.0456	0.2106	−1.5788	−0.0018	0.0004	$741	$859
2	2.1346	−0.0865	0.4679	−0.3975	−0.0017	−0.0004	−$531	$1,148
3	1.3547	0.2021	1.0214	−6.6239	−0.0022	0.0020	$1,166	−$625
4	0.5329	0.2866	0.5329	−0.2676	−0.0016	0.0004	−$48	−$1,627
5	2.5024	0.0169	0.8357	−3.6619	−0.0034	−0.0003	$386	$532
6	4.1013	−0.0232	0.4347	−5.1306	−0.0010	0.0009	−$754	$1,029
7	0.5554	0.1889	0.8888	−1.5558	−0.0024	0.0011	$680	−$207
8	3.0384	0.0038	1.0384	0.0768	−0.0004	0.0014	−$7,066	$91
9	0.7431	0.0076	0.4097	−2.5139	−0.0055	−0.0018	$657	$550
10	−0.1041	0.0229	−0.4375	−2.2082	−0.0040	0.0011	$818	$454
11	2.9604	−0.0040	1.2937	−5.4125	0.0015	0.0013	$670	$477
Total							−$3,278	$2,680

Note that the example is presented here primarily with the objective of illustrating the analysis methodology. We acknowledge the severe limitations of the pilot study for the purposes of decision making in practice. First, the limited number of questions do not allow us to reliably estimate preferences at the individual level. Second, the estimates were obtained from decision makers; if we accept the welfare economics foundation, the estimates should be obtained from a random sample of the general population. We will discuss means of dealing with these limitations in the next section. These limitations, not withstanding the example, are useful for the purposes of illustrating how conjoint analysis methodology can be used to derive a loss function based in CBA and how the Bayesian posterior probabilities can be used to compute the expected loss associated with a decision. This then can allow the decision maker to choose the decision with the least expected loss.

DISCUSSION

Little work has been done in examining how uncertainty surrounding the cost-effectiveness estimates can be incorporated into the healthcare decision-making process. Sufficient controversy remains regarding the value of this information to the decision maker and the decision model in which this information would play a valuable role. Bayesian analysis is a promising alternative approach to classical statistics in incorporating uncertainty into health economic analysis. Bayesian decision analysis provides an analytical framework for identifying the optimal decision by combining the probability distribution around the outcome parameters with estimates of the loss function.

In published Bayesian decision analyses in health economics (6), the loss function is represented in dollar terms, where a value such as $50,000 per QALY was used to convert QALY gains into dollars. This approach allows one to easily estimate the loss function from available QALY estimates. However, there are several theoretical and practical problems associated with using a loss function based on QALYs. QALYs may not truly capture the utility associated with a health outcome (8). Further, the QALY approach provides a loss function consistent with the welfare economic foundation only under unrealistic conditions (17). From a practical perspective, it may be difficult for patients to determine how many years of life they are willing to give up to have the medication in caplet form instead of tablet. The incremental value of such change in form of medication may play in important role in the loss function but cannot be easily capture in a QALY-based approach.

https://doi.org/10.1017/S0266462301104046 Published online by Cambridge University Press

A loss function based on CBA is more closely related to welfare economics, which is recommended by the panel on cost-effectiveness as the appropriate theoretical foundation for health economic analysis (9). CBA has not found wide use in health economics, partly due to practical problems associated with estimating the loss function for a large number of outcomes. In this paper we described a methodology—conjoint analysis—that can be used to estimate the CBA loss function in practice. The strength of the conjoint methodology lies in the fact that it allows one to express the loss function as a function of cost and outcome. Thus, conjoint analysis allows one to easily compute the loss function for a range of outcomes. Conjoint analysis may also provide a more effective means of capturing the value of less important aspects of the intervention, such as mode of administration or dosing frequency.

In this paper, we illustrated the use of the conjoint methodology through the use of a small pilot study. The pilot study design suffered from several limitations—the small number of attributes and levels, the assumption of attribute independence, and the constraints on the loss function imposed by the estimation model. However, the purpose of the pilot studies was to explore the use of the conjoint analysis technique to estimate the loss function, and not to provide a definitive loss function that can be used in future analyses.

In our pilot study we surveyed a small number of decision makers. If the analysis is based in welfare economics, the conjoint survey should be administered to a representative sample of the general population (11). However, the welfare economics foundation of health economics has not found universal acceptance among economists. The extra-welfarist approach has been proposed as an alternative (7;28). In this approach, QALYs are assumed to have the same value for each individual. Further, the loss function need not be based on a survey of the general population but can be based on preferences obtained from decision makers. Even if we adopt the extra-welfarist approach, conjoint analysis can still be used as an alternative to the standard gamble or time trade-off techniques to estimate QALYs. If life-years gained (in perfect health) is included as an attribute in the conjoint study, dividing the utility difference by the coefficient of this attribute (instead of the coefficient of cost) will yield a loss function in QALYs, since this division converts the difference in utility into life-year units.

The drawback of the conjoint methodology is the complexity involved in designing the conjoint survey, administering the survey, and reliably estimating individual utility function. However, advances made in design and analysis of conjoint surveys has made this technique significantly easier to use than in the past. Off-the-shelf software is available to design and administer conjoint surveys. The use of techniques such as hierarchical Bayes enable more accurate and reliable estimation of individual utility functions from limited survey data (1). The use of Bayesian decision analysis as a means of capturing uncertainty surrounding cost and efficacy estimates in healthcare decision making and the use of conjoint analysis to estimate loss functions that fully represent the impact of product attributes on patient and societal welfare bear further investigation.

REFERENCES

1. Allenby GM, Arora N, Ginter JL. Incorporating prior knowledge into the design and analysis of conjoint studies. *Journal of Marketing Research*. 1995;17:152-162.
2. Bala MV, Mauskopf JO. The estimation and use of confidence intervals in economic analysis. *Drug Information Journal*. 1999;33:841-848.
3. Bala MV, Wood LL, Zarkin GA, et al. Are health states timeless? The case of the standard gamble method. *J Clin Epidemiol*. 1999;52:1047-1053.
4. Briggs AH. A Bayesian approach to stochastic cost-effectiveness analysis. *Health Econ*. 1999; 8:257-261.
5. Cattin P, Wittink RR. Commercial use of conjoint analysis: A survey. *Journal of Marketing*. 1982; 46:44-53.

6. Claxton K. The irrelevance of inference: A decision-making approach to the stochastic evaluation of health care technologies. *J Health Econ*. 1999;18:341-364.
7. Culyer AJ, Evans RG. Mark Pauly on welfare economics: Normative rabbits from positive hats. *J Health Econ*. 1996;15:243-251.
8. Gafni A. The standard gamble method: What is being measured and how it is interpreted. *Health Serv Res*. 1994;29:207-224.
9. Garber AM, Weinstein MC, Torrance GW, Kamlet MS. Theoretical foundations of cost-effectiveness analysis. In: MR Gold, JE Siegel, LB Russell, MC Weinstein, eds. *Cost-Effectiveness in Health and Medicine*. New York: Oxford University Press; 1996:25-53.
10. Gelman A, Carlin JB, Stern HS, Rubin DB. *Bayesian data analysis*. New York: Chapman & Hall, 1995.
11. Gold MR, Siegel JE, Russell LB, Weinstein MC, eds. *Cost-effectiveness in health and medicine*. New York: Oxford University Press; 1996.
12. Green PE, Rao VR. Conjoint measurement for quantifying judgmental data. *Journal of Marketing Research*. 1971;8:355-363.
13. Green PE, Srinivasan V. Conjoint analysis in consumer research: Issues and outlook. *Journal of Consumer Research*. 1978;5:103-123.
14. Green PE, Srinivasan V. Conjoint analysis in marketing: New developments with implications for research and practice. *Journal of Marketing*. 1990;54:3-19.
15. Greene WH. Econometric analysis. Upper Saddle River, NJ: Prentice-Hall, Inc.; 1997.
16. Hillier FS, Lieberman GJ. Introduction to operation research. San Francisco: Holden-Day Inc.; 1980.
17. Johannesson M. The relationship between cost-effectiveness analysis and cost-benefit analysis. *Soc Sci Med*. 1995;41:483–489.
18. Johannesson M, Weinstein MC. On the decision rules of cost-effectiveness analysis. *J Health Econ*. 1993;12:459–467.
19. Laska EM, Meisner M, Siegel C. Statistical inference for cost-effectiveness ratios. *Health Econ*. 1997;6:229-242.
20. Laska EM, Meisner M , Siegel C, Stinnett AA. Ratio-based and net benefit-based approaches to health care resource allocation: Proofs of optimality and equivalence. *Health Econ*. 1999;8:171-174.
21. Luce RD, Raiffa H. *Games and decisions: Introduction and critical survey*. New York: John Wiley & Sons, Inc.; 1964.
22. Manning WG, Fryback DG, Weinstein MC. Reflecting uncertainty in cost-effectiveness analysis. In: MR Gold, JE Siegel, LB Russell, MC Weinstein, eds. *Cost-effectiveness in health and medicine*. New York: Oxford University Press; 1996:247-275.
23. O'Hagan A, Stevens JW, Montmartin J. Inference for the cost-effectiveness acceptability curve and cost-effectiveness ratio. *PharmacoEconomics*. 2000;17:339-349.
24. Pliskin J, Shepard D, Weinstein MC. Utility functions for life-years and health status. *Operations Research*. 1980;28:206-224.
25. Ryan M, Hughes J. Using conjoint analysis to assess women's preferences for miscarriage management. *Health Econ*. 1997;6:261-273.
26. Stinnett AA, Mullahy J. A new framework for the analysis of uncertainty in cost-effectiveness analysis. *Med Decis Making*. 1998;18(suppl):S68-S80.
27. Urban GL, Hauser JR. *Design and marketing of new products*. Englewood Cliffs, NJ: Prentice-Hall, Inc.; 1980.
28. Wagstaff A. QALYs and the equity-efficiency tradeoff. In: *Cost-benefit analysis*. New York: Cambridge University Press;1994:428-447.