Risk & Sustainable Management Group

Risk and Uncertainty Working Paper: R08#2

Research supported by an Australian Research Council Federation Fellowship
http://www.arc.gov.au/grant_programs/discovery_federation.htm

Bounded rationality and small worlds

By

Simon Grant
Rice University

and

John Quiggin
Australian Research Council Federation Fellow, University of Queensland

Schools of Economics and Political Science
University of Queensland
Brisbane, 4072
rsmg@uq.edu.au
http://www.uq.edu.au/economics/rsmg

THE UNIVERSITY OF QUEENSLAND AUSTRALIA

Bounded rationality and small worlds*

Simon Grant
Rice University

John Quiggin
University of Queensland

9 June 2008

Abstract

We consider conditions under which the representation of the world available to a boundedly rational decision-maker, whose awareness increases over time, constitutes an adequate 'small world' (in the sense of Savage 1954) for the assessment of a given decision. Equivalently, we consider whether boundedly rational decision-makers who gradually become aware of all relevant contingencies, can pursue a strategy that is sequentially consistent. We derive conditions on beliefs and preferences that yield a separation between the set of propositions of which the boundedly rational decision-maker is aware and those of which she is unaware and show that these conditions are sufficient to ensure sequential consistency.

JEL Classification: D80, D82

*We thank Joseph Halpern, Jeff Kline, Mamoru Kaneko and Nancy Wallace for helpful comments and criticism.

1 Introduction

Bayesian decision theory and its generalizations provide a powerful set of tools for analyzing problems involving state-contingent uncertainty. In problems of this class, decision-makers begin with a complete specification of uncertainty in terms of a state space (a set of possible states of the world). The ultimate problem may be formulated as a choice among a set of acts, represented as mappings from the state space to a set of possible outcomes. In many applications, there is an intermediate stage in which the decision-maker may obtain information in the form of a signal about the state of the world, represented by a refinement of the state space. That is, any possible realization of the signal means that the true state of the world must lie in some subset of the state space.

The starting point of Bayesian analysis is the state space, representing all possible contingencies. A fundamental difficulty with such a state-contingent models of decision-making under uncertainty is that, in reality, decision-makers are boundedly rational and do not possess a complete state-contingent description of the uncertainty they face. Decision makers cannot foresee and consider all the contingencies relevant to their decisions (Grant & Quiggin, 2007a&b, Heifetz, Meier & Schipper, 2006, Halpern & Rego, 2006a).

In this paper, we consider conditions under which the representation of the world available to a boundedly rational decision-maker, whose awareness increases over time, constitutes an adequate 'small world' (in the sense of Savage 1954) for the assessment of a given decision. Equivalently, we consider whether boundedly rational decision-makers who gradually become aware of all relevant contingencies, can pursue a strategy that is sequentially consistent. Here sequential consistency means that, reconsidering past decisions in the light of increased awareness about the possible states of the world, but disregarding information received after the decision was made, the individual would still regard the decisions as ex ante optimal.

The paper is organized as follows. We first briefly outline in section 2 a model we developed in Grant and Quiggin (2007a) that provides a representation in which the state of the world is represented by the truth values for a set of propositions (the syntactic representation). We use this representation to describe a game with nature, and derive the expected utility of possible strategies from two viewpoints. The external viewpoint is that of an unboundedly rational (but not, in general, perfectly informed) decision-maker with access to a complete set of states of the world and an associated set of propositions rich enough to describe all possible states. The second is that of a boundedly rational decision-maker with limited awareness. In Section 3 we derive conditions on beliefs and preferences that yield a separation between the set of propositions of which the boundedly rational decision-maker is aware and those of which she is unaware

In4 we present a dynamic model in which individual awareness increases over time, reaching the maximal (relevant) level of awareness when the game concludes. We derive our main result, showing that the conditions of Section 3 are sufficient to ensure sequential consistency. We conclude with a discussion of some of the implications of our analysis.

2 Structure and notation

We adapt the model of choice under uncertainty developed by Grant and Quiggin (2007a) in which an individual does not necessarily possess a complete description of the world. As we discuss below, the underpinnings of this model can be embedded in a dynamic tree structure that can be viewed as an extensive-form game between Nature and our boundedly rational decision-maker, where the awareness of the decision-maker increases gradually through learning and discovery. But for our purposes here, it is sufficient for us to model all decisions and beliefs in terms of binary (elementary) propositions that either nature or the decision-maker determines the truth value. The key distinction will be between the external viewpoint and the limited or restricted viewpoint of the decision-maker.

Let the set of states of the world from the external viewpoint be Ω. We focus on the representation $\Omega = 2^{\mathbf{P}^0} \times 2^{\mathbf{P}^1}$, where $\mathbf{P}^0 = \{p_1^0, \ldots, p_M^0\}$ is a finite set of 'elementary' propositions about the world that are determined by nature, and $\mathbf{P}^1 = \{p_1^1, \ldots, p_N^1\}$ is a finite set of 'elementary propositions (i.e. [binary] decisions) that the individual controls. Each proposition in $\mathbf{P}^0$ is a statement such as 'The fourth named storm of the year is a force five hurricane and makes landfall at Galveston.' Each proposition in $\mathbf{P}^1$ is a statement such as 'The decision-maker buys flood insurance for her house in Houston.' Thus, an exhaustive description of the state of the world from the external or objective viewpoint, consists of an evaluation of each of the propositions in $\mathbf{P}^0$ and $\mathbf{P}^1$. With each proposition p_n^i and each possible state of the world ω in Ω, a fully informed observer can associate a truth value $V\left(p_n^i;\omega\right)\in\{0,1\}$, which will be 1 if p_n^i is true and 0 if p_n^i is false at ω.

By way of contrast to the external viewpoint, we shall consider a decision maker who has only *limited* awareness of the uncertainty embodied in Ω. In particular, there are propositions both in $\mathbf{P}^0$ and in $\mathbf{P}^1$ that she does not or cannot consider. Let $\mathbf{P}_{\mathrm{R}}^0=\left\{p_1^0,,\ldots,p_{M_{\mathrm{R}}}^0\right\}$ and $\mathbf{P}_{\mathrm{R}}^1=\left\{p_1^1,,\ldots,p_{N_{\mathrm{R}}}^1\right\}$ denote the 'restricted' sets of propositions that she explicitly considers are under the control of nature and are under her control, respectively. Correspondingly, let $\mathbf{P}_{\mathrm{U}}^0=\left\{p_{M_{\mathrm{R}}+1}^0,,\ldots,p_M^0\right\}$ and $\mathbf{P}_{\mathrm{U}}^1=\left\{p_{N_{\mathrm{R}}+1}^1,,\ldots,p_N^1\right\}$, be the sets of propositions that are controlled respectively, by nature and by her own decisions that she currently does not consider when formulating her plan of action.

From the viewpoint of our partially aware decision-maker, any state of her (restricted) world, can be characterized by pair of rational numbers $(s_{\mathrm{R}}^0,s_{\mathrm{R}}^1)$, where s_{R}^0 is an element of $S_{\mathrm{R}}^0\subseteq\left[0,1-2^{-M_{\mathrm{R}}}\right)\cap\mathbb{Q}$, and s_{R}^1 is an element of $S_{\mathrm{R}}^1\subseteq\left[0,1-2^{-N_{\mathrm{R}}}\right)\cap\mathbb{Q}$. Formally, the 'set of strategies for player i, $i=0$ (nature),1 (decision-maker) is,

$$S_{\mathrm{R}}^i:=\left\{s_{\mathrm{R}}^i\in\mathbb{Q}:s_{\mathrm{R}}^i=\sum_{p_n^i\in\mathbf{P}_{\mathrm{R}}^i}2^{-n}\times V\left(p_n^i;\omega\right),\text{ for some }\omega\in\Omega\right\}.$$

The set of *unconsidered* decisions of nature and the set of unconsidered decisions of the decision-maker herself can be represented as elements of $S_{\mathrm{U}}^0\subseteq\left[0,2^{-M_{\mathrm{R}}}-2^{-M}\right)\cap\mathbb{Q}$ and $S_{\mathrm{U}}^1\subseteq\left[0,2^{-N_{\mathrm{R}}}-2^{-N}\right)\cap\mathbb{Q}$, respectively, where for each ω in Ω, we have $s_{\mathrm{U}}^0\in S_{\mathrm{U}}^0$ and $s_{\mathrm{U}}^1\in S_{\mathrm{U}}^1$, given by

$$S_{\mathrm{U}}^i:=\left\{s_{\mathrm{U}}^i\in\mathbb{Q}:s_{\mathrm{U}}^i=\sum_{p_n^i\in\mathbf{P}_{\mathrm{U}}^i}2^{-n}\times V\left(p_n^i;\omega\right),\text{ for some }\omega\in\Omega\right\},\ i=0,1.$$

A state of the world ω in Ω, may be thus viewed as being jointly determined by a 'complete strategy' chosen by nature, that is, a state of *nature* $s^0\in S^0$, where

$$S^0=S_{\mathrm{R}}^0+S_{\mathrm{U}}^0=\left\{s^0:s^0=s_{\mathrm{R}}^0+s_{\mathrm{U}}^0\text{ for some }s_{\mathrm{R}}^0\in S_{\mathrm{R}}^0,\ \&\text{ for some }s_{\mathrm{U}}^0\in S_{\mathrm{U}}^0\right\},$$

and a 'complete strategy' choice of the individual $s^1 \in S^1$, where

$$S^1 = S^1_{\mathrm{R}} + S^1_{\mathrm{U}} = \left\{ s^1 : s^1 = s^1_{\mathrm{R}} + s^1_{\mathrm{U}} \text{ for some } s^1_{\mathrm{R}} \in S^1_{\mathrm{R}}, \ \& \text{ for some } s^1_{\mathrm{U}} \in S^1_{\mathrm{U}} \right\}.$$

We shall refer to $(S^0, S^1) \subset ([0,1) \cap \mathbb{Q})^2$ as the (normal or strategic) game-form associated with Ω and $(S^0_{\mathrm{R}}, S^1_{\mathrm{R}})$ as the restricted game-form for the decision maker with limited awareness.

2.1 Decision-making in the 'Large' and in the 'Small'

For a finite set E, let $\Delta(E)$ denote the set of probability distributions defined on E.

Fix a set of states of the world Ω, with associated sets of propositions $\mathbf{P}^0$ and $\mathbf{P}^1$.

A (fully aware) subjective expected utility maximizer (SEUM) decision-maker for this *large* world is characterized by a belief $\sigma^0 \in \Delta(S^0)$, a consequence function $c : S^0 \times S^1 \to \mathcal{C}$, where $\mathcal{C}$ is a space of consequences, and a utility index over consequences, $u : \mathcal{C} \to \mathbb{R}$. Choices for this individual are then ranked according to their subjective expected utility. That is, for any pair of strategies s^1 or $\tilde{s}^1$ in S^1, s^1 is at least as good as $\tilde{s}^1$ if and only if the subjective expected utility of the former is greater than or equal to the subjective expected utility of the latter, that is,

$$\sum_{s^0 \in S^0} u\left(c\left(s^0, s^1\right)\right) \sigma^0\left(s^0\right) \geq \sum_{s^0 \in S^0} u\left(c\left(s^0, \tilde{s}^1\right)\right) \sigma^0\left(s^0\right).$$

For a decision-maker who is only aware of propositions $\mathbf{P}^0_{\mathrm{R}}$ and $\mathbf{P}^1_{\mathrm{R}}$, to complete the description of her as a subjective expected utility maximizer for the restricted game-form $(S^0_{\mathrm{R}}, S^1_{\mathrm{R}})$, requires specifying a belief $\sigma^0_{\mathrm{R}} \in \Delta(S^0_{\mathrm{R}})$, a consequence function $c_{\mathrm{R}} : S^0_{\mathrm{R}} \times S^1_{\mathrm{R}} \to \mathcal{C}_{\mathrm{R}}$, where $\mathcal{C}_{\mathrm{R}}$ is the space of consequences of which she is aware, and a utility index over that space of consequences, $u_{\mathrm{R}} : \mathcal{C}_{\mathrm{R}} \to \mathbb{R}$. In the terminology of Savage, the restricted game-form $(S^0_{\mathrm{R}}, S^1_{\mathrm{R}})$ may be regarded as a *small world* within which decision analysis may be applied to choose among available strategies $\hat{s}^1_{\mathrm{R}}$ ranked according to their subjective expected utility:

$$\sum_{s^0_{\mathrm{R}} \in S^0_{\mathrm{R}}} u_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}}, \hat{s}^1_{\mathrm{R}}\right)\right) \sigma^0_{\mathrm{R}}\left(s^0_{\mathrm{R}}\right).$$

In the next section we shall specify circumstances in which even from the fully informed perspective (S^0, S^1), the small world $(S^0_{\mathrm{R}}, S^1_{\mathrm{R}})$ is an appropriate choice for modelling the decision among available strategies $\hat{s}^1_{\mathrm{R}}$ in S^1_{R}.

3 Consistent Small-world Bayesian Decisions

Suppose the 'true' or 'objective' uncertainty corresponds to Ω, and the fully aware SEU maximizer in the large world (S^0, S^1) is characterized by $(\sigma^0, c(.,.), u(.))$. We shall consider a (restricted) game-form given by $(S^0_{\mathrm{R}}, S^1_{\mathrm{R}})$. Notice the marginal distributions over s^0_{R} and s^0_{U}, derived from σ^0, are given by

$$\sigma^0_{\mathrm{R}}\left(s^0_{\mathrm{R}}\right) = \sum_{s^0_{\mathrm{U}} \in S^0_{\mathrm{U}}} \sigma^0\left(s^0_{\mathrm{R}} + s^0_{\mathrm{U}}\right) \text{ and } \sigma^0_{\mathrm{U}}\left(s^0_{\mathrm{U}}\right) = \sum_{s^0_{\mathrm{R}} \in S^0_{\mathrm{R}}} \sigma^0\left(s^0_{\mathrm{R}} + s^0_{\mathrm{U}}\right), \text{ respectively.}$$

Now consider a less than fully aware decision-maker who is a (small-world) SEU maximizer characterized by $(\sigma^0_{\mathrm{R}}, c_{\mathrm{R}}(.,.), u_{\mathrm{R}}(.))$. An 'optimal' choice $\hat{s}^1_{\mathrm{R}}$ given her limited awareness is given by

$$\hat{s}^1_{\mathrm{R}} \in \underset{s^1_{\mathrm{R}} \in S^1_{\mathrm{R}}}{\operatorname{argmax}} \sum_{s^0_{\mathrm{R}} \in S^0_{\mathrm{R}}} u_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}}, s^1_{\mathrm{R}}\right)\right) \sigma^0_{\mathrm{R}}\left(s^0_{\mathrm{R}}\right). \tag{1}$$

The question we address is the following. Under what conditions can we be assured that the 'optimal' choice $\hat{s}^1_{\mathrm{R}}$ in the small world $(S^0_{\mathrm{R}}, S^1_{\mathrm{R}})$ for the SEU maximizer $(\sigma^0_{\mathrm{R}}, c_{\mathrm{R}}(.,.), u_{\mathrm{R}}(.))$, would be part of the optimal choice in the large world (S^0, S^1) for the fully aware SEU maximizer $(\sigma^0, c(.,.), u(.))$. If these conditions are satisfied we say that the small world model is **consistent** with the large world model.

The first condition is the requirement that the consequence resulting from the decisions of nature and the individual of which the individual is aware, is separable from the consequence resulting from the decisions of nature and the individual of which the individual is unaware. Moreover, the utility index over this pair of consequences must have the so-called multiplicative form of Keeney & Raiffa (1976).

Definition 1 (Multiplicative Separable Utility) *The fully aware SEU maximizer's utility over consequences in the large world is said to be* multiplicatively separable *with respect to the small-world SEU maximizer's utility* *if there exists a consequence space* $\mathcal{C}_{\text{U}}$*, a consequence function* $c_{\text{U}}: S^0_{\text{U}} \times S^1_{\text{U}} \rightarrow \mathcal{C}_{\text{U}}$*, and a utility function* $u_{\text{U}}: \mathcal{C}_{\text{U}} \rightarrow \mathbb{R}$*, s.t. for all* $s^0_{\text{R}} \in S^0_{\text{R}}$*, all* $s^1_{\text{R}} \in S^1_{\text{R}}$*, all* $s^0_{\text{U}} \in S^0_{\text{U}}$ *and all* $s^1_{\text{U}} \in S^1_{\text{U}}$*,*

$$\begin{aligned} u\left(c\left(s^0_{\text{R}} + s^0_{\text{U}}, s^1_{\text{R}} + s^1_{\text{U}}\right)\right) &= u_{\text{R}}\left(c_{\text{R}}\left(s^0_{\text{R}}, s^1_{\text{R}}\right)\right) + u_{\text{U}}\left(c_{\text{U}}\left(s^0_{\text{U}}, s^1_{\text{U}}\right)\right) \\ &\quad + k u_{\text{R}}\left(c_{\text{R}}\left(s^0_{\text{R}}, s^1_{\text{R}}\right)\right) u_{\text{U}}\left(c_{\text{U}}\left(s^0_{\text{U}}, s^1_{\text{U}}\right)\right), \end{aligned}$$

where k *is a constant satisfying:*

$$\begin{aligned} 1 + k u_{\text{R}}(c_{\text{R}}) &> 0, \text{ for all } c_{\text{R}} \in \mathcal{C}_{\text{R}}, \\ \text{and } 1 + k u_{\text{U}}(c_{\text{U}}) &> 0, \text{ for all } c_{\text{U}} \in \mathcal{C}_{\text{U}}. \end{aligned}$$

The second condition is the requirement that nature's decisions over propositions in $\mathbf{P}^0_{\text{R}}$ are independently distributed with respect to the propositions in $\mathbf{P}^0_{\text{U}}$.

Definition 2 (Belief Independence) *For all* $s^0_{\text{R}} \in S^0_{\text{R}}$ *and all* $s^0_{\text{U}} \in S^0_{\text{U}}$*,*

$$\sigma^0\left(s^0_{\text{R}} + s^0_{\text{U}}\right) = \sigma^0_{\text{R}}\left(s^0_{\text{R}}\right)\sigma^0_{\text{U}}\left(s^0_{\text{U}}\right).$$

To see that these two conditions are jointly sufficient, it is enough to show that for any $\hat{s}^1_{\text{R}}$ in S^1_{R} that satisfies (1) we have

$$\max_{s^1 \in S^1} \sum_{s^0 \in S^0} u\left(c\left(s^0, s^1\right)\right)\sigma^0\left(s^0\right) = \max_{s^1_{\text{U}} \in S^1_{\text{U}}} \sum_{s^0 \in S^0} u\left(c\left(s^0, \hat{s}^1_{\text{R}} + s^1_{\text{U}}\right)\right)\sigma^0\left(s^0\right). \quad (2)$$

For the case $k = 0$, we only require **Multiplicative Separable Utility,** since we have

$$\begin{aligned} &\max_{s^1 \in S^1} \sum_{s^0 \in S^0} u\left(c\left(s^0, s^1\right)\right)\sigma^0\left(s^0\right) \\ =\ &\max_{s^1_{\text{U}} \in S^1_{\text{U}}} \sum_{s^0_{\text{U}} \in S^0_{\text{U}}} \max_{s^1_{\text{R}} \in S^1_{\text{R}}} \sum_{s^0_{\text{R}} \in S^0_{\text{R}}} \left[u_{\text{R}}\left(c_{\text{R}}\left(s^0_{\text{R}}, s^1_{\text{R}}\right)\right) + u_{\text{U}}\left(c_{\text{U}}\left(s^0_{\text{U}}, s^1_{\text{U}}\right)\right)\right]\sigma^0\left(s^0_{\text{R}} + s^0_{\text{U}}\right) \\ =\ &\max_{s^1_{\text{U}} \in S^1_{\text{U}}} \sum_{s^0_{\text{U}} \in S^0_{\text{U}}} u_{\text{U}}\left(c_{\text{U}}\left(s^0_{\text{U}}, s^1_{\text{U}}\right)\right)\sigma^0_{\text{U}}\left(s^0_{\text{U}}\right) + \sum_{s^0_{\text{U}} \in S^0_{\text{U}}} \max_{s^1_{\text{R}} \in S^1_{\text{R}}} u_{\text{R}}\left(c_{\text{R}}\left(s^0_{\text{R}}, s^1_{\text{R}}\right)\right)\sigma^0_{\text{R}}\left(s^0_{\text{R}}\right), \end{aligned}$$

and so (2) holds as required.

For the case $k \neq 0$, set

$$\begin{aligned} &\hat{u}_{\text{R}}(c_{\text{R}}) := 1 + k u_{\text{R}}(c_{\text{R}}) \\ \text{and}\qquad &\hat{u}_{\text{U}}(c_{\text{U}}) := 1 + k u_{\text{U}}(c_{\text{U}}). \end{aligned}$$

Now,

$$\begin{aligned}
&\max_{s^1\in S^1}\sum_{s^0\in S^0} u\left(c\left(s^0,s^1\right)\right)\sigma^0\left(s^0\right)\\
=\ &\max_{s^1_{\mathrm{U}}\in S^1_{\mathrm{U}}}\sum_{s^0_{\mathrm{U}}\in S^0_{\mathrm{U}}}\max_{s^1_{\mathrm{R}}\in S^1_{\mathrm{R}}}\sum_{s^0_{\mathrm{R}}\in S^0_{\mathrm{R}}}\left[u_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}},s^1_{\mathrm{R}}\right)\right)+u_{\mathrm{U}}\left(c_{\mathrm{U}}\left(s^0_{\mathrm{U}},s^1_{\mathrm{U}}\right)\right)\right.\\
&\left.+ku_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}},s^1_{\mathrm{R}}\right)\right)u_{\mathrm{U}}\left(c_{\mathrm{U}}\left(s^0_{\mathrm{U}},s^1_{\mathrm{U}}\right)\right)\right]\sigma^0\left(s^0_{\mathrm{R}}+s^0_{\mathrm{U}}\right).
\end{aligned}$$

For $k>0$, we can divide through by $1/k$ and we obtain,

$$\begin{aligned}
=\ &\frac{1}{k}\left(\max_{s^1_{\mathrm{U}}\in S^1_{\mathrm{U}}}\sum_{s^0_{\mathrm{U}}\in S^0_{\mathrm{U}}}\max_{s^1_{\mathrm{R}}\in S^1_{\mathrm{R}}}\sum_{s^0_{\mathrm{R}}\in S^0_{\mathrm{R}}}\left[ku_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}},s^1_{\mathrm{R}}\right)\right)+ku_{\mathrm{U}}\left(c_{\mathrm{U}}\left(s^0_{\mathrm{U}},s^1_{\mathrm{U}}\right)\right)\right.\right.\\
&\left.\left.+k^2u_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}},s^1_{\mathrm{R}}\right)\right)u_{\mathrm{U}}\left(c_{\mathrm{U}}\left(s^0_{\mathrm{U}},s^1_{\mathrm{U}}\right)\right)\right]\sigma^0\left(s^0_{\mathrm{R}}+s^0_{\mathrm{U}}\right)\right)\\
=\ &\frac{1}{k}\left[\max_{s^1_{\mathrm{U}}\in S^1_{\mathrm{U}}}\sum_{s^0_{\mathrm{U}}\in S^0_{\mathrm{U}}}\max_{s^1_{\mathrm{R}}\in S^1_{\mathrm{R}}}\sum_{s^0_{\mathrm{R}}\in S^0_{\mathrm{R}}}\left[\hat{u}_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}},s^1_{\mathrm{R}}\right)\right)\hat{u}_{\mathrm{U}}\left(c_{\mathrm{U}}\left(s^0_{\mathrm{U}},s^1_{\mathrm{U}}\right)\right)\right]\sigma^0_{\mathrm{R}}\left(s^0_{\mathrm{R}}\right)\sigma^0_{\mathrm{U}}\left(s^0_{\mathrm{U}}\right)\right]\\
=\ &\frac{1}{k}\left[\max_{s^1_{\mathrm{U}}\in S^1_{\mathrm{U}}}\sum_{s^0_{\mathrm{U}}\in S^0_{\mathrm{U}}}\hat{u}_{\mathrm{U}}\left(c_{\mathrm{U}}\left(s^0_{\mathrm{U}},s^1_{\mathrm{U}}\right)\right)\left[\max_{s^1_{\mathrm{R}}\in S^1_{\mathrm{R}}}\sum_{s^0_{\mathrm{R}}\in S^0_{\mathrm{R}}}\hat{u}_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}},s^1_{\mathrm{R}}\right)\right)\sigma^0_{\mathrm{R}}\left(s^0_{\mathrm{R}}\right)\right]\sigma^0_{\mathrm{U}}\left(s^0_{\mathrm{U}}\right)\right]\\
=\ &\frac{1}{k}\left[\max_{s^1_{\mathrm{U}}\in S^1_{\mathrm{U}}}\sum_{s^0_{\mathrm{U}}\in S^0_{\mathrm{U}}}\hat{u}_{\mathrm{U}}\left(c_{\mathrm{U}}\left(s^0_{\mathrm{U}},s^1_{\mathrm{U}}\right)\right)\sigma^0_{\mathrm{U}}\left(s^0_{\mathrm{U}}\right)\right]\left[\max_{s^1_{\mathrm{R}}\in S^1_{\mathrm{R}}}\sum_{s^0_{\mathrm{R}}\in S^0_{\mathrm{R}}}\hat{u}_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}},s^1_{\mathrm{R}}\right)\right)\sigma^0_{\mathrm{R}}\left(s^0_{\mathrm{R}}\right)\right],
\end{aligned}$$

and again (2) holds as required.
For $k<0$, we can divide through by $1/k<0$, and change the $\max s$ to $\min s$ which yields,

$$\begin{aligned}
=\ &\frac{1}{k}\left(\min_{s^1_{\mathrm{U}}\in S^1_{\mathrm{U}}}\sum_{s^0_{\mathrm{U}}\in S^0_{\mathrm{U}}}\min_{s^1_{\mathrm{R}}\in S^1_{\mathrm{R}}}\sum_{s^0_{\mathrm{R}}\in S^0_{\mathrm{R}}}\left[ku_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}},s^1_{\mathrm{R}}\right)\right)+ku_{\mathrm{U}}\left(c_{\mathrm{U}}\left(s^0_{\mathrm{U}},s^1_{\mathrm{U}}\right)\right)\right.\right.\\
&\left.\left.+k^2u_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}},s^1_{\mathrm{R}}\right)\right)u_{\mathrm{U}}\left(c_{\mathrm{U}}\left(s^0_{\mathrm{U}},s^1_{\mathrm{U}}\right)\right)\right]\sigma^0\left(s^0_{\mathrm{R}}+s^0_{\mathrm{U}}\right)\right)\\
=\ &\frac{1}{k}\left[\min_{s^1_{\mathrm{U}}\in S^1_{\mathrm{U}}}\sum_{s^0_{\mathrm{U}}\in S^0_{\mathrm{U}}}\min_{s^1_{\mathrm{R}}\in S^1_{\mathrm{R}}}\sum_{s^0_{\mathrm{R}}\in S^0_{\mathrm{R}}}\left[\hat{u}_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}},s^1_{\mathrm{R}}\right)\right)\right]\left[\hat{u}_{\mathrm{U}}\left(c_{\mathrm{U}}\left(s^0_{\mathrm{U}},s^1_{\mathrm{U}}\right)\right)\right]\sigma^0_{\mathrm{R}}\left(s^0_{\mathrm{R}}\right)\sigma^0_{\mathrm{U}}\left(s^0_{\mathrm{U}}\right)\right]\\
=\ &\frac{1}{k}\left[\min_{s^1_{\mathrm{U}}\in S^1_{\mathrm{U}}}\sum_{s^0_{\mathrm{U}}\in S^0_{\mathrm{U}}}\hat{u}_{\mathrm{U}}\left(c_{\mathrm{U}}\left(s^0_{\mathrm{U}},s^1_{\mathrm{U}}\right)\right)\left[\min_{s^1_{\mathrm{R}}\in S^1_{\mathrm{R}}}\sum_{s^0_{\mathrm{R}}\in S^0_{\mathrm{R}}}\hat{u}_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}},s^1_{\mathrm{R}}\right)\right)\sigma^0_{\mathrm{R}}\left(s^0_{\mathrm{R}}\right)\right]\sigma^0_{\mathrm{U}}\left(s^0_{\mathrm{U}}\right)\right]
\end{aligned}$$

$$= \frac{1}{k}\left[\min_{s^1_{\mathrm{U}}\in S^1_{\mathrm{U}}}\sum_{s^0_{\mathrm{U}}\in S^0_{\mathrm{U}}}\hat{u}_{\mathrm{U}}\left(c_{\mathrm{U}}\left(s^0_{\mathrm{U}},s^1_{\mathrm{U}}\right)\right)\sigma^0_{\mathrm{U}}\left(s^0_{\mathrm{U}}\right)\right]\left[\min_{s^1_{\mathrm{R}}\in S^1_{\mathrm{R}}}\sum_{s^0_{\mathrm{R}}\in S^0_{\mathrm{R}}}\hat{u}_{\mathrm{R}}\left(c_{\mathrm{R}}\left(s^0_{\mathrm{R}},s^1_{\mathrm{R}}\right)\right)\sigma^0_{\mathrm{R}}\left(s^0_{\mathrm{R}}\right)\right],$$

and again (2) holds as required.

One example, common in applications, in which preferences admit a multiplicative separable utility representation, is where the consequences are monetary amounts and so can be added together, and the individual's risk preferences over money lotteries exhibit constant absolute risk aversion (CARA).

Example 1 (Monetary Consequences and CARA risk preferences) *Suppose,* $\mathcal{C} = \mathcal{C}_{\mathrm{R}} = \mathcal{C}_{\mathrm{U}} = \mathbb{R}_+$, $u(c) = u_{\mathrm{R}}(c) = u_{\mathrm{U}}(c) = 1 - \exp(-\alpha c)$ *and* $c\left(s^0_{\mathrm{R}} + s^0_{\mathrm{U}}, s^1_{\mathrm{R}} + s^1_{\mathrm{U}}\right) = c_{\mathrm{R}}\left(s^0_{\mathrm{R}}, s^1_{\mathrm{R}}\right) + c_{\mathrm{U}}\left(s^0_{\mathrm{U}}, s^1_{\mathrm{U}}\right)$.

To see that monetary consequences with CARA risk preferences generate multiplicatively separable utility, notice that for $k = -1$,

$$\begin{aligned}
& u_{\mathrm{R}}(c_{\mathrm{R}}) + u_{\mathrm{U}}(c_{\mathrm{U}}) - u_{\mathrm{R}}(c_{\mathrm{R}})\,u_{\mathrm{U}}(c_{\mathrm{U}}) \\
= \; & 1 - \exp(-\alpha c_{\mathrm{R}}) + 1 - \exp(-\alpha c_{\mathrm{U}}) - [1 - \exp(-\alpha c_{\mathrm{R}})]\,[1 - \exp(-\alpha c_{\mathrm{U}})] \\
= \; & 1 - \exp(-\alpha c_{\mathrm{R}})\exp(-\alpha c_{\mathrm{U}}) = 1 - \exp(-\alpha\,[c_{\mathrm{R}} + c_{\mathrm{U}}]) = u(c_{\mathrm{R}} + c_{\mathrm{U}})
\end{aligned}$$

An important point to note here is that discussion in terms of 'bounded rationality' does not imply (though it does not rule out) a focus on heuristics inconsistent with standard decision theory. On the contrary, the limits associated with bounded rationality are even sharper in relation to the application of standard decision theory than in other cases. A heuristic decision process may take account, in some form, of a broad and loosely defined set of considerations more or less relevant to the decision in question. By contrast, the requirements for a formal decision process, beginning with the assignment of a prior probability distribution over the state space, and proceeding to updating posterior decision probabilities and contingent strategies are so demanding that in practice, the number of propositions taken into account is not merely finite (this aspect of boundedness is logically necessary) but commonly quite small. Hence, the derivation of conditions under which a particular small world is appropriate should be of particular interest to highly, but nevertheless boundedly rational decision-makers.

4 Dynamics and Sequentially Consistent Bayesian Decisions

Our analysis so far has focused on the comparisons between a fully aware decision-maker and one with limited awareness. It is straightforward to extend the analysis to the case of a set of decision-makers, ordered in terms of awareness from the least aware to most aware. Such an ordering arises naturally in the course of a history in which individuals become aware of (or discover) new propositions over time, but do not forget propositions of which they are already aware.

More generally, as in Halpern and Rego (2006b) or Grant and Quiggin (2007b) we may consider extensive form games with Nature (or with other players, but we will not pursue the multiagent case further here) in which awareness increases as the game progresses, for example because an individual is presented with a choice she or he had not previously anticipated. In this case, the awareness of the decision-maker at any point in time will, in general, depend on the history of the game up to that time. The possible awareness states of the player are then partially ordered and (since awareness increases along any given history) this ordering is consistent with the ordering of partial histories generated by the dynamic structure of the game.

To put this into the notation introduced above we will consider a special case in which the increase in awareness depends only on calendar time and not the particular history of play up to that point in time. The line of argument developed below, however, can be readily generalized to accommodate history dependent increases in awareness as well.

Fix an objective state-space $\Omega = 2^{\mathbf{P}^0} \times 2^{\mathbf{P}^1}$ with its associated truth valuation $V(p_n^i; \omega)$. Let $\{\mathbf{P}_1^0, \ldots, \mathbf{P}_T^0\}$ and $\{\mathbf{P}_1^1, \ldots, \mathbf{P}_T^1\}$ be T element (ordered) partitions of $\mathbf{P}^0$ and $\mathbf{P}^1$, respectively.[1] Associated with each partition element $\mathbf{P}_t^i$ is a 'partial' strategy set S_t^i where s_t^i in $\mathbf{P}_t^i$ corresponds to the rational

[1]By ordered we simply mean the partitions respect the indexing of propositions in $\mathbf{P}^i$, $i = 0, 1$. That is, if $p_n^i \in \mathbf{P}_t^i$ and $p_{n'}^i \in \mathbf{P}_{t'}^i$, then $t > t'$ implies $n > n'$. In addition, the fact that both partitions have the same number of partitions is without loss of generality since (with slight abuse of notation) we allow some of the partition elements to be empty.

number given by

$$S_t^i := \left\{ s_t^i \in \mathbb{Q} : s_t^i = \sum_{p_n^i \in \mathbf{P}_t^i} 2^{-n} \times V\left(p_n^i; \omega\right), \text{ for some } \omega \in \Omega \right\}, \ i = 0, 1.$$

The interpretation is that at the point in time $t \in \{1, \ldots, T\}$, the individual is aware of nature's propositions $\mathbf{P}_1^0 \cup \mathbf{P}_2^0 \ldots \cup \mathbf{P}_t^0$ and her own propositions in $\mathbf{P}_1^1 \cup \mathbf{P}_2^1 \ldots \cup \mathbf{P}_t^1$. Correspondingly $\left(\sum_{\tau=1}^t S_\tau^0, \sum_{\tau=1}^t S_\tau^1\right)$ is the (restricted) game-form for the game against nature she perceives herself to be playing. Associated with each $t \in \{1, \ldots, T\}$, is a *prior* belief $\sigma^t \in \Delta\left(\sum_{\tau=1}^t S_\tau^0\right)$, a consequence function $c^t : \left(\sum_{\tau=1}^t S_\tau^0\right) \times \left(\sum_{\tau=1}^t S_\tau^1\right) \to \mathcal{C}^t$ where $\mathcal{C}^t$ is the space of consequences of which she is aware at point t, a a utility index $u^t : \mathcal{C}^t \to \mathbb{R}$.

Thus, by repeated application of the results derived above we may derive conclusions concerning the dynamics of choice under conditions of limited, but increasing awareness. If all the models $\left(\sum_{\tau=1}^t S_\tau^0, \sum_{\tau=1}^t S_\tau^1\right)$ are consistent with the complete model $\left(\sum_{\tau=1}^t S_T^0, \sum_{\tau=1}^t S_T^1\right)$ (and therefore with each other) we say that the individual's model of the world is **sequentially consistent**.

The first restriction on the consequences and utility is the multivariate analog of multiplicative separable utility.

Definition 3 (Sequential Multiplicative Separable Utility) *For each $t = 2, \ldots, T$, there exists a consequence space $\mathcal{C}_t$, a consequence function $c_t : S_t^0 \times S_t^1 \to \mathcal{C}_t$, and a utility function $u_t : \mathcal{C}_t \to \mathbb{R}$, s.t. for all $s_1^0 \in S_1^0$, ... all $s_t^0 \in S_t^0$, all $s_1^1 \in S_1^1$, ..., and all $s_t^1 \in S_t^1$,*

$$\begin{aligned} u^t\left(c^t\left(\sum_{\tau=1}^t S_\tau^0, \sum_{\tau=1}^t S_\tau^1\right)\right) &= u^{t-1}\left(c^{t-1}\left(\sum_{\tau=1}^{t-1} S_\tau^0, \sum_{\tau=1}^{t-1} S_\tau^1\right)\right) + u_t\left(c_t\left(s_t^0, s_t^1\right)\right) \\ &\quad + k_t u^{t-1}\left(c^{t-1}\left(\sum_{\tau=1}^{t-1} S_\tau^0, \sum_{\tau=1}^{t-1} S_\tau^1\right)\right) u_t\left(c_t\left(s_t^0, s_t^1\right)\right), \end{aligned}$$

where k_t is a constant satisfying:

$$\begin{aligned} 1 + k_t u^{t-1}\left(c^{t-1}\right) &> 0, \text{ for all } c^{t-1} \in \mathcal{C}^{t-1}, \\ \text{and } 1 + k_t u_t\left(c_t\right) &> 0, \text{ for all } c_t \in \mathcal{C}_t. \end{aligned}$$

The second is simply the multivariate extension of the independence of the prior beliefs over the partial strategy sets of nature.

To state it, notice that $\sigma^T \in \Delta\left(\sum_{\tau=1}^t S_\tau^0\right)$, and the marginal $\sigma_t^T \in \Delta\left(S_t^0\right)$, is given by

$$\sigma_t^T\left(s_t^0\right) = \sum_{\tau \neq t} \sum_{s_\tau^0 \in S_\tau^0} \sigma^T\left(s_t^0 + \sum_{\tau \neq t} s_\tau^0\right).$$

Definition 4 (Multivariate Belief Independence) *1.* $\sigma^1\left(s_1^0\right) \equiv \sigma_1^T\left(s_1^0\right)$. *2. For each* $t = 2, \ldots, T$, *and for all* $s_1^0 \in S_1^0$, *…* *all* $s_t^0 \in S_t^0$, *all* $s_1^1 \in S_1^1$, *…, and all* $s_t^1 \in S_t^1$,

$$\sigma^t\left(\sum_{\tau=1}^{t-1} s_\tau^0 + s_t^0\right) = \sigma^{t-1}\left(\sum_{\tau=1}^{t-1} s_\tau^0\right)\sigma_t^T\left(s_t^0\right).$$

Given the starting point $t = 1$, the individual perceiving (S_1^0, S_1^1) as the game-form of the game she is playing against nature, with associated belief $\sigma^1\left(s_1^0\right) \equiv \sigma_1^T\left(s_1^0\right)$ (by multivariate belief independence), consequence function $c^1 : S_1^0 \times S_1^1 \to \mathcal{C}^1$ and utility $u^1 : \mathcal{C}^1 \to \mathbb{R}$, chooses a strategy $\hat{s}_1^1$ that maximizes her (perceived) subjective expected utility. At point $t = 2$, the individual now is made aware of more of nature's choices and possible her own decisions, making $(S_1^0 + S_2^0, S_1^1 + S_2^1)$, the new game-form, with associated *prior* belief $\sigma^2\left(s_1^0 + s_2^0\right) = \sigma_1^T\left(s_1^0\right)\sigma_2^T\left(s_2^0\right)$ and consequence function $c^2 : (S_1^0 + S_2^0) \times (S_1^1 + S_2^1) \to \mathcal{C}^2$ and utility $u^2 : \mathcal{C}^2 \to \mathbb{R}$. Applying the analysis from the previous section, it readily follows that (sequential multivariate) separable utility along with (multivariate) belief independence means that if the decision maker were selecting her ex ante optimal prior strategy given her new increased level of awareness, she could so by selecting an ex ante strategy $\hat{s}_1^1 + s_2^1$. That is, she need not reoptimize over S_1^1. Moreover, with (multivariate) belief independence she can simply play the *continuation* strategy $\hat{s}_1^1$ entails at point 2 given her planned responses to new information about nature's choice of s_1^0 that is revealed at point 2. And so, on for each point thereafter from point $t = 3$, until point $t = T$, at which point she has become fully aware of the nature of the uncertainty she is facing and her options.

5 Implications and concluding comments

The conditions derived above are quite stringent, which raises the questions of how boundedly rational decision-makers should act. One way to address this question is related to the work of Bordley and Hazen (1992) who consider a single decision-maker and derive conditions similar to those presented above to determine when it would be appropriate to apply expected utility theory in the context of a restricted model similar to $(S^0_{\text{R}}, S^1_{\text{R}})$. Bordley and Hazen argue that, if these conditions are not satisfied, the induced preferences over strategies in $(S^0_{\text{R}}, S^1_{\text{R}})$ may be represented by non-expected utility preferences.

By contrast, the analysis here compares the perspective of boundedly rational decision-makers with that which they would take if they were fully aware. This approach has a range of implications. First, consider the perspective of an external observer, with the possibility of intervening to affect the choices of a boundedly aware decision-maker. Such an intervention might simply involve making the decision-maker aware of some previously unconsidered possibilities, or it might involve actions aimed at encouraging some choices and discouraging others. Under the conditions derived above, increasing the awareness of the decision-maker will not affect decisions, but will add to the decision-makers computational burden and is thus undesirable. Similarly unless the external decision-maker has private and noncommunicable information, as distinct from greater awareness of the possible set of states, intervening directly will reduce the decision-makers welfare. Conversely, where the conditions derived above are not satisfied, intervention may improve welfare.

As is shown by Grant and Quiggin (2007b, Proposition 11, p19) a boundedly aware decision-maker cannot know with certainty that there exist propositions of which they are unaware[2]. On the other hand, inductive reasoning may be used to justify the belief that a given small-world model is incomplete, and that expected utility may need to be modified (leading to non-expected utility or multiple priors models) or supplemented with heuristics derived from experience of decisions made under conditions of bounded awareness

[2]Here knowledge is interpreted in the modal-logical sense appropriate to a state-space model of the world. A proposition is known to be true if it is true in every state of the world that may possibly hold.

(the precautionary principle is a prominent example.of such a heuristic).

In summary, Bayesian decision theory provides an appealing basis for reasoning and choice for an unboundedly rational individual, capable of formulating a prior distribution over all possible events, and updating it in the light of new information. In practice, however, boundedly rational individuals can apply Bayesian reasoning only within 'small worlds' in the sense described by Savage (1954). That is, a boundedly rational Bayesian will define particular subproblems for which she judges that a well-defined prior over relevant states (the projections of events in the larger world) is available, and will then apply Bayesian decision theory to these subproblems. The limitation to small worlds raises the problem of determining conditions under which Bayesian updating is valid, and what response is reasonable if these conditions are not satisfied.

In this paper, we have presented a dynamic model within which both the discovery of new propositions and the updating of probability beliefs takes place over time. Using this model, we have considered independence conditions for newly discovered propositions under which the restricted Bayesian approach to probability updating is valid.

References

[1] Bordley, R. and Hazen, G. (1992), 'Nonlinear utility models arising from unmodelled small world intercorrelation', *Management Science*, 38(7), 1010-17.

[2] Grant, S. and J. Quiggin (2007a): "Conjectures, Refutations and Discoveries: Incorporating New Knowledge in Models of Belief and Choice under Uncertainty" in *Uncertainty and Risk: Mental, Formal, Experimental Representations,* eds M. Abdellaoui, R. Duncan Luce, M. Machina and B. Munier, Springer Verlang, Berlin.

[3] Grant, S. and J. Quiggin (2007b):"Awareness and Discovery," mimeo Rice University, [http://www.ruf.rice.edu/~econ/faculty/Grant/sgrant/GrantQuigginAwareness0708.pdf].

[4] Halpern, J. (2003), *Reasoning About Uncertainty.* MIT Press, Cambridge, MA.

[5] Halpern, J., and L. C. Rêgo (2006a), 'Reasoning About Knowledge of Unawareness', *Proceedings of the Conference on Knowledge Representation and Reasoning.*

[6] Halpern, J., and L. C. Rêgo (2006b), 'Extensive Games with Possibly Unaware Players', Conference on Autonomous Agents and Multiagent Systems.

[7] Heifetz, A., M. Meier, and B. Schipper (2006): "Interactive Unawareness," *Journal of Economic Theory* **130**, 78-94.

[8] Keeney, R.L. and H. Raiffa (1976): *Decisions with Multiple Objectives.* John Wiley & Sons, Inc, New York.

[9] Savage, L.J. (1954): *The Foundations of Statistics.* 1972 Dover Publications Edition, New York.