European Journal of Operational Research 221 (2012) 222–230

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

Nonparametric predictive utility inference

B. Houlding [a,*], F.P.A. Coolen [b]

[a] Discipline of Statistics, Trinity College Dublin, Ireland
[b] Department of Mathematical Sciences, Durham University, UK

ARTICLE INFO

Article history:
Received 18 April 2011
Accepted 13 March 2012
Available online 20 March 2012

Keywords:
Decision analysis
Nonparametrics
Predictive inference
Exchangeability
Uncertain utility

ABSTRACT

We consider the natural combination of two strands of recent statistical research, *i.e.*, that of decision making with uncertain utility and that of Nonparametric Predictive Inference (NPI). In doing so we present the idea of Nonparametric Predictive Utility Inference (NPUI), which is suggested as a possible strategy for the problem of utility induction in cases of extremely vague prior information. An example of the use of NPUI within a motivating sequential decision problem is also considered for two extreme selection criteria, *i.e.*, a rule that is based on an attitude of extreme pessimism and a rule that is based on an attitude of extreme optimism.

1. Introduction

The Bayesian paradigm, *e.g.*, de Finetti (1974), coupled with the expected utility hypothesis of Bernoulli (1738), provides a transparent and attractive methodology for solving problems of decision making under uncertainty. In this approach preferences over a set of possible decisions are reconstructed by taking into account both the probability that each decision leads to a particular outcome, and the relative preference for obtaining that outcome as measured by its utility value. Furthermore, if the outcome that pertains from a particular decision depends on the value of an unknown random quantity, then the probabilities associated with the set of possible decision outcomes are typically assumed to be subject to an assigned prior parametric distribution. Learning then occurs following observation of data that has a probabilistic dependence with the unknown random quantity of interest, and the usual 'posterior is proportional to likelihood times prior' of Bayes' Theorem is employed.

However, implicit within this theory (and hence necessary for its application) is the assumption that the Decision Maker (DM) knows her preferences, meaning that she can assign an appropriate utility function (with domain the full set of all possible decision outcomes) for use within the problem. In applications this is usually achieved by either assuming a fixed utility form, *e.g.*, a logarithmic utility function for monetary returns, or by selecting specific utility values for particular and relevant decision outcomes. As such, classical Bayesian subjective expected utility theory does not permit inherent uncertainty in preferences over decisions. It also does not allow the learning of utility and specifies that the DM will never be surprised by the utility of a realized outcome.

Nevertheless, not for all situations is the assumption of a known preference relation over outcomes deserved, and often a DM may instead need to learn her preferences through suitable experimentation. Indeed, a DM may consider it inappropriate to assign a particular and fixed utility value for any outcome that is novel or unfamiliar, choosing instead to only do so after direct experience or exposure. Such cases of utility uncertainty motivate so-called adaptive utility theory, *e.g.*, Cyert and DeGroot (1975), Houlding (2008), Houlding and Coolen (2007, 2011), which generalizes the traditional utility concept by only requiring the utility function be known up to the value of some uncertain utility parameter. The principal idea of adaptive utility is then to treat the uncertain utility parameter in the same manner that unknown random quantities are typically treated in standard Bayesian statistical inference, *i.e.*, they are subjected to a parametric learning model in accordance with Bayes' Theorem. Yet, and despite adaptive utility theory explicitly permitting a DM to remain uncommitted to a presumed known and correct utility function, its previous use has required knowledge of a precise and meaningful prior distribution concerning true preferences, something that is unlikely to be considered either reasonable or justifiable when selecting from outcomes that include (initially) new and foreign possibilities.

Rather than assuming a precise prior distribution over an uncertain utility parameter, interest here is in the use of the Nonparametric Predictive Inference (NPI) technique of Coolen (1996,

* Corresponding author. Address: Rm 105, Discipline of Statistics, Lloyd's Institute, Trinity College Dublin, Dublin 2, Ireland. Tel.: +353 (0) 1896 2933; fax: +353 (0) 1 677 0711.

E-mail address: brett.houlding@tcd.ie (B. Houlding).

0377-2217/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2012.03.024

2006), Augustin and Coolen (2004), which is a low structure statistical technique arising naturally as a result of Hill's $A_{(n)}$ assumption, Hill (1968, 1988, 1993). Given a known ordered series of utility values that are considered subject to a post-data exchangeability assumption, Nonparametric Predictive Utility Inference (NPUI) proceeds by assigning equal mass to the probability that a new utility value falls within any of the intervals formed by the known ordered utility, leading to the quantification of such utility uncertainty through interval probability.

An outline of the remainder is as follows; Section 2 reviews the concept of uncertain utility and briefly describes how this can be taken into account through adaptive utility theory, whilst Section 3 reviews the statistical technique of NPI. In Section 4 the two strands of uncertain utility and nonparametric predictive inference are combined to formally introduce the NPUI model, with an illustrative example being given in Section 5. Finally, Section 6 concludes with a discussion of possible future directions.

2. Uncertain utility

The assumption that a DM can accurately identify the utility value of any considered decision outcome is prevalent within the theory and application of Bayesian decision making under uncertainty. Often this will be through the use of an assumed model for the utility of outcomes from a continuous domain, *e.g.*, a logarithmic model for monetary returns. Alternatively, specific utility values may be considered appropriate following suitable introspection concerning the decision problem. In either case, the possibility that a DM may not *a priori* know their preference for a given outcome is often ignored.

In contrast, the theory of adaptive utility, as introduced by Cyert and DeGroot (1975) and further developed by Houlding et al. (2008), Houlding and Coolen (2007, 2011), generalises classical Bayesian decision theory, suggesting a methodology for decision selection even when the DM is unable to fully specify her preferences. In this setting the DM is permitted to be uncertain over her true preferences, with the appropriate utility function over decision outcomes only being known up to the value of some uncertain parameter θ. Such a parameter is referred to as the DM's *state of mind*, and it can be used to model uncertainty over any aspect of the DM's preferences, with interesting examples including vectors of unknown trade-off weights or unknown level of risk aversion. Notationally, such a dependence between the utility function $u(\cdot)$ and the state of mind θ was displayed via the inclusion of a conditioning argument, *e.g.*, $u(\cdot|\theta)$.

Instead of assuming that the utility function is fully known, adaptive utility theory makes use of a probabilistic specification concerning the uncertain state of mind θ. Bayesian updating can then occur once the DM receives additional information concerning her true preferences, and previous examples considered for such utility related information include noise corrupted observations of the true utility value, or the sign of the difference between the prior expectation of the utility value for a given outcome and the utility value that was actually received (*i.e.*, an indication of elation or disappointment). An adaptive utility function ${}_{a}u(\cdot)$ was then defined as the expectation of the possible utility values with respect to beliefs over the state of mind, *i.e.*, ${}_{a}u(\cdot) = E_{\theta}[u(\cdot|\theta)]$.

In essence, the use of adaptive utility theory is analogous to the use of a hierarchical prior within robust Bayesian analysis, *e.g.*, Berger (1993). A utility value, once scaled to fall within the interval [0, 1], corresponds to a probability, with the utility of decision outcome o being that probability p which makes the DM indifferent between receiving o for sure, or selecting the decision which results in most preferred outcome o^* with probability p and least preferred outcome o_* otherwise, see DeGroot (1970). Adaptive utility theory allows the DM to be uncertain of the value of p that results in her indifference, instead considering a non-degenerate prior subjective probability distribution for its value.

That such a probability distribution concerning utility values may change following updating in light of additional information motivates the name adaptive utility theory, for if the probability distribution did alter, then so would the expected utility return, *i.e.*, the adaptive utility value. In other words, the adaptive utility value of an unfamiliar decision outcome will 'adapt' in light of additional information concerning preferences. This is in contrast to classical utility theory in which the utility values of all decision outcomes are considered known and fixed. Yet, explicitly permitting utility uncertainty does not alter the suggested decision selection within a one-off decision problem, as the strategy arising from the adaptive utility setting is the same as that which would arise from traditional theory if the adaptive utility values were assumed equal to the true utility values. However, in a sequential decision problem the acceptance that certain decision outcomes do not have known utility can alter the optimal selection strategy.

Notable alternative theories that similarly seek to incorporate uncertain preference within a decision making paradigm include the approaches of Farrow and Goldstein (2006) and Ben-Haim et al. (2009). Rather than expressing uncertainty over an unknown utility parameter through the use of a precise prior probability distribution, Farrow and Goldstein allow the DM to remain non-committed and to instead only provide an upper and lower bound for the true utility value (through the declaration of lower and upper bounds on trade-off parameters in a multi-attribute utility hierarchy). In this respect the approach of Farrow and Goldstein is similar to the NPUI approach presented here. However, differences result in that the use of the NPI statistical technique would appear to be a simpler approach that is data driven. In particular, the NPUI approach considered here does not require any explicit declaration of subjective judgments and/or expressions concerning the utility of a novel outcome other than that of a usually reasonable and objective post-data exchangeability assumption.

In contrast, the approach of Ben-Haim et al. is to instead specify an Info-Gap model of utility uncertainty, and is relevant for situations of severe uncertainty whereby, other than the specification of a best point-estimate guess, nothing else can be elicited. In particular, there is no probabilistic specification of how accurate such a best point-estimate guess may be. Instead a nested subset of possible horizons of uncertainty is specified and the decision selected which is deemed most robust in that it guarantees a specified minimum critical return for the largest horizon of uncertainty. Unlike the imprecise theory of Farrow and Goldstein, or the NPI method used here, the Info-Gap approach of Ben-Haim et al. is non-probabilistic, and can not quantify, or even give bounds on, the probability of an outcome for any quantity that is subject to 'Info-Gapping'.

There is also a long and developed literature on theories that seek to take into account the descriptive behaviour of decision makers following observations such as that identified in the Ellsberg Paradox, see for example Ellsberg (1961), Schmeidler (1989), Epstein (1999), Ghiradato et al. (2004), Klibanoff et al. (2005), Halevy (2007) and the references therein. These typically deal with what has been described as 'ambiguity aversion', which has arisen following observational studies where individuals are found to be less prepared to take part in bets when the stated potential outcomes depended on the occurrence or non-occurrence of a vague or unfamiliar event, compared to other bets whose stated potential outcomes depended on the occurrence or non-occurrence of events with which the individuals held greater experiences.

Here we do not consider situations where the likely outcome of a decision is unknown (though that would be a straightforward generalisation of the material presented), but instead with