Supervisors: Dr Nikolaos Argyris (Dept. of Statistics), Prof. Juergen Branke (Warwick Business School)

Topic: Constructing utility functions from inconsistent preference data for Decision Support.

It is routinely observable, both in experiments and in practical situations, that many of the assumptions underpinning *rational choice* are in fact violated by observable choice behaviour. This poses a major problem for methods for Decision Support (DS), which, broadly speaking, rely on using normative models to extrapolate from observed preferences to recommendations. When such preference information is incompatible with the assumed underlying model, a situation referred to as *Inconsistency*, DS methods need reconcile the two in an "acceptable" way. Thus the question arises, as to what is the best way to "resolve" preference Inconsistency.

This problem arises in different domains. One example is Decision Analysis methodologies, where the aim is to learn a single agent's preferences and guide them to an "optimal" choice. Another example is the field of Machine Learning, where various methods can be used to extract preferences from the observed choices of many agents (e.g. purchasing decisions) and then construct models of typical-user behaviour or recommender-systems. Finally, the fields of Marketing Science and Econometrics have also developed methods for extracting preferences from observed choices.

The proposed project will explore the question of how to best resolve preference Inconsistency, but in doing so will focus on a particular family of methods applied to a particular type of problem. The methods in focus will be those based on the construction of representative *utility functions* from observed preferences, as typical in the field of Decision Analysis (DA). The problem in focus will be that of constructing an *optimal ranking* of a set of multi-attribute objects (although the related problem of *optimal choice* from the same set may be chosen instead). The overall work will involve two main components: a) a literature review, and b) computational tests.

The starting point for the literature review will be the relatively small number of approaches for inconsistency resolution within the context of DA. Depending on time, the review may also encompass approaches in related domains (e.g. Machine Learning, Econometrics). The review will identify a small subset of appropriate methods to be tested through computational tests, as well as a number of criteria to be used in evaluating the output of these methods. The computational part of the work will first involve simulating datasets of inconsistent preference information. This will be done via generating multi-attribute objects and then simulating sets of inconsistent preference information over these objects (e.g. through distorting the scores of *a-priori* defined utility functions). Following this, the different methods will be applied to the simulated inconsistent data to generate different utility functions which "approximately" represent the inconsistent preferences. The last step would be to compare these approximations. For ordinal data this may be done by comparing the approximate rankings with the true ranking. For cardinal data the differences between true and approximate utility scores will also be compared.

It is envisaged that the short review and the project's findings will be written up in the form paper to be submitted to a Journal in Decision Analysis, Operations Research or a related field. The research topic offers potential for pursuing a PhD on "Learning Preferences", within or across the domains of Decision Analysis, Machine Learning and Statistics. Promising extensions include the aggregation of preferences across a group of individuals, and interactive approaches that allow to elicit additional information where it is most helpful.

Contact: Nikolaos Argyris, Dept. of Statistics, n.argyris@warwick.ac.uk