

SEQUENTIAL TWO-PLAYER GAMES WITH AMBIGUITY¹

Jürgen Eichberger and David Kelsey
Alfred Weber Institut Department of Economics
Universität Heidelberg The University of Birmingham
Germany United Kingdom

February 2004

Abstract

If players' beliefs are strictly non-additive, the Dempster-Shafer updating rule can be used to define beliefs off the equilibrium path. We define an equilibrium concept in sequential two-person games where players update their beliefs with the Dempster-Shafer updating rule. We show that in the limit as uncertainty tends to zero, our equilibrium approximates Bayesian Nash equilibrium. We argue that our equilibrium can be used to define a refinement of Bayesian Nash equilibrium by imposing context-dependent constraints on beliefs under uncertainty.

¹ Financial assistance from the ESRC senior research fellowship scheme, award no. H52427502595, the School of Social Science at the University of Birmingham, the Department of Economics at the University of Melbourne and The British Academy is gratefully acknowledged. For comments we would like to thank Simon Grant, Hans Haller, Youngse Kim, Bart Lipman, Frank Milne, Shasi Nandeibam, Hyun Shin, Peter Sinclair, Willy Spanjers, Martin Summer, Peter Wakker and participants in seminars at Queen's university, the universities of Melbourne and Birmingham, the Econometric Society World meetings in Tokyo 1995 and the ESRC game theory workshop, Kenilworth 1997. We thank two anonymous referees of this journal for detailed comments which helped to improve the the final version.

JEL Classification: C72, D81

Keywords: Uncertainty-aversion, capacity, Dempster-Shafer rule, bargaining, signalling game.

Running Head: SEQUENTIAL GAMES WITH AMBIGUITY

Address:

Professor Dr. Jürgen Eichberger

Wirtschaftstheorie I

Alfred Weber Institut

RUPRECHT-KARLS-UNIVERSITÄT HEIDELBERG

Grabengasse 14

D-69117 HEIDELBERG

GERMANY

1. INTRODUCTION

Economists have made a distinction between risk (where probabilities are objectively known) and ambiguity (where probabilities are unknown). Until recently it was not clear how to model this formally. Schmeidler (1989) has proposed an axiomatic decision theory, which is able to model ambiguity. In this theory, the decision-maker's beliefs are represented by a capacity (non-additive subjective probability) and (s)he is modelled as maximising the expected value of utility with respect to the capacity. Ambiguity is represented by strictly non-additive capacities. The expectation is expressed as a Choquet integral (Choquet, 1953-4). Schmeidler's theory will henceforth be referred to as Choquet expected utility (CEU).

A number of researchers have applied CEU (or related theories) to games². Most of these papers consider strategic (normal) form games. Lo (1999) suggests an equilibrium concept for extensive form games under ambiguity. Since he uses the related multiple-prior expected utility theory to model ambiguity, he discusses a number of conceptual problems which arise in the context of dynamic games if players face strategic ambiguity. The paper contains many instructive examples but no general theorems about existence of equilibrium if players face ambiguity.

Approaching the problem of extensive form games in a very general way, Lo (1999) cannot exploit one of the strengths of non-additive probabilities, namely that unlike additive probabilities, they can be updated after events with a capacity value of zero. In the present paper, we apply CEU to sequential two-player games. This class of extensive form games

² See Dow and Werlang (1994), Eichberger and Kelsey (2000), Hendon et al. (1994), Klibanoff (1996), Lo (1996) and Marinacci (2000).

comprises many important game-theoretic models in economics such as signalling games, two-stage industrial organisation models or bargaining problems. Restricting oneself to this class of games allows us to ignore some of the consistency problems encountered in Lo (1999).

In extensive form games, updating of beliefs on newly acquired information is important. If beliefs are represented by additive probability distributions, then Bayesian updating is the natural method to incorporate the information obtained from the observed moves of the opponents. Bayesian updating however is possible only at information sets which have a positive probability of being reached. As is well-known, play at information sets off the equilibrium path can have a major effect on the equilibrium itself. Thus it is important to determine players' beliefs at such information sets. Because Bayesian updating puts no restrictions on such beliefs, a multiplicity of equilibria is compatible with Bayesian beliefs.

Games with incomplete information are usually plagued by a large number of Bayes-Nash equilibria. Signalling games in particular have typically an excessively large number of equilibria because the signal space is large compared to the type space, which implies that most actions will not be observed in equilibrium. The multiplicity of equilibria depends on the lack of constraints on out-of-equilibrium strategies. There is a huge literature in game-theory which tries to impose further constraints on beliefs by additional rules about how a player should interpret out-of-equilibrium moves in equilibrium. Such constraints on beliefs refine the set of Bayes-Nash equilibria. Compare Mailath (1992) for a survey of refinements in the context of signalling games. Most refinements have been based on forward or backward induction arguments. A common criticism of such arguments is that, if

the initial situation is indeed an equilibrium, then players should conclude from a deviation that the opponent is not rational or does not understand the structure of the game.

In this paper, we propose a definition of equilibrium where players have non-additive beliefs and use an updating rule proposed by Dempster and Shafer in the literature for capacities. This equilibrium notion comprises Bayes-Nash equilibrium as a special case. The Dempster-Shafer updating rule, which is part of our equilibrium concept, has well-defined updated capacities off the equilibrium path as long as there is ambiguity. Capacities can be further constrained by adequate assumptions about beliefs without affecting consistency of beliefs in an equilibrium under ambiguity. Hence, there is room to put constraints on beliefs which may be specific to situation one wants to model. For example, one can exogenously determine the degree of ambiguity or one can restrict beliefs to agree with an additive prior distribution, if one wants to model a situation where players are completely confident about the distribution of types but ambiguous about their opponents' strategic behaviour. It is possible to control for the ambiguity of a situation in experiments in order to see how it affects decision behaviour. For individual decision situations, such experiments have been performed (Camerer and Weber, 1992). There are few experiments so far, which focus on strategic ambiguity, but we are confident that such tests can be conducted.

One can parametrise the notion of ambiguity and demonstrate existence of equilibrium for any exogenously determined level of ambiguity. This opens up the possibility to study sequences of equilibria under ambiguity which converge to a Bayesian equilibrium as ambiguity vanishes. Assumptions about beliefs under ambiguity will determine which Bayesian equilibrium will be selected. An interesting aspect of this approach is, even in a Bayesian

equilibrium, beliefs off the equilibrium path may be represented by capacities which are not additive. The greater freedom of modelling beliefs under ambiguity provides a novel and useful modelling device for economic applications.

In section 2 we introduce the CEU model and demonstrate some properties of CEU and the Dempster–Shafer updating rule. Section 3 introduces our solution concept for two-stage games under ambiguity and relates it to some existing solution concepts. Section 4 studies limits of sequences of equilibria as ambiguity vanishes. We show that ambiguous beliefs can select among Bayesian equilibria. Section 6 concludes the paper. All proofs are gathered in an appendix.

2. CEU PREFERENCES AND DS-UPDATING

In this section we consider a finite set S of states of nature. A subset of S is referred to as an event. The set of possible outcomes or consequences is denoted by X . An act is a function from S to X . The space of all acts is denoted by $A(S) := \{a \mid a : S \rightarrow X\}$. The decision-maker’s preferences over $A(S)$ are denoted by \succsim .

A capacity or non-additive probability on S is a real-valued function ν on the subsets of S , which satisfies the following properties:

- (i) $A \subseteq B$ implies $\nu(A) \leq \nu(B)$;
- (ii) $\nu(\emptyset) = 0$, $\nu(S) = 1$.

The capacity is said to be convex if for all $A, B \subseteq S$, $\nu(A \cup B) \geq \nu(A) + \nu(B) - \nu(A \cap B)$. Representing beliefs by a convex capacity is compatible with experimental evidence (see Camerer und Weber, 1992) and is commonly used in applications of CEU to model ambiguity averse behaviour. *We shall assume that all capacities are convex.*

We shall use capacities to represent the beliefs of players. In game-theoretic applications, the opponents' strategy combinations will be the relevant states for a player. It is possible to define an expected value with respect to a capacity to be a Choquet integral (Choquet, 1955).

For any function $\phi : S \rightarrow \mathbb{R}$ and any outcome $x \in X$ let $\bar{B}(x|\phi) := \{s \in S : \phi(s) \geq x\}$ be the event in which ϕ is greater than or equal to x . Similarly, denote by $B(x|\phi) := \{s \in S : \phi(s) > x\}$ the event in which ϕ produces a strictly greater outcome than x . The Choquet integral of ϕ with respect to the capacity ν is defined as

$$(1) \quad \int \phi \, d\nu := \sum_{x \in \phi(S)} x \cdot [\nu(\bar{B}(x|\phi)) - \nu(B(x|\phi))],$$

where the summation is over the range of the act, $\phi(S) := \{x \in X \mid \exists s \in S, \phi(s) = x\}$.

We shall assume that preferences may be represented by Choquet expected utility (CEU) with respect to a capacity, i.e.

$$a \succcurlyeq b \Leftrightarrow \int u(a(s)) \, d\nu(s) \geq \int u(b(s)) \, d\nu(s).$$

Such preferences have been axiomatised by Gilboa (1987) and Sarin and Wakker (1992).

Definition 2.1 *The degree of ambiguity of capacity ν is defined by*

$$\lambda(\nu) := 1 - \min_{A \subseteq S} (\nu(A) + \nu(S \setminus A)).$$

This definition is adapted from Dow and Werlang (1992). It has been justified epistemologically by Mukerji (1997). The degree of ambiguity is a measure of the deviation from additivity. For an additive probability $\lambda(\nu) = 0$, while for complete ambiguity $\lambda(\nu) = 1$. The following result confirms that the degree of ambiguity is a reasonable measure of de-

viation from additivity, for convex capacities³.

Lemma 2.2 *If a convex capacity ν has zero degree of ambiguity then it is additive.*

2.1 The support of a capacity

In game theory, players are assumed to maximise their expected payoffs. Strategy choices are considered in equilibrium if beliefs are consistent with actual behaviour. The strongest form of consistency, Nash equilibrium, requires players' beliefs to coincide with their actual behaviour. In an alternative and equivalent definition of Nash equilibrium the strategies in the support of the opponents' beliefs about a player's behaviour must be best responses of that player. In other words, players expect their opponents to play only best-response strategies.

If decision makers' ambiguity is modelled by capacities then there are several concepts of a support which all coincide with the usual notion of support in the case of additive capacities. In this paper we will use the following definition.

Definition 2.3 *A support of a capacity ν is an event $A \subseteq S$ such that $\nu(S \setminus A) = 0$ and $\nu(S \setminus B) > 0$, for all events $B \subset A$.*

This definition of the support is due to Dow and Werlang (1994). Above we define the support of a capacity to be a minimal set whose complement has a capacity value of zero. This is equivalent to the usual definition of support (i.e. a minimal set of probability one) for an additive capacity but will generally yield a smaller set if the capacity is not additive.

³ The lemma is false if convexity is not assumed. A counter-example would be the class of symmetric capacities studied by Gilboa (1989) and Nehring (1994).

With this support notion every capacity has a support. However it has been criticised because the support is not necessarily unique and states outside the support may affect decision making if a bad outcome occurs on them. In Eichberger and Kelsey (2001) we provide an extensive discussion of various support notions for capacities suggested in the literature⁴. In particular, we show that the support is unique if and only if one requires in addition $\nu(B) > 0$, for all events B in the support. Adding this requirement to Definition 2.3 guarantees a unique support but there are convex capacities for which no such support exists. In game-theoretic applications, the lack of uniqueness poses no problem because we show the existence of an equilibrium in which beliefs have a unique support. Moreover, our results and examples all have unique supports, which satisfy this additional restriction.

More substantial is the objection to Definition 2.3 that states outside the support are not Savage-null. An event E is Savage-null if outcomes on E never affect a decision, i.e. if $a_{Ec} \sim b_{Ec}$ for all acts a, b, c , where a_{Ec} denotes an act which yields $a(s)$ for all states in E and $c(s)$ in all other states. We believe that this argument is not appropriate in game-theoretic applications. We will argue this case below in context with the game-theoretic equilibrium concept, which we advance in this paper.

2.2 CEU preferences and DS-updating

In sequential games players may receive information about the opponents by observing their moves in earlier stages of the game. In particular in signalling games, second-stage players will try to infer information about characteristics of their opponents from the signals which they receive. It is therefore necessary to specify a rule for how to revise beliefs represented

⁴ Haller (2000), Marinacci (2000) and in particular Ryan (1998) discuss and argue for other support notions.

by capacities in the light of new information.

If beliefs are additive, Bayes' rule is the unique updating rule which maintains additivity. As in the case of the support, with non-additive capacities there are several updating procedures, which all coincide with Bayesian updating in the case of additivity. Gilboa and Schmeidler (1993) provide an exposition and an axiomatic treatment from a behavioural perspective. In this paper we choose the Dempster-Shafer belief revision rule (see Shafer, 1976).

Definition 2.4 Dempster-Shafer revision

The Dempster-Shafer revision of capacity ν given event $E \subset S$ is

$$(2) \quad \nu(F|E) := \frac{\nu((F \cap E) \cup (S \setminus E)) - \nu(S \setminus E)}{1 - \nu(S \setminus E)}.$$

The axiomatisation by Gilboa and Schmeidler (1993) shows that the Dempster-Shafer rule (DS-rule) may be interpreted as a pessimistic updating rule. If one views capacities as constraints on a set of additive probability distributions then it is equivalent to a maximum likelihood updating procedure.

For extensive-form games the DS-rule is particularly interesting, since it may be defined even when $\nu(E) = 0$. If the event E , about which the decision maker obtains information, was ambiguous, $\nu(E) + \nu(S \setminus E) < 1$, then the DS-rule will be well-defined even if it has a prior capacity value of zero. Thus, it may be possible to update non-additive beliefs on events with a capacity value of zero. We will argue in Section 4 that this property of DS-updating provides an approach to equilibrium selection based on ambiguity of players. Contrary to the refinements of Bayes-Nash equilibrium based on second-order reasoning

about out-of-equilibrium moves, which dominate the literature, ambiguity-related refinements can be given a behavioural content which is independent of the equilibrium notion.

3. SEQUENTIAL TWO-PLAYER GAMES

In this paper we will consider two-player games with complete and incomplete information, where players move sequentially. Without loss of generality, we will assume throughout that player 1 moves first and that player 2 knows the move of player 1 when she makes her move. Player 1 may have one of several types which are described by a finite set T . If T contains a single type the game has complete information, otherwise it is of incomplete information. Beliefs about types will be described below. Both players choose actions from finite action sets A^i , $i = 1, 2$. Their payoffs are described by the utility functions $u^i(s, a, t)$, $i = 1, 2$.

Strategies of player 1 coincide with actions, $S^1 := A^1$. In contrast, player 2 who observes the action of player 1 can condition her moves on this observation. Hence, $S^2 := \{s^2 \mid s^2 : S^1 \rightarrow A^2\}$ denotes the strategy set of player 2, which is also finite because the action sets of both players are finite. We will denote by $s^2(s^1) \in A^2$ the action, which player 2 will choose in response to s^1 according to her strategy s^2 .

Both players hold beliefs about the opponent's behaviour which are represented by convex capacities. Player 1 has beliefs ν^2 about the strategies in S^2 , which player 2 will choose. A belief ν^2 on S^2 of player 1 about player 2's strategy induces a set of beliefs $\{\tilde{\nu}_j^2 \mid s_j^1 \in S^1\}$ about the actions in A^2 , which player 2 will choose in response to a strategy $s_j^1 \in S^1$: $\tilde{\nu}_j^2(E) := \nu^2(\{s^2 \in S^2 \mid s^2(s_j^1) \in E\})$. For notational convenience we will state definitions and results in terms of ν^2 though the respective statements translate easily into

statements about the set of beliefs about actions $\{\tilde{\nu}_j^2 | s_j^1 \in S^1\}$.

Player 2 has to form beliefs about strategic behaviour of the possible types of player 1. Beliefs of player 2 about player 1's type-contingent strategy choices are represented by the capacity ν^1 on $S^1 \times T$. These beliefs represent jointly this player's ambiguity about type and strategy⁵. In game-theoretic applications it is usually assumed that prior beliefs about types are common knowledge and additive. We can constrain ν^1 to be compatible with an exogenous prior distribution over types, which is represented by a (possibly additive) capacity μ on the type space T .

Definition 3.1 *A capacity ν on $S^1 \times T$ agrees with the capacity μ on T if, for any subset T' of T , $\nu(S^1 \times T') = \mu(T')$.*

Whether the prior belief of player 2, ν^1 , agrees with a prior distribution on types or not, once player 2 observes the action s^1 , which player 1 chooses she will have to revise her beliefs in the light of this information. To simplify notation, we will write

$$(3) \quad \nu^1(T'|s^1) \quad : \quad = \nu^1(S^1 \times T' | \{s^1\} \times T) \\ = \frac{\nu^1(\{\{s^1\} \times T'\} \cup ((S^1 \setminus \{s^1\}) \times T)) - \nu^1((S^1 \setminus \{s^1\}) \times T)}{1 - \nu^1((S^1 \setminus \{s^1\}) \times T)}$$

to denote the DS-updated capacity of the event $T' \subset T$ if the action s^1 has been observed. This DS-update is well-defined if $\nu^1((S^1 \setminus \{s^1\}) \times T) < 1$ holds. This condition is satisfied if

⁵ Here we follow the approach of Milgrom and Roberts (1986) for the representation of games under incomplete information.

- either player 2 feels ambiguity about player 1's choice of the strategy s^1 , i.e.,
the event $\{s^1\} \times T$ is ambiguous, i.e. if $\nu^1(\{s^1\} \times T) + \nu^1((S^1 \setminus \{s^1\}) \times T) < 1$,
- or player 2 is certain that player 1 will play strategy s^1 with positive probability, i.e.,
the event $\{s^1\} \times T$ is unambiguous and $\nu^1(\{s^1\} \times T) > 0$.

Finally, let

$$(4) \quad P^1(s^1|t, \nu^2) := \int u^1(s^1, s^2(s^1), t) d\nu^2$$

and, for a capacity ρ on T ,

$$(5) \quad P^2(a|s^1, \rho) := \int u^2(s^1, a, t) d\rho$$

be the CEU-payoff of player 1 and 2 respectively. The belief ρ about player 1's type will either be ν^1 or $\nu^1(\cdot|s^1)$, depending on whether beliefs are formed by DS-updating or not.

3.1 Dempster-Shafer equilibrium

In order to see the relationship between the Dempster-Shafer equilibrium concept, which will be proposed in this section, with the familiar notion of a Perfect Bayesian Equilibrium, consider the case of additive capacities π^1 and π^2 representing players' beliefs.

A *Perfect Bayesian Equilibrium (PBE)* with prior distribution p on T is

- a probability distribution π^1 on $S^1 \times T$ such that $\sum_{s^1 \in S^1} \pi(s^1, t) = p(t)$ for all $t \in T$,
- a probability distribution π^2 on S^2 ,
- and a family of probability distributions $\{\mu(\cdot|s^1)\}_{s^1 \in S^1}$ on T such that:

$$\begin{aligned}
\text{a) } & (s^1, t) \in \text{supp } \pi^1 && \text{then } s^1 \in \arg \max_{\tilde{s}^1 \in S^1} \sum_{s^2 \in S^2} \pi^2(s^2) \cdot u^1(\tilde{s}^1, s^2(\tilde{s}^1), t), \\
\text{b) } & s^2 \in \text{supp } \pi^2 && \text{then } s^2(s^1) \in \arg \max_{\tilde{a} \in A^2} \sum_{t \in T} \mu(t|s^1) \cdot u^2(s^1, \tilde{a}, t), \\
& && \text{for all } s^1 \in S^1, \\
\text{c) } & \mu(t|s^1) = \frac{\pi(s^1, t)}{\sum_{t \in T} \pi(s^1, t)} && \text{if } \sum_{t \in T} \pi(s^1, t) > 0.
\end{aligned}
\tag{PBE}$$

The standard interpretation is that $\pi^1(s^1|t) := \pi(s^1, t)/p(t)$, $\pi^2(a|s^1) := \pi^2(\{s^2 \in S^2 \mid s^2(s^1) = a\})$ are behaviour strategies and $\mu(\cdot|s^1)$ are beliefs at the information set reached after move s^1 . Behaviour strategies are identified with beliefs of players. Note that the belief interpretation requires π^1 to be the belief of player 2 about the behaviour of player 1 and vice versa.

Without the family of probability distributions $\{\mu(\cdot|s^1)\}_{s^1 \in S^1}$, conditions PBE-a and PBE-b define a *Bayesian equilibrium* if one uses $\pi(s^1, t)$ instead of $\mu(t|s^1)$ in PBE-b. Obviously, all PBE are Bayesian equilibria. All a PBE requires in addition to the conditions of a Bayesian equilibrium is optimality of behaviour at information sets off the equilibrium path, i.e., after moves s^1 such that $\sum_{t \in T} \pi(s^1, t) = 0$, with respect to some arbitrary additive belief $\mu(t|s^1)$. Hence, it only rules out strictly dominated actions at such information sets. Since beliefs at information sets off the equilibrium path are arbitrary, there are usually many PBE depending on the beliefs $\mu(t|s^1)$ assumed at information sets off the equilibrium path.

This multiplicity of PBE poses a serious problem in games, where there are few types of player 1 and many more strategies of this player, as it is typically the case in signalling

games. The literature is therefore particularly rich in refinements for signalling games. Mailath (1992) provides a good survey of the refinements applied in the context of signalling games. We show below in Example 4.2 how with ambiguity aversion, a quite natural assumption on beliefs can reduce this multiplicity of equilibria.

Games in strategic form where the beliefs of players are represented by capacities have been studied by Dow and Werlang (1994), Marinacci (2000) and Eichberger and Kelsey (2000). In the strategic form, beliefs (ν^1, ν^2) form an equilibria under uncertainty if these beliefs have supports containing only pure strategies which are optimal for the respective player given these beliefs.

In this paper we study sequential two-player games where the action of player 1 conveys information to player 2. This requires a reformulation of the equilibrium concept. In contrast to the equilibrium notion in the strategic form, we will require that player 2's strategy consists of actions that are optimal at each information set, i.e. after observing the action of player 1.

Definition 3.2 Dempster-Shafer Equilibrium (DSE)

A Dempster-Shafer equilibrium consists of capacities ν^1 on $S^1 \times T$, ν^2 on S^2 and a family of capacities $\{\rho(\cdot|s^1)\}_{s^1 \in S^1}$ on T such that there are supports $\text{supp } \nu^1$ and $\text{supp } \nu^2$ satisfying:

- a) $(s^1, t) \in \text{supp } \nu^1$ then $s^1 \in \arg \max_{\tilde{s}^1 \in S^1} P^1(\tilde{s}^1|t, \nu^2)$,
- b) $s^2 \in \text{supp } \nu^2$ then $s^2(s^1) \in \arg \max_{\tilde{a} \in A^2} P^2(\tilde{a}|s^1, \rho(\cdot|s^1))$ (DSE)
for all $s \in S^1$,
- c) $\rho(T'|s^1) = \nu^1(T'|s^1)$ if $\nu^1((S^1 \setminus \{s^1\}) \times T) < 1$.

In general, the support of a capacity need not be unique. Hence, a DSE requires us to

specify a set of beliefs and some associated support. In most applications the capacities considered will have a unique support and this seemingly arbitrary choice of support for a given capacity poses no problem⁶. Condition DSE-a of Definition 3.2 guarantees that only optimal type-contingent strategies of player 1 will be included in the support of player 2's beliefs. Similarly, by condition DSE-b there are only strategies in the support of player 1's beliefs which prescribe optimal behaviour after a strategy of player 1. We call the equilibrium a Dempster-Shafer equilibrium (DSE) because, according to DSE-c, beliefs of player 2 about the type of player 1 are obtained by the DS-updating rule.

Remark: A Dempster-Shafer equilibrium is defined as a set of beliefs over type-contingent strategies (ν^1, ν^2) and a set of updated beliefs after strategy choices of player 1, $\{\rho(\cdot|s^1)\}_{s^1 \in S^1}$. By Condition DSE-c the updated beliefs $\{\rho(\cdot|s^1)\}_{s^1 \in S^1}$ are however a derived concept. When referring to a DSE we therefore often mention only (ν^1, ν^2) , if there is no danger of confusion. ■

The degree of ambiguity $\lambda(\nu^i)$, formally defined in the previous section, is a property of the equilibrium beliefs (ν^1, ν^2) . We will demonstrate below in Proposition 3.5 that one can take this degree of ambiguity as an exogenous parameter and deduce equilibrium beliefs. The DSE concept is useful for economic applications because one can study games under different degrees of ambiguity. Nash equilibrium is a special case of an equilibrium under no ambiguity. Since the DS-updated capacity $\rho(\cdot|s^1)$ is a derived concept its degree of

⁶ There are, however, interesting cases where supports are not unique, yet where there is only one support for each capacity which is consistent with the optimality conditions DSE-a and DSE-b.

ambiguity $\lambda(\nu^1(\cdot|s^1))$ is also a derived property⁷.

We define the degree of ambiguity of a game to be the maximal degree of ambiguity of the equilibrium beliefs. Formally, we will say that

- a Dempster-Shafer equilibrium (ν^1, ν^2) has degree of ambiguity $\lambda \in [0, 1]$ if $\lambda := \max\{\lambda(\nu^1), \lambda(\nu^2)\}$, and
- a Dempster-Shafer equilibrium (ν^1, ν^2) agrees with the additive prior distribution p on T if $\nu^1(S^1 \times \{t\}) = p(t)$ for all $t \in T$.

If the beliefs of a DSE are additive, i.e., if there is no ambiguity, and if there is a common prior distribution over types then a DSE is a Bayesian equilibrium.

The following proposition relates the Dempster-Shafer equilibrium concept to the Bayesian equilibrium and the perfect Bayesian equilibrium notions.

Proposition 3.3

- a) A Dempster-Shafer Equilibrium (ν^1, ν^2) with a degree of ambiguity $\lambda = 0$, for which the belief of player 2, ν^1 , agrees with the additive prior distribution p on T , is a *Bayesian Equilibrium*.
- b) Consider a Dempster-Shafer Equilibrium (ν^1, ν^2) with a degree of ambiguity $\lambda = 0$, for which the belief of player 2, ν^1 , agrees with the additive prior distribution p on T . If for each strategy $s^1 \in S^1$ there exists a type $t \in T$ such that $(s^1, t) \in \text{supp } \nu$, then the Dempster-Shafer Equilibrium (ν^1, ν^2) is a *perfect Bayesian equilibrium*.

Discussion of the DSE concept. Proposition 3.3 shows that Bayesian equilibrium and

⁷ For specific types of capacities one can prove implications for the degree of ambiguity of their DS-update. Compare Eichberger and Kelsey (1999).

perfect Bayesian equilibrium are special cases of a DSE if there is no ambiguity. There is mounting experimental evidence that Nash equilibrium and its refinements do not yield good predictions of actual behaviour in all games. Situations in which one finds consistent deviations from the Nash equilibrium hypothesis include bargaining, e.g. the ultimatum bargaining experiments (Roth, 1995), coordination problems (Ochs, 1995), public goods provision (Ledyard, 1995), and signalling games (Brandts and Holt, 1992 and 1993). These findings pose a challenge to theory and call for the investigation of modified equilibrium concepts. Some of these anomalies can be explained by altruistic preferences, e.g. in the cases of public goods and bargaining. Even in these cases however, it is difficult to account for all the observed phenomena by modifying preferences alone.

The approach we propose in this paper focuses on ambiguity about the behaviour of the opponent players. We do not give up the idea that players maximise their expected payoff but we investigate how ambiguity about the strategic behaviour of opponents affects the equilibria of games. Ambiguity about the opponents' behaviour may arise by a number of reasons.

Traditional game theory maintains that players deduce beliefs about the opponents' behaviour from firm knowledge about the preferences of opponents. In games with incomplete information one replaces knowledge about other players' payoffs, with knowledge about the probability distribution over possible types of payoffs. Though logically sound and completely consistent, modelling games of incomplete information by probability distributions over type spaces assumes that players have extremely high computational abilities.

Ambiguity about the strategic behaviour of the opponent players reflects the difficulty of settling one's beliefs firmly on a particular probability distribution over types and their behaviour. This does not imply that players do not care about the motivation of the opponents or that they do not consider the possibility of different types of opponents. It means however that their behaviour may be influenced by the fact that they do not feel certain about such inferences.

Ambiguous beliefs represented by capacities, allow us to model players who hold and process information about their opponents in order to predict their behaviour but who, depending on the situation, may feel more or less certain about these predictions. If ambiguity is about two or more possible characteristics of an opponent then beliefs should be modelled by a capacity over the relevant type space, if ambiguity concerns the correct description of the situation in general it is best modelled by ambiguity about the opponents' strategy choices. Equilibrium concepts for ambiguous players, like the DSE suggested here, provide also a unified framework, in which the completely consistent beliefs of Nash equilibrium analysis as well as behaviour influenced by ambiguity can be analysed.

3.1.1 Ambiguity about the strategy choice of the opponents

In traditional game-theoretic reasoning, players trust completely their reasoning about the rationality of their opponents. If players believe that an opponent's strategy is strictly dominated, then they will act on the presumption that this player will never choose this strategy. Similarly, strategies of the opponents which are not in the support of the capacity representing a player's belief should not influence this player's behaviour.

The following example illustrates that these properties are not true for DS-equilibria.

In Example 3.1 there are two DS-equilibria, one which describes behaviour similar to the backward-induction Nash equilibrium. The second DSE shows that strategies of the opponent with bad outcomes may influence the decision of a player, even if the opponents strategy is strictly dominated and not in the support of this player's beliefs.

Example 3.1 Frivolous lawsuits⁸

BECHUK (1988) studies legal disputes where the plaintiff threatens to go to court even if the expected value of the court case is negative in order to extract a settlement offer from the defendant. Figure 1 represents a stylised version of this situation.

Insert FIGURE 1 here

Once the potential plaintiff has threatened to go to court the defendant, D , can make an settlement offer o which will be accepted or refuse to make an offer no , in which case the plaintiff, P , has to decide whether to drop the case, d , or to go to court, c . The payoffs reflect the incentives of the players. If the defendant makes no offer, no , and the potential plaintiff decides not to file the suit, d , both players receive a payoff of 1. If the plaintiff goes to court, c , both players obtain -1 , which reflects the negative expected value of the court case. A settlement offer, which is accepted, yields the plaintiff a payoff of 3 and the defendant a payoff of 0. The settlement yields the plaintiff a higher payoff than not going to court, the incentive for the frivolous suit⁹.

This is a game with complete information where player 1, the defendant, has a single type t , the notation of which is suppressed. Hence, $(\rho(t|no), \rho(t|o)) = (1, 1)$ in any DSE. Equilibrium beliefs for the two types of DS-equilibria of this game are given in the following

⁸ The structure of this game corresponds to the well-known entry game in Industrial Organisation.

⁹ This example is a slightly modified and parametrised version of the model in Bebchuk (1988), p. 441.

table.

DS-equilibria $\{(\nu_1^D, \nu_1^P), (\nu_2^D, \nu_2^P)\}$

event	$\nu_1^D(\cdot)$	$\nu_2^D(\cdot)$	event	$\nu_1^P(\cdot)$	$\nu_2^P(\cdot)$
$\{no\}$	$\alpha^D > 0$	0	$\{d\}$	$\alpha^P \geq \frac{1}{2}$	$\alpha^P < \frac{1}{2}$
$\{o\}$	0	$\alpha^D > 0$	$\{c\}$	0	0
S^D	1	1	S^P	1	1
\emptyset	0	0	\emptyset	0	0
support:	$\{no\}$	$\{o\}$	support:	$\{d\}$	$\{d\}$

The first DSE (ν_1^D, ν_1^P) describes behaviour which is similar to the backward induction equilibrium in the analysis without ambiguity. The equilibrium beliefs need however not be additive. The plaintiff decides not to go to trial. Whether the defendant D will refuse a settlement or not depends on the strength of the belief α^P that the plaintiff will drop the case. For $\alpha^P \geq 1/2$ the defendant will make no offer otherwise a settlement offer is made. Filing the suit becomes weakly optimal only if the plaintiff is completely certain that the defendant will not refuse a settlement, $\nu_i^D(\{no\}) = 0$.

The second DSE (ν_2^D, ν_2^P) shows behaviour which cannot occur in a Nash equilibrium. In this case, the plaintiff plans to drop the case if challenged. The defendant, on the other hand, feels sufficiently ambiguous, $\nu_2^P(\{d\}) = \alpha^P < 1/2$, about the prediction that the plaintiff will not file a suit, $\text{supp } \nu_2^P = \{d\}$, and will offer settlement. Such behaviour cannot be supported by a Nash equilibrium. Yet it does not appear unreasonable. After careful consideration of the situation, the defendant may well recognise the incentives of the plaintiff to drop the action. This conclusion depends however on the correct perception of the situation as modelled by the game. If the defendant feels sufficiently uncertain about this information she may be justified in offering settlement. Moreover, since the defendant does not challenge the potential plaintiff, no information about the actual behaviour of the plaintiff is generated. So one is also justified to call such a situation an equilibrium. ■

From Part (b) of Definition 3.2 it is clear that no strictly dominated strategy will be chosen in a DSE. Therefore, one may be led to conclude that all DS-equilibria are backward induction equilibria. This conclusion is false however, as the DSE (ν_2^D, ν_2^P) in Example 3.1 demonstrates. Ambiguity may prevent players from choosing strategies which expose them to situations where they might be hurt by a strictly dominated choice of the opponents. This may be so, even if they do not expect the opponents to play strictly dominated strategies, if they do not trust this conclusion sufficiently.

Example 3.1 shows also that the DSE concept can describe behaviour, which is inconsistent with the strict consistency requirements of a Nash equilibrium. Indeed, behaviour, as in equilibrium (ν_2^D, ν_2^P) , can occur only if the worst outcome of an interaction can influence the decision of a player even if it is on events which are outside the support. If one would constrain the support notion to make events outside the support Savage-null, i.e., irrelevant for the player, then this equilibrium would disappear.

This is obvious from the following lemma which has been proved in Ryan (1998) (Lemma 1, p.34).

Lemma 3.4 *Let ν be a capacity on a set S . An event $E \subset S$ is Savage-null if and only if $\nu(S \setminus E) = 1$.*

Applying Lemma 3.4 to the equilibrium in Example 3.1 would imply $\alpha^D = \alpha^P = 1$. Hence, DSE (ν_2^D, ν_2^P) would no longer be possible and only the DSE (ν_1^D, ν_1^P) corresponding to a Nash equilibrium would survive. Indeed, Lo (1996) (Corollary of Proposition 4, p. 468) shows that this is true for all two-player games. Adopting such a strong notion of support therefore defeats the objective of modelling ambiguity of players.

DSE offers more possibilities to model economic situations than traditional Bayesian analysis because consistency requirements on beliefs are weaker. In our opinion this additional freedom is useful for modelling economic situations since it allows us to include aspects of the economic environment, which are precluded by Bayesian analysis, but which are supported by experimental evidence or other robust findings. It is beyond the scope of this paper to investigate these applications in depth.

3.1.2 Ambiguity and Pessimism

With convex capacities, as we assume throughout this paper, ambiguity aversion is built into the concept of the Choquet integral. This pessimism concerns however only events on which there is ambiguity. DSE leaves us with more modelling options. DSE allows us to distinguish between the preference of players for unambiguous choices and their pessimism in the face of ambiguity. For example, if one would like to restrict ambiguity to an opponent's strategic behaviour and considers information about types as hard, one can model this by a capacity which agrees with an additive prior distribution. In this case, pessimism is restricted to the behaviour of the opponent but not to the probability over types. The following example which is due to Ryan (2002) illustrates such a modelling option¹⁰.

Example 3.2 (Ryan 2002)

Consider a signalling game with two players, $i = 1, 2$, where player 1 can be one of two types $T = \{t_1, t_2\}$. It is known that each type occurs with probability $\frac{1}{2}$. Action sets for the two players are $A^1 = \{R, L\}$ and $A^2 = \{U, D\}$. Figure 2 represents the game.

¹⁰ This game has been advanced in Ryan (2002) as an argument for a stronger support notion called “robust support”. In Eichberger and Kelsey (2001) we provide a detailed and more formal discussion of this approach.

Insert FIGURE 2 here

It is easy to check that this game has a unique perfect Bayesian equilibrium where

- player 1 of type t_1 chooses L ,
- player 1 of type t_2 chooses R ,
- player 2 chooses U in response to L and R .

Ryan (2002, p.12) argues that this equilibrium describes the only sensible behaviour in this game because each type of player 1 has a strictly dominant strategy and player 2, knowing the strategy of both types of player 1, maximises her payoff by choosing U . Moreover, the move U is also recommended if player 2 is ambiguity averse since it guarantees the certain payoff of 1 no matter what type player 1 turns out to be.

For a degree of ambiguity $\alpha < 1/2$, the following beliefs (ν^1, ν^2) are a DSE which agrees with the additive prior distribution $(p(t_1), p(t_2)) = (0.5, 0.5)$:

$$\nu^1(\{(t_1, L)\}) = \nu^1(\{(t_2, R)\}) = \alpha \cdot \frac{1}{2},$$

$$\nu^1(\{(t_1, L), (t_2, L)\}) = \nu^1(\{(t_1, L), (t_2, L)\}) = \alpha \cdot \frac{1}{2},$$

$$\nu^1(\{(t_1, L), (t_1, R)\}) = \nu^1(\{(t_2, R), (t_2, L)\}) = \frac{1}{2},$$

$$\nu^1(\{(t_1, L), (t_2, L), (t_2, R)\}) = \nu^1(\{(t_1, L), (t_2, L), (t_2, R)\}) = (1 + \alpha) \cdot \frac{1}{2},$$

$$\nu^1(E) = 0 \quad \text{for all other } E \subset T \times S^1,$$

and ν^2 is an additive probability distribution with $\nu^2(\{(D, D)\}) = 1$.

One checks easily that, for $\alpha > 0$, $\text{supp } \nu^1 = \{(t_1, L), (t_2, R)\}$. The DS-updates are additive and always well-defined: $\rho(t_1|L) := \nu^1(t_1|L) = 1/(2 - \alpha)$ and $\rho(t_1|R) := \nu^1(t_1|R) = (1 - \alpha)/(2 - \alpha)$.

Computing the CEU payoffs for these beliefs yields $P^1(L|t_1, \nu^2) = P^1(R|t_2, \nu^2) = 1$,

$P^1(R|t_1, \nu^2) = P^1(L|t_2, \nu^2) = 0$ and

$$P^2(\tilde{a}|s^1, \nu^1) = \begin{cases} 1 & \text{for } \tilde{a} = U \\ 4 \cdot \frac{1-\alpha}{2-\alpha} & \text{for } \tilde{a} = D \end{cases} \quad \text{for } s^1 = L$$

$$P^2(\tilde{a}|s^1, \nu^1) = \begin{cases} 1 & \text{for } \tilde{a} = U \\ 4 \cdot \frac{1-\alpha}{2-\alpha} & \text{for } \tilde{a} = D \end{cases} \quad \text{for } s^1 = R$$

For $\alpha < 2/3$, D is the best response of player 2 no matter which strategy player 1 chooses, and L and R are the best strategies for player 1 of type t_1 and t_2 respectively. Since the support of ν^1 is $\{(t_1, L), (t_2, R)\}$ and of ν^2 is $\{D, D\}$, playing D is a Dempster-Shafer equilibrium.

For $\alpha \geq 2/3$, (U, U) is the best response of player 2 and the associated DSE yields the same behaviour as the perfect Bayesian equilibrium. ■

Notice that the DS-updates $\nu^1(t_1|s^1)$ equal the prior distribution for $\alpha = 0$, the case of complete strategic ambiguity $\lambda(\nu^1) = 1$. For $\alpha = 1$, that is complete strategic certainty $\lambda(\nu^1) = 0$, the DS-updates correspond to the Bayesian updates, $\nu^1(t_1|L) = 1$ and $\nu^1(t_1|R) = 0$.

If there is no strategic ambiguity, then player 1's moves provide a perfect signal for the type of player 1, i.e. $\text{supp } \nu^1 = \{(t_1, L), (t_2, R)\}$, and player 2 will respond by choosing U in response. With complete ambiguity about player 1's strategy choice, player 2 will assess the likelihood of the two types with the prior probability $1/2$. Based on the expected payoff with respect to the unambiguous prior distribution player 2 will find action D optimal, and not U . Whether the action U or D is chosen depends therefore on the degree of ambiguity which player 2 feels about the deduced equilibrium behaviour. If ambiguity is low, $\lambda(\nu^1) = 1 - \alpha < 1/3$, then player 2 will choose U and if ambiguity is high she will choose D . The

critical level which determines when the behaviour of player 2 will change depends, of course, on the payoff of actions.

If a player feels great ambiguity regarding the strategy of the opponent but not with respect to the prior type distribution, then DS-updating on the observed actions leads the player to disregard the ambiguous strategy and to decide based on the unambiguous prior. Faced with strategic information, a player who is extremely ambiguous about strategic information and unambiguous about type information will revert to the unambiguous information of the prior distribution¹¹.

One can, of course, question the assumption about the unambiguous prior distribution. Indeed, we do not require that beliefs do in general agree with unambiguous priors. It is an easy exercise to check that, with complete ambiguity about the prior distribution, the argument that pessimism commends to play U in order to secure the constant payoff of 1 is correct.

3.2 Existence and properties of DSE

Since Bayesian equilibria are DS-equilibria with a degree of ambiguity $\lambda = 0$, existence of a DSE is guaranteed under the usual conditions. It is not clear however whether there exist DS-equilibria for arbitrary degrees of ambiguity λ and arbitrary prior beliefs about types. Proposition 3.5 shows that DS-equilibria exist under the usual assumptions for any degree of ambiguity.

¹¹ Ryan (2002) considers the case $\alpha = 1/2$. In this case, the DSE predicts player 2 to play D . He sees a tension between the interpretation of Choquet preferences as pessimistic and the preference for the action D which is risky rather than playing U , an action yielding a constant outcome of 1.

Proposition 3.5 *For any degree of ambiguity $\lambda \in (0, 1)$ and any additive prior probability distribution p on T , there exists a Dempster-Shafer equilibrium with this degree of ambiguity λ which agrees with the distribution p on T .*

Proposition 3.5 shows that the DSE concept can be applied in all cases, in which standard Nash equilibria exist. Moreover, it shows that one can choose the degree of ambiguity λ exogenously as a characteristic of a situation and still obtain DS-equilibria. This property is particularly important in economic applications where one wants to study the impact of ambiguity on the behaviour of agents¹²

Games with complete information, i.e., with a type space containing a single type, $|T| = 1$, form an important special case to which one can apply DSE. The DSE (ν_2^D, ν_2^P) in Example 3.1 shows that behaviour in DS-equilibria does not necessarily correspond to behaviour in backward induction equilibria.

Backward induction in the presence of Knightian uncertainty has also been discussed by Dow and Werlang (1994). This paper shows that if there is ambiguity, there are non-backward induction equilibria in the finitely repeated prisoner's dilemma game. These equilibria arise with large degrees of ambiguity, which is compatible with our analysis. Dow and Werlang (1994) analyse games in normal form. Our theory confirms their analysis with an extensive form solution concept based on Knightian uncertainty. We believe that an extensive form solution concept is preferable, since DSE requires that equilibrium strategies are optimal when each move is made. A solution defined on the normal form cannot do this.

¹² Eichberger and Kelsey (2001, 2003) study applications to economic problems.

4. DSEL AS A NASH EQUILIBRIUM REFINEMENT

In Example 3.1 the DSE with little or no ambiguity selects the backward induction equilibrium. With no ambiguity $\lambda = 0$, Dempster-Shafer equilibria, where each strategy of player 1 is played by some type, are perfect Bayesian equilibria (Proposition 3.3). These results suggest that DSE equilibria with ambiguity may provide reasonable restrictions on beliefs, which in the limit as ambiguity vanishes select Bayesian equilibria which are robust with respect to ambiguity. In order to explore this possibility more formally we define a *Dempster-Shafer Equilibrium Limit (DSEL)*.

Ambiguity may affect beliefs over types and beliefs over strategy choice. We do not want to exclude the possibility that ambiguity extends also to ambiguity about types, but we will require only ambiguity about strategies in order to allow for capacities which agree with an additive prior distribution.

Condition A A DSE (ν^1, ν^2) is subject to *strategic ambiguity* if

$$\nu^1(\{s^1\} \times T) + \nu^1((S^1 \setminus \{s^1\}) \times T) < 1$$

holds for all $s^1 \in S^1$.

If a DSE is subject to strategic ambiguity, then the degree of ambiguity is strictly positive, even if the equilibrium beliefs agree with an additive prior distribution over types. Thus, Condition A allows us to consider sequences of Dempster-Shafer equilibria with positive degree of ambiguity which agree with a given additive prior distribution over types.

Definition 4.1 Dempster-Shafer Equilibrium Limit (DSEL)

A set of beliefs $(\bar{\nu}^1, \bar{\nu}^2)$ and updated beliefs $\{\bar{\rho}(\cdot | s^1)\}_{s^1 \in S^1}$ is a *Dempster-Shafer Equilib-*

rium Limit (DSEL) if it is the limit of a sequence of strategically ambiguous Dempster-Shafer equilibria $((\nu_n^1, \nu_n^2), \{\rho_n(\cdot|s^1)\}_{s^1 \in S^1})$ such that the degree of ambiguity λ tends to zero as n tends to infinity.

By Proposition 3.5 there exists a DSE $((\nu_n^1, \nu_n^2), \{\rho_n(\cdot|s^1)\}_{s^1 \in S^1})$ for any degree of ambiguity $\lambda_n > 0$, where, for all $s^1 \in S^1$, $\rho_n(\cdot|s^1)$ is well-defined by the DS-updates $\nu_n^1(\cdot|s^1)$. By convexity of the capacities, $\nu_n^i(E) \in [0, 1]$, for all events $E \subseteq S^i$, $i = 1, 2$, and for all n . Since we consider finite games, the sequence (ν_n^1, ν_n^2) is contained in $[0, 1]^m$, where $m = |T| + |S^1| + |S^2|$. Hence, for any sequence $\lambda_n \rightarrow 0$ there must be a converging subsequence $(\nu_n^1, \nu_n^2) \rightarrow (\bar{\nu}^1, \bar{\nu}^2)$ and $\rho_n(\cdot|s^1) \rightarrow \bar{\rho}(\cdot|s^1)$ for all $s^1 \in S^1$. Thus, DSEL is always well-defined.

Notice that a DSEL requires also to specify a sequence of updated capacities $\{\rho_n(\cdot|s^1)\}_{s^1 \in S^1}$. Since we impose strategic ambiguity there exists always a supporting sequence of Dempster-Shafer equilibria for which the updated beliefs $\{\rho_n(\cdot|s^1)\}_{s^1 \in S^1}$ are well defined by the DS-updates. Even if DS-updates are well-defined along the sequence of DS-equilibria, a DSEL $((\bar{\nu}^1, \bar{\nu}^2), \{\bar{\rho}(\cdot|s^1)\}_{s^1 \in S^1})$ may have non-additive DS-updates $\bar{\rho}(\cdot|s^1)$ for strategies s^1 , which are not played in equilibrium.

The following example shows that beliefs off the equilibrium path need not be additive. In particular, the sequence of DS-equilibria supporting a DSEL

- can be strategically ambiguous,
- agree with an additive prior distribution and
- have well-defined DS-updates.

Notice that a DSEL also requires the specification of a sequence of updated capacities. Yet, as the degree of ambiguity converges to zero, additive beliefs $(\bar{\nu}^1, \bar{\nu}^2)$ obtain in the limit, but DS-updates $\bar{\rho}(\cdot|s^1)$ may remain non-additive if strategy s^1 is not played in the DS-equilibria of the supporting sequence.

Example 4.1 *Consider the signalling game, in Figure 3.*

Insert FIGURE 3 here

The strategy set of player 1 is $A^1 = \{L, R\}$ and of player 2, $A^2 = \{u, m, d\}$. There are two types of player 1, $T = \{t_1, t_2\}$, which occur with probability p_1 and p_2 , $p_1 > p_2$ for concreteness.

In any perfect Bayesian equilibrium both types of player 1 choose R, since any belief $\mu(\cdot|R)$ makes d strictly dominated for player 2. But if player 2 plays d with probability zero, strategy R strictly dominates strategy L for player 1. Hence, there is a unique perfect Bayesian equilibrium where player 2 chooses u.

There are two types of DSEL agreeing with the additive prior distribution (p_1, p_2) . The completely additive DSEL $\left((\tilde{\nu}^1, \tilde{\nu}^2), (\tilde{\rho}(\cdot|R), \tilde{\rho}(\cdot|L)) \right)$, $\tilde{\nu}^1(t_1, R) = p_1$, $\tilde{\nu}^1(t_2, R) = p_2$, $\tilde{\rho}(t_1|R) = p_1$, $\tilde{\rho}(t_2|R) = p_2$, $\tilde{\nu}^2(u) = 1$, is behaviourally equivalent to the perfect Bayesian equilibrium. This DSEL is supported by a sequences of strategically ambiguous E-capacities which agree with the prior distribution. For details about the construction of such a sequence see Eichberger and Kelsey (1999).

There is however another DSEL $(\bar{\nu}^1, \bar{\nu}^2)$ where the updated beliefs of player 1 are not

additive. For any $\alpha < 1/3$, consider the following sequence of DS-equilibria:

E	$\nu_n^1(E)$	E	$\nu_n^2(E)$
$\{(t_1, L)\}$	$(1 - \lambda_n) \cdot p_1$	$\{u\}$	0
$\{(t_1, R)\}$	0	$\{m\}$	0
$\{(t_2, L)\}$	$(1 - \lambda_n) \cdot p_2$	$\{d\}$	$(1 - \lambda_n)$
$\{(t_2, R)\}$	0	$\{u, m\}$	0
$\{(t_1, L), (t_1, R)\}$	p_1	$\{u, d\}$	$(1 - \lambda_n)$
$\{(t_1, L), (t_2, L)\}$	$1 - \lambda_n$	$\{m, d\}$	$(1 - \lambda_n)$
$\{(t_1, L), (t_2, R)\}$	$(1 - \lambda_n) \cdot p_1$	S	1
$\{(t_1, R), (t_2, L)\}$	$(1 - \lambda_n) \cdot p_2$	\emptyset	0
$\{(t_1, R), (t_2, R)\}$	0	$\text{supp } \nu^2$	$\{d\}$
$\{(t_2, R), (t_2, L)\}$	p_2		
$\{(t_1, L), (t_2, L), (t_2, R)\}$	$1 - \lambda_n + \lambda_n \cdot \alpha$		
$\{(t_1, R), (t_2, L), (t_2, R)\}$	p_2		
$\{(t_2, L), (t_1, L), (t_1, R)\}$	$1 - \lambda_n + \lambda_n \cdot \alpha$		
$\{(t_2, R), (t_1, L), (t_1, R)\}$	p_1		
S	1		
\emptyset	0		
$\text{supp } \nu^1$	$\{(t_1, L), (t_2, L)\}$		

The DS-updated capacities $\rho_n(t_1|R) = \nu_n^1(t_1|R) = \alpha$ and $\rho_n(t_2|R) = \nu_n^1(t_2|s_1) = \alpha$ are well-defined, but strictly non-additive.

Computing the CEU payoffs for the DSE (ν_n^1, ν_n^2) , one easily checks that $P^1(L|t_i, \nu_n^2) = 1 > 0 = P^1(R|t_i, \nu_n^2)$ and, for $\alpha < 1/3$, $P^2(d|R, \rho_n) = 4 > 3 + 3 \cdot \alpha = P^2(u|R, \rho_n) = P^2(m|R, \rho_n)$ for all n . Hence, the DSEL $((\bar{\nu}^1, \bar{\nu}^2), (\bar{\rho}(\cdot|R), \bar{\rho}(\cdot|L)))$,

$$\bar{\nu}^1(t_1, L) = p_1, \bar{\nu}^1(t_2, L) = p_2, \quad \text{additive}$$

$$\bar{\rho}(t_1|R) = \alpha, \bar{\rho}(t_2|R) = \alpha, \quad \text{non-additive}$$

$$\bar{\nu}^2(d) = 1, \quad \text{additive}$$

follows from this sequence of DS-equilibria as $\lambda_n \rightarrow 0$. ■

The DSEL $(\bar{\nu}^1, \bar{\nu}^2)$ is interesting since beliefs off the equilibrium path are non-additive, even though beliefs on the equilibrium path are additive. Since perfect Bayesian equilibrium requires that beliefs be additive at all information sets, the expected payoff from u

dominates the payoff from d . DSEL, however, allows strict non-additivity off the equilibrium path, so that the certain payoff of 4 obtained from strategy d becomes more attractive. It is plausible that a player who has observed an out-of-equilibrium move will have some doubts about his original theory of how the game is played. This could cause him to become ambiguity-averse as represented by the non-additivity of the updated beliefs. DSEL allows us to model ambiguity of a player as a consequence of having to update beliefs on events with a capacity weight of zero.

Example 4.1 shows also that there are few constraints on the DS-updates. Indeed, DSE, and therefore DSEL, allow us to impose constraints on players' beliefs directly and to deduce equilibrium beliefs satisfying these constraints. This opens the opportunity to design experiments where ambiguity is manipulated independently from the equilibrium play which one wants to test.

4.1 Properties of DSEL

In this section, we will compare the concept of a DSEL with Bayesian and perfect Bayesian equilibrium. Since Bayesian and perfect Bayesian equilibria have an additive prior distribution over types as a defining criterion we will restrict attention to Dempster-Shafer equilibria which agree with an additive prior distribution throughout this section.

The capacities $(\bar{\nu}^1, \bar{\nu}^2)$ of a DSEL are additive. So it is not difficult to prove that a DSEL is a Bayesian equilibrium.

Proposition 4.2 *A DSEL which agrees with an additive prior distribution over types is a Bayesian equilibrium.*

All DSEL are Bayesian equilibria. The potential of strategic ambiguity to select among the set of Bayesian equilibrium lies in the updated beliefs. Beliefs are generated by the DS-updating rule in combination with constraining assumptions about equilibrium beliefs. A DSE does not tie down the equilibrium beliefs as much as Nash equilibrium does. Hence, there is room for game-specific constraints on beliefs and attitudes towards ambiguity. Depending on the application one can focus on the consequences of the degree of ambiguity aversion, of ambiguity about types or of other characteristics of beliefs. The DS-updates inherit their properties from these fundamental assumptions. To the extent that one can control for the degree of ambiguity, ambiguity aversion and other characteristics of an environment one may be able to test equilibrium properties in experiments.

One of the weakest refinements of Bayesian equilibria is a perfect Bayesian equilibrium. By making updated beliefs part of the equilibrium concept it guarantees optimising behaviour at all information sets whether or not they will be reached in equilibrium. Since perfect Bayesian equilibrium puts no constraint on out-of-equilibrium beliefs, it eliminates only equilibria relying on strictly dominated strategies at information sets off the equilibrium path.

A DSEL allows for beliefs off the equilibrium path which are strictly non-additive. Hence, Example 4.1 shows that a DSEL need not be a perfect Bayesian equilibrium. One may however conjecture that a DSEL with additive updates $\{\bar{p}(\cdot|s^1)\}_{s^1 \in S^1}$ at all information sets is a perfect Bayesian equilibrium. We will show below in Proposition 4.3 that this is the case, indeed.

One may also conjecture that the restrictions on beliefs induced by the sequence of state-

gically ambiguous Dempster-Shafer equilibria would rule out DSEL with additive updates $\{\bar{\rho}(\cdot|s^1)\}_{s^1 \in S^1}$ at all information sets where a player uses a weakly dominated action. This is however not true. For every perfect Bayesian equilibrium it is possible to construct a sequence of strategically ambiguous DS-equilibria, which agree on the additive prior over types and converges to this perfect Bayesian equilibrium. This is almost obvious if all information sets will be reached in the perfect Bayesian equilibrium. If there are information sets following actions which are not played in a perfect Bayesian equilibrium, then one can find a sequence of DS-equilibria in which the DS-updates are not defined at these information sets. Hence, one can assign the off-the-equilibrium-path beliefs of the perfect Bayesian equilibrium to those DS-equilibria. Thus, one can obtain even a perfect Bayesian equilibrium where player 2 chooses weakly dominated strategies off the equilibrium path as a DSEL.

Proposition 4.3 Perfect Bayesian equilibrium and DSEL

- (i) *Every perfect Bayesian equilibrium is a DSEL.*
- (ii) *A DSEL which agrees with an additive prior distribution is a perfect Bayesian equilibrium if all updates are additive.*

DS-updates of a DSEL can, but need not, be additive. Proposition 4.3 shows that additive limits of the DS-updates is the crucial condition for the two concepts to coincide. If DS-updates do not converge to additive probability distributions off the equilibrium path, then strategic ambiguity, modelled by the DS equilibrium concept, provides a refinement of Nash equilibrium based on other principles than standard refinements in the literature.

Mailath (1992) provides an excellent survey of the refinements most commonly used in

signalling games. They all operate by restricting out-of-equilibrium beliefs. Justification for such restrictions is obtained by forward or backward induction arguments. There is an obvious tension in such arguments because out-of-equilibrium behaviour is constrained by reasoning about behaviour which will never be observed.

The DSEL provides an alternative approach to equilibrium selection. Modelling ambiguity about the equilibrium strategy choices directly avoids the tension in the interpretation of out-of-equilibrium beliefs. Moreover, there are behavioural theories behind the DS-updating rule (Gilboa and Schmeidler, 1993) and the Choquet expected utility model (Gilboa, 1987; Sarin and Wakker, 1993). Assumptions about the behavioural foundations of this decision and updating model can and have been tested independently from the equilibrium notion (Camerer and Weber, 1992).

4.2 Out-of-equilibrium beliefs

Refinement of the set of equilibria can be obtained by imposing additional restrictions on the players' non-additive beliefs. As in the standard refinement literature, one can strengthen or weaken the robustness requirement imposed on Bayesian equilibrium by putting further constraints on the sequence of ambiguous DS-equilibria which support it. In contrast to this literature such assumptions are in principle testable.

DS-equilibria which are not perfect Bayesian equilibria are plausible, since they correspond to cases in which player 2 is ambiguity-averse after observing an unexpected move. Example 4.1 illustrates the potential of DSEL to select among Bayesian equilibria based on ambiguity about the behaviour in case of an unexpected out-of-equilibrium move. Yet, even if we do not want to rely on non-additive beliefs off the equilibrium path, DSEL offers

quite intuitive out-of-equilibrium beliefs.

It is impossible to develop here a complete theory of reasonable refinements based on ambiguity, but the following example may provide some intuition. It is a simplified version of the education-signalling model introduced by Spence (1973). With plausible restrictions on the out-of-equilibrium beliefs ambiguity may select the pooling equilibrium as the a unique DSEL¹³. The intuition about beliefs is as follows. A DS-equilibrium which agrees with an additive prior distribution over types models a situation where a player feels ambiguity about the opponents' behaviour but not about the prior distribution over types. This is a natural assumption if past experience has provided information about the frequency of types but if there is no well-established way of signalling private information. In such a situation signalling is endogenous equilibrium behaviour. An out-of-equilibrium move indicates a break-down of the implicit understanding of equilibrium behaviour. In such a case, it appears quite reasonable to return to the "firm" information about the prior distribution over types.

Example 4.2 education signalling¹⁴

Consider two workers with different productivity, $\phi_H > \phi_L$. A worker's productivity is private information but it is common knowledge that the proportion of high-productivity workers is p . Workers will apply for a position in the firm with a wage proposal. A worker

¹³ Notice that Example 4.2 is no contradiction to the result in Proposition 4.3. There are other DSELS corresponding to the typical perfect Bayesian equilibria of the Spence signalling model.

¹⁴ This is a highly stylised version of the education signalling model by Spence (1973). For simplicity of the exposition, we have assumed that the education level is not a choice variable. A more general treatment of the education signalling model can be found in Eichberger and Kelsey (1999a).

can ask for a high wage $w_H = \phi_H$, a low wage $w_L = \phi_L$, or the average wage $\bar{w} = p \cdot \phi_H + (1 - p) \cdot \phi_L$. The firm can only choose to accept the application, a , or to reject it, r . In order to qualify for a high wage w_H , a worker must present an education certificate. The strategy w_H implies that the worker has obtained this education certificate. High-productivity workers can obtain the certificate at no cost, while low-productivity workers incur a cost of -2 . We will assume throughout that the education costs of the low-productivity worker are not justified by the productivity and wage difference, $0 < w_H - w_L = \phi_H - \phi_L < 2$.

Figure 4 illustrates the situation.

Insert FIGURE 4 here

In the notation of Section 3, the game is described by $S^1 = \{w_H, w_L, \bar{w}\}$ and $S^2 = \{s^2(s^1) \mid s^1 \in S^1\}$ and $T = \{H, L\}$. It is easy to see that there are exactly two perfect Bayesian equilibria:

	Worker	Firm	out-of-equilibrium beliefs
(i)	$((w_H, H), (w_L, L))$	$(s^2(w_H), s^2(w_L), s^2(\bar{w})) = (a, a, r)$	$\mu(H \bar{w}) = 0$.
(ii)	$((\bar{w}, H), (\bar{w}, L))$	$(s^2(w_H), s^2(w_L), s^2(\bar{w})) = (r, a, a)$	$\mu(H w_H) = \mu(H w_L) = 0$.

For notational convenience, we have only noted the equilibrium strategies. In terms of beliefs, a perfect Bayesian equilibrium $(\pi^1, \pi^2, \{\mu(\cdot|s^1)\})$ is described by $\pi^1(w^H, H) = 1$, $\pi^1(w^L, L) = 1$ and $\pi^2(a, a, r) = 1$ in case (i) and $\pi^1(\bar{w}, H) = p$, $\pi^1(\bar{w}, L) = 1 - p$ and $\pi^2(r, a, a) = 1$ in case (ii). Beliefs about all other strategies are zero.

Equilibrium (i) is the Pareto-optimal separating equilibrium selected by the intuitive criterion. Equilibrium (ii) is the Pareto-optimal pooling equilibrium, which does not satisfy the “intuitive” belief condition that $w_H > \bar{w}$ could only come from the high-productivity type since only this player would gain from such a deviation relative to the equilibrium payout. If we assume that the prior distribution of types is hard knowledge, while equilibrium infer-

ences about behaviour are ambiguous, then only behaviour of the pooling equilibrium (ii) can arise in a DSEL. We will formalise this assumption about the beliefs by an E -capacity¹⁵. E -capacities are a modification of simple capacities (or distorted probabilities) which have a constant degree of ambiguity and which allow for marginal distributions which are additive.

Fix a common degree of ambiguity λ_n for both players. Denote by $\mu_o^t(E)$ the capacity which equals 1 for $E = S^1 \times T$ and 0 otherwise and by $\mu_o(F)$ the capacity equalling 1 for $F = S^2$ and 0 otherwise. A compact way to write the E -capacities based on additive probability distributions π^1 and π^2 is

$$(6) \quad \begin{aligned} \bullet \quad \nu_n^1(E) &:= \lambda_n \cdot [p \cdot \mu_H(E) + (1-p) \cdot \mu_L(E)] + (1-\lambda_n) \cdot \pi^1(E) \quad \text{for } E \subseteq S^1 \times T, \\ \bullet \quad \nu_n^2(F) &:= \lambda_n \cdot \mu(F) + (1-\lambda_n) \cdot \pi^2(F) \quad \text{for } F \subseteq S^2. \end{aligned}$$

E -capacities are a convex combination between an additive probability distribution π^i and a weighted average of the capacities μ_o^t with weights equal to the probabilities of the prior distribution. Notice that $\nu_n^1(E) + \nu_n^1((S^1 \times T) \setminus E) = 1 - \lambda_n$ and $\nu_n^2(F) + \nu_n^2(S^2 \setminus F) = 1 - \lambda_n$ holds for all events $E \neq S^1 \times T$ and $F \neq S^2$. Thus, there is strict ambiguity if $\lambda_n > 0$ holds. E -capacities have also the property that $\text{supp } \nu^i = \text{supp } \pi^i$. The strategies in the support of the capacity, i.e., the strategies of the opponents which must be optimal, are the strategies in the additive part of the E -capacity. Using Equation 3, one can compute¹⁶ the DS-update of ν_n^1 as

$$\nu_n^1(H|s^1) = \frac{\lambda_n \cdot p + (1-\lambda_n) \cdot \pi^1(s^1, H)}{\lambda_n + (1-\lambda_n) \cdot [\pi^1(s^1, H) + \pi^1(s^1, L)]}.$$

For $\lambda_n > 0$, $\nu_n^1(H|s^1)$ is well defined even if $\pi^1(s^1, H) + \pi^1(s^1, L) = 0$ holds, i.e., if no type plays strategy s^1 in equilibrium. Notice also that for $\pi^1(s^1, H) + \pi^1(s^1, L) = 0$,

¹⁵ Eichberger and Kelsey (1999) provide a thorough study of the properties of E -capacities and an axiomatisation.

¹⁶ An explicit computation is in Eichberger and Kelsey (1999).

$\nu_n^1(t|s^1) = p(t)$ coincides with the prior distribution. This means that a player who observes an out-of-equilibrium move s^1 will update her beliefs to the prior distribution. This property of an E -capacity which agrees with a prior distribution appears sensible if one views the knowledge about the type distribution as firm compared to the beliefs about strategy choices which represent just a consistency requirement for beliefs and optimal actions. It is easy to check that $\nu_n^1 \rightarrow \pi^1$ and $\nu_n^2 \rightarrow \pi^2$. Notice, however, that $\nu_n^1(H|s^1) \rightarrow p$ for all s^1 . Hence, in the limit, we have out-of-equilibrium beliefs $\mu(H|w_H) = \mu(H|w_L) = p$. It remains to show that the beliefs in Equation 6 based on the additive probability distributions $\pi^1(\bar{w}, H) = p$, $\pi^1(\bar{w}, L) = 1 - p$ and $\pi^2(r, a, a) = 1$ form a DSE. This is easily established since $\text{supp } \nu_n^1 = \{(\bar{w}, H), (\bar{w}, L)\}$, $\text{supp } \nu_n^2 = \{(r, a, a)\}$ and the expected payoff functions are:

$$P^1(s^1|H, \nu_n^2) = \begin{cases} w_H \cdot \nu_n^1(\bar{s}^2(w_H) = a) & = 0 & \text{for } s^1 = w_H \\ w_L \cdot \nu_n^1(\bar{s}^2(w_L) = a) & = (1 - \lambda_n) \cdot w_L & \text{for } s^1 = w_L \\ \bar{w} \cdot \nu_n^1(\bar{s}^2(\bar{w}) = a) & = (1 - \lambda_n) \cdot \bar{w} & \text{for } s^1 = \bar{w} \end{cases},$$

$$P^1(s^1|L, \nu_n^2) = \begin{cases} (w_H - 2) \cdot \nu_n^1(\bar{s}^2(w_H) = a) & = 0 & \text{for } s^1 = w_H \\ w_L \cdot \nu_n^1(\bar{s}^2(w_L) = a) & = (1 - \lambda_n) \cdot w_L & \text{for } s^1 = w_L \\ \bar{w} \cdot \nu_n^1(\bar{s}^2(\bar{w}) = a) & = (1 - \lambda_n) \cdot \bar{w} & \text{for } s^1 = \bar{w} \end{cases}$$

and

$$P^2(\tilde{a}|s^1, \nu_n^1) = \begin{cases} [\phi_L - w_H] \cdot [1 - p] & \text{for } \tilde{a} = a \\ 0 & \text{for } \tilde{a} = r \end{cases} \quad \text{if } s^1 = w_H$$

$$\begin{cases} [\phi_H - w_L] \cdot p & \text{for } \tilde{a} = a \\ 0 & \text{for } \tilde{a} = r \end{cases} \quad \text{if } s^1 = w_L$$

$$\begin{cases} p \cdot \phi_H + (1 - p) \cdot \phi_L - \bar{w} = 0 & \text{for } \tilde{a} = a \\ 0 & \text{for } \tilde{a} = r \end{cases} \quad \text{if } s^1 = \bar{w}$$

Hence, choosing \bar{w} is optimal for the worker of either type and accepting a wage offer \bar{w} is optimal for the firm. This establishes that the beliefs in Equation 6 form a DSE for any n . The resulting DSEL is therefore

	Worker	Firm	out-of-equilibrium beliefs
(iii)	$(\bar{w}, H), (\bar{w}, L)$	$(s^2(w_H), s^2(w_L), s^2(\bar{w})) = (r, a, a)$	$\mu(H w_H) = \mu(H w_L) = p$.

■

The equilibrium selection in the DSEL of Example 4.2 depends on the joint assumptions of an unambiguous prior distribution over types and a degree of ambiguity aversion λ_n which, for each step n , is the same for all events $E \subset S^1 \times T$. Constant ambiguity aversion controls for distorted beliefs¹⁷. The result that out-of-equilibrium beliefs coincide with the additive prior distribution is driven by the assumption that the prior distribution is unambiguously known. If this is the case, it makes sense for a player to revert to the unambiguous information as implied by DS-updating, whenever an out-of-equilibrium move occurs which invalidates the equilibrium behaviour prediction. In the game of Example 4.2, this assumption about the prior distribution rules out the separating equilibrium. The separating equilibrium would require complete trust in the equilibrium behaviour because a low-productivity worker has an incentive to break away from the separating equilibrium and to propose the average wage rather than the low wage. To argue that the firm should assume that an average wage offer could only come from the low-productivity type would mean that the firm feels no ambiguity about the behaviour of the workers.

In contrast to Bayesian equilibrium, DSE has a updating rule which works also for out-of-equilibrium beliefs if there is some ambiguity about strategy choices. Reasonable assumptions about beliefs can be imposed directly. For example, partial information can be assumed as in the case of a well-known prior distribution in Example 4.2. Whether this is an appropriate assumption or not can be assessed independent from the equilibrium, which is an advantage in economic applications.

¹⁷ To establish the result of Proposition 4.3 that every perfect Bayesian equilibrium is a DSEL we had to relax this assumption of a constant degree of ambiguity in each step of the belief sequence supporting the DSEL.

5. CONCLUSION

We have applied the theory of Knightian uncertainty to sequential games. The evidence suggests that there are occasions in which individuals have large degrees of ambiguity. Despite this we believe that an interesting case is when the degree of ambiguity is small. Under this assumption, we have shown that our definition of equilibrium is a refinement of Bayesian equilibrium, which is similar in spirit to perfect Bayesian equilibrium but does not exactly coincide with it. Since DSEL is a special case of Bayesian equilibrium, no irrational behaviour is introduced by considering non-additive beliefs. As we have shown, even in the limit as beliefs converge to additive probabilities, significant deviations from behaviour under subjective expected utility are possible off the equilibrium path. We believe this is one of our main innovations.

Appendix

PROOFS

Lemma 2.2: If a convex capacity ν has zero degree of ambiguity then it is additive.

Proof. Suppose that ν is not additive, then there exist $A, B \subset S$, such that $A \cap B = \emptyset$ and $\nu(A \cup B) > \nu(A) + \nu(B)$. Let $C = S \setminus (A \cup B)$. Then since the degree of ambiguity is zero:

$$(A-1) \quad 1 = \nu(A \cup B) + \nu(C) > \nu(A) + \nu(B) + \nu(C).$$

By convexity, $1 = \nu((A \cup B) \cup (A \cup C)) \geq \nu(A \cup B) + \nu(A \cup C) - \nu(A)$. Since the degree of ambiguity is zero, $\nu(A \cup B) = 1 - \nu(C)$ and $\nu(A \cup C) = 1 - \nu(B)$. Substituting, we obtain $1 \geq 1 - \nu(C) + 1 - \nu(B) - \nu(A)$, but this contradicts A-1. ■

Proposition 3.3:

- a) A Dempster-Shafer Equilibrium (ν^1, ν^2) with a degree of ambiguity $\lambda = 0$, for which the belief of player 2, ν^1 , agrees with the additive prior distribution p on T is a *Bayesian Equilibrium*.
- b) Consider a Dempster-Shafer Equilibrium (ν^1, ν^2) with a degree of ambiguity $\lambda = 0$, for which the belief of player 2, ν^1 , agrees with the additive prior distribution p on T . If for each strategy $s^1 \in S^1$ there exists a type $t \in T$ such that $(s^1, t) \in \text{supp } \nu$, then the Dempster-Shafer Equilibrium (ν^1, ν^2) is a *perfect Bayesian equilibrium*.

Proof. Part (a): Since the DSE (π^1, π^2) has a degree of ambiguity $\lambda = 0$, by Lemma 2.2, π^1 and π^2 must be additive probability distributions. Since the DSE agrees with an additive

prior distribution p on T , $\sum_{s^1 \in S^1} \pi^1(s^1, t) = p(t)$ for all $t \in T$. Hence, condition DSE-a of Definition 3.2 can be written as

$$s^1 \in \arg \max_{\tilde{s}^1 \in S^1} \sum_{s^2 \in S^2} \pi^2(s^2) \cdot u^1(\tilde{s}^1, s^2(s^1), t)$$

for all $s^1 \in S^1$ with $\pi^1(s^1, t) > 0$ and all $t \in T$. Condition DSE-b requires

$$s^2(s^1) \in \arg \max_{\tilde{a} \in A^2} \sum_{t \in T} \rho(t|s^1) \cdot u^2(s^1, \tilde{a}, t)$$

for all $s^2 \in S^2$ with $\pi^2(s^2) > 0$. By Condition DSE-c and Equation 3,

$$\rho(t|s^1) := \pi^1(S^1 \times \{t\} | \{s^1\} \times T) = \frac{\pi^1(s^1, t)}{\sum_{t \in T} \pi^1(s^1, t)}$$

provided $\sum_{t \in T} \pi^1(s^1, t) \neq 0$. Note that beliefs off the equilibrium path $\rho(t|s^1)$ are arbitrary and need not even be additive. For actions $s^1 \in S^1$ such that $\sum_{t \in T} \pi^1(s^1, t) = 0$ all actions $a \in A^2$ are optimal. Hence, Part (a) of Proposition 3.3 defines a Bayesian equilibrium with mixed strategies (π^1, π^2) .

Part (b): If, in addition, for each strategy $s^1 \in S^1$ there exists a type $t \in T$ such that $(s^1, t) \in \text{supp } \pi^1$, then $\sum_{t \in T} \pi^1(s^1, t) \neq 0$ for all $s^1 \in S^1$ and $\rho(t|s^1) = (\pi^1(s^1, t)) / (\sum_{t \in T} \pi^1(s^1, t))$ is defined at all information sets. In this case, (π^1, π^2) is a perfect Bayesian equilibrium.

■

Lemma 3.4: Let ν be a capacity on a set S . An event $E \subset S$ is Savage-null if and only if $\nu(S \setminus E) = 1$.

Proof. An event E is Savage-null if for any three outcomes $x, y, z \in X$ the CEU value of the acts $x_E y$ and $z_E y$ are equal, i.e.

$$u(x) \cdot \nu(E) + u(y) \cdot [1 - \nu(E)] = u(y) \cdot \nu(S \setminus E) + u(z) \cdot [1 - \nu(S \setminus E)]$$

where we assume, without loss of generality, $u(x) > u(y) > u(z)$. This equality can hold for arbitrary outcomes $x, y, z \in X$ with this order if and only if $\nu(S \setminus E) = 1$ and $\nu(E) = 0$.

■

Proposition 3.5: For any degree of ambiguity $\lambda \in (0, 1)$ and any additive prior probability distribution p on T , there exists a Dempster-Shafer equilibrium with this degree of ambiguity λ , which agrees with the distribution p on T .

Proof. The proof uses the special form of an E-capacity, which is extensively discussed in Eichberger and Kelsey (1999). E-capacities are modifications of an additive probability distribution with a constant degree of ambiguity and, possibly, some additive marginal distributions. If there are no additive marginals then E-capacities are *simple capacities*. Moreover, the support of an E-capacity coincides with the support of its additive part. Hence, for given prior distributions of types and given degrees of ambiguity, E-capacities are completely described by their additive part. Given beliefs modelled by E-capacities, one can use standard arguments to show that there is a Nash-equilibrium in mixed strategies for the modified game where the Choquet payoff functions are viewed as functions of the additive part of the E-capacities.

Fix $\lambda^1, \lambda^2 \in (0, 1)$ and any additive probability distribution p on T . For any finite set X denote by $\Delta(X)$ the set of additive probability distributions on X . Let $\Delta_p(S^1 \times T)$ be the set of additive probability distributions on $S^1 \times T$ with marginal distribution p on T . This set is non-empty, compact and convex.

For any $E \subset S^1 \times T$ let the capacity ν_t be defined as

$$\nu_t(E) := \begin{cases} 1 & \text{if } S^1 \times \{t\} \subset E \\ 0 & \text{otherwise} \end{cases},$$

and for any $\pi \in \Delta_p(S^1 \times T)$ consider the capacity

$$(A-2) \quad \nu_\pi^1(E) := \lambda^1 \cdot \sum_{t \in T} p(t) \cdot \nu_t(E) + (1 - \lambda^1) \cdot \pi(E).$$

For any set $T' \subset T$, the DS-update of the capacity defined in Equation A-2 is (see Eichberger and Kelsey, 1999, Lemma 4.2, p. 132)

$$(A-3) \quad \nu_\pi^1(T'|s^1) = \frac{\lambda^1 \cdot \sum_{t \in T'} p(t) + (1 - \lambda^1) \cdot \sum_{t \in T'} \pi(s^1, t)}{\sum_{t \in T} \pi(s^1, t) + \lambda^1 \cdot \left[1 - \sum_{t \in T} \pi(s^1, t) \right]}.$$

For $\lambda^1 > 0$, $\nu_\pi^1(T'|s^1)$ is a continuous function of π .

For any $s^1 \in S^1$ and any $\pi \in \Delta_p(S^1 \times T)$, let

$$(A-4) \quad \rho(\pi, s^1) = \arg \max_{\sigma \in \Delta(A^2)} \sum_{a \in A^2} \sigma(a) \cdot P^2(a|s^1, \nu_\pi^1)$$

be the set of best behaviour strategies on A^2 for given history s^1 and given belief ν_π^1 based on π .

From the definition of the Choquet integral in Equation 1, it is clear that $P^2(a|s^1, \nu_\pi^1)$ is a continuous function of the capacity ν_π^1 . From the DS-update in Equation A-3 we know that ν_π^1 is a continuous function of π . Thus, $\sum_{a \in A^2} \sigma(a) \cdot P^2(a|s^1, \nu_\pi^1)$ is a continuous function of π . Hence, by Berge's maximum theorem (e.g., Takayama, 1985, p. 254), $\rho(\pi, s^1)$ is a upper-hemi-continuous correspondence. Since $\Delta(A^2)$ is a convex set and since $\sum_{a \in A^2} \sigma(a) \cdot P^2(a|s^1, \nu_\pi^1)$ is linear in σ , the correspondence $\rho(\pi, s^1)$ is also convex-valued.

For any $t \in T$ and a vector $\mu = (\mu(\cdot|s^1))_{s^1 \in S^1}$ of additive probability distributions $\mu(\cdot|s^1) \in \Delta(A^2)$ define the capacities $\nu_\mu^2(E|s^1) := (1 - \lambda^2) \cdot \mu(E|s^1)$. The capacity $\nu_\mu^2(\cdot|s^1)$ is continuous in μ . Let

$$(A-5) \quad \psi(\mu, t) = \arg \max_{\sigma \in \Delta(S^1)} \sum_{s^1 \in S^1} \sigma(s^1) \cdot P^1(s^1|t, \nu_\mu^2(\cdot|s^1))$$

be the best-response correspondence for player 1. Finally, let

$$(A-6) \quad \phi(\mu) := \{\pi \in \Delta_p(S^1 \times T) \mid \pi(s^1, t) = \sum_{t \in T} p(t) \cdot \sigma_t(s^1), \sigma_t \in \psi(\mu, t)\}.$$

Since the Choquet integral $P^1(s^1|t, \nu_\mu^2(\cdot|s^1))$ is continuous in $\nu_\mu^2(\cdot|s^1)$, which in turn is continuous in μ , we can conclude that $\sum_{s^1 \in S^1} \sigma(s^1) \cdot P^1(s^1|t, \nu_\mu^2(\cdot|s^1))$ is continuous in μ . Moreover, the objective function is linear in σ . Applying Berge's maximum theorem again, we conclude that the correspondences $\psi(\mu, t)$ are upper-hemi-continuous and convex-valued. The correspondence $\phi(\mu)$ defined in Equation A-6 is a convex set for each μ and clearly also upper-hemi-continuous.

Consider the mapping $\Theta : \Delta_p(S^1 \times T) \times \Delta(A^2)^{|S^1|} \rightarrow \Delta_p(S^1 \times T) \times \Delta(A^2)^{|S^1|}$ defined by

$$(\pi, \mu) \mapsto \Theta(\pi, \mu) := \times_{s^1 \in S^1} \rho(\pi, s^1) \times \phi(\mu).$$

As the Cartesian product of upper-hemi-continuous and convex-valued correspondences Θ is itself a upper-hemi-continuous and convex-valued correspondence. Moreover, the set $\Delta_p(S^1 \times T) \times \Delta(A^2)^{|S^1|}$ is compact and convex. Hence one can apply Kakutani's fixed-point theorem to establish the existence of $(\pi^*, \mu^*) \in \Theta(\pi^*, \mu^*)$.

Define the additive probability distribution τ^* on S^2 , the set of strategies of player 2, by $\tau^*(s^2) := \prod_{s^1 \in S^1} \mu^*(s^2(s^1)|s^1)$ for all $s^2 \in S^2$. For all $E \subset S^2$, let $\nu_{\tau^*}^2(E) := (1 - \lambda^2) \cdot \tau^*(E)$.

We claim that the capacities $(\nu_{\pi^*}^1, \nu_{\tau^*}^2)$ are a DSE. To see this, note first that $\text{supp } \nu_{\pi^*}^1 = \{s^1 \in S^1 \mid \pi^*(s^1) > 0\}$ and $\text{supp } \nu_{\tau^*}^2 = \{s^2 \in S^2 \mid \tau^*(s^2) > 0\}$ (Eichberger and Kelsey, 1999, Lemma 2.2, p. 121). Hence, $(s^1, t) \in \text{supp } \nu_{\pi^*}^1$ implies $\pi^*(s^1, t) = p(t) \cdot \sigma_t(s^1) > 0$ and $\sigma_t \in \psi(\mu^*, t)$. Hence, s^1 must maximise $P^1(s^1|t, \nu_{\mu^*}^2(\cdot|s^1))$. Similarly, $s^2 \in \text{supp } \nu_{\tau^*}^2$ implies $\tau^*(s^2) > 0$. Hence, $\mu^*(s^2(s^1)|s^1) > 0$ for all $s^1 \in S^1$. Therefore, $s^2(s^1)$ must be a maximiser of $P^2(a|s^1, \nu_{\pi^*}^1)$.

Finally, by construction, $\lambda^1(\nu_{\pi^*}^1) = \lambda^1$ and $\lambda^2(\nu_{\tau^*}^2) = \lambda^2$. Since λ^1 and λ^2 were chosen arbitrarily, the existence result follows for any $\lambda \in (0, 1)$. ■

Proposition 4.2: A DSEL which agrees with an additive prior distribution over types is a *Bayesian equilibrium*.

Proof. Note first that $\lambda = 0$ in the limit implies that the limit capacities $(\bar{\nu}^1, \bar{\nu}^2)$ are additive probability distributions. Hence, $\text{supp } \bar{\nu}^1 = \{(s^1, t) \in S^1 \times T \mid \bar{\nu}^1(s^1, t) > 0\}$ and $\text{supp } \bar{\nu}^2 = \{s^2 \in S^2 \mid \bar{\nu}^2(s^2) > 0\}$. Moreover, $\nu_n^1(s^1, t) \rightarrow \bar{\nu}^1(s^1, t) > 0$ and $\nu_n^2(s^2) \rightarrow \bar{\nu}^2(s^2) > 0$, imply $\nu_n^1(s^1, t) > 0$ and $\nu_n^2(s^2) > 0$ for n large enough. Note also that $P^1(s^1|t, \nu^2)$ is continuous in ν^2 and $P^2(a|s^1, \rho)$ is continuous in ρ . If $\bar{\nu}^1(\{s^1\} \times T) > 0$, i.e., if the DS-update $\bar{\nu}^1(t|s^1)$ is defined, then $\bar{\rho}(\cdot|s^1) = \bar{\nu}^1(\cdot|s^1)$.

Suppose now that $(\bar{\nu}^1, \bar{\nu}^2)$ is not a Bayesian equilibrium. Then there exists $(s^1, t) \in \text{supp } \bar{\nu}^1$ such that $P^1(s^1|t, \bar{\nu}^2) < P^1(\tilde{s}^1|t, \bar{\nu}^2)$ for some $\tilde{s}^1 \in S^1$ and/or $s^2 \in \text{supp } \bar{\nu}^2$ such that $P^2(s^2(s^1)|s^1, \bar{\rho}) < P^2(\tilde{a}|s^1, \bar{\rho})$ for some $\tilde{a} \in A^2$. By continuity of $P^1(s^1|t, \nu^2)$ and $P^2(a|s^1, \rho)$ in ρ and ν^2 , respectively, we can conclude that $P^1(s^1|t, \nu_n^2) < P^1(\tilde{s}^1|t, \nu_n^2)$ for some $\tilde{s}^1 \in S^1$ and/or $P^2(s^2(s^1)|s^1, \rho_n) < P^2(\tilde{a}|s^1, \rho_n)$ for some $\tilde{a} \in A^2$. Since $(s^1, t) \in \text{supp } \bar{\nu}^1$ and $s^2 \in \text{supp } \bar{\nu}^2$ imply that $\nu_n^1(s^1, t) > 0$ and $\nu_n^2(s^2) > 0$ for n large, $(s^1, t) \in \text{supp } \nu_n^1$ and $s^2 \in \text{supp } \nu_n^2$ follows. Hence, (ν_n^1, ν_n^2) is not a DSE. ■

Proposition 4.3: *Perfect Bayesian equilibrium and DSEL*

- (i) Every *perfect Bayesian equilibrium* is a DSEL.
- (ii) A DSEL which agrees with an additive prior distribution is a *perfect Bayesian equilibrium* if all updates are additive.

Proof. *Part (i).* The proof is constructive. For a given perfect Bayesian equilibrium $(\pi^1, \pi^2, \{\mu(\cdot|s^1)\}_{s^1 \in S^1})$

we show that there is a sequence of appropriately modified E-capacities¹⁸ $(\nu_n^1, \nu_n^2, \{\rho_n(\cdot|s^1)\}_{s^1 \in S^1})$ which converges to the perfect Bayesian equilibrium. The trick is to construct this sequence such that the Dempster-Shafer-update of any strategy s^1 which is not played by any type in the perfect Bayesian equilibrium, $\pi^1(\{s^1\} \times T) = 0$, is not defined, i.e., $\nu_n^1(S^1 \setminus \{s^1\} \times T) = 1$. Hence, one can choose the update $\rho_n(\cdot|s^1)$ arbitrarily, in particular equal to the update $\mu(\cdot|s^1)$ of the perfect Bayesian equilibrium. This sequence of capacities $(\nu_n^1, \nu_n^2, \{\rho_n(\cdot|s^1)\}_{s^1 \in S^1})$ is a DSE by standard continuity arguments.

Suppose there is a perfect Bayesian equilibrium $\pi^1 \in \Delta(S^1) \times \Delta(T)$ with $\pi^1(S^1 \times T') = \sum_{t \in T'} p(t)$ for all $T' \subseteq T$, $\pi^2 \in \Delta(S^2)$ with (additive) out-of-equilibrium beliefs $\mu(\cdot|s^1) \in \Delta(T)$ for all $s^1 \in S^1$. Define a sequence of DS-equilibria $((\nu_n^1, \nu_n^2), \{\rho_n(\cdot|s^1)\}_{s^1 \in S^1})$ as follows:

Consider sequences $\lambda_n^1 > 0$ and $\lambda_n^2 > 0$ which converge to zero. Denote by $S_+^1 \subseteq S^1$ the set of strategies with $\sum_{t \in T} \pi^1(s^1, t) \neq 0$. For any non-empty $E \subseteq S^1$ and $F \subseteq T$ let

$$\nu_n^1(E \times F) = \begin{cases} 1 & \text{if } S_+^1 \times T \subseteq E \times F \\ \lambda_n^1 \cdot \sum_{t \in T} p(t) \cdot \nu_t(E \times F) + (1 - \lambda_n^1) \cdot \pi^1(E \times F) & \text{otherwise} \end{cases},$$

where

$$\nu_t(E \times F) := \begin{cases} 1 & \text{if } S^1 \times \{t\} \subset E \times F \\ 0 & \text{otherwise} \end{cases}.$$

It is easy to check that ν_n^1 is a capacity which, by construction, agrees with the prior distribution $p(t)$. Moreover, $\nu_n^1 \rightarrow \pi^1$.

For $s^1 \in S_+^1$; $S_+^1 \times T \not\subseteq S^1 \setminus \{s^1\} \times T$ and $\nu_n^1(S^1 \setminus \{s^1\} \times T) < 1$. Hence, the DS-updates $\nu_n^1(T'|s^1)$ are well-defined and converge to the Bayesian updates. The updates $\rho_n(\cdot|s^1) = \nu_n^1(\cdot|s^1)$ are well-defined and converge to $\mu(\cdot|s^1)$.

On the other hand, for $s^1 \notin S_+^1$ $S_+^1 \times T \subseteq S^1 \setminus \{s^1\} \times T$ and the DS-updates are not defined. Hence, we can choose $\rho_n(\cdot|s^1) = \mu(\cdot|s^1)$ for all n in this case.

¹⁸ For more on E-capacities, their properties and updates, see Eichberger and Kelsey (1999).

Finally, for any subset $E \subset S^2$, $\nu_n^2(E) = (1 - \lambda_n) \cdot \sum_{s^2 \in E} \pi^2(s^2)$.

It is easy to check that the capacities ν_n^1, ν_n^2 are strategically ambiguous and that $\text{supp } \nu_n^1 = \{(s^1, t) | \pi^1(s^1, t) > 0\}$ and $\text{supp } \nu_n^2 = \{s^2 | \pi^2(s^2) > 0\}$. By continuity of P^1 and P^2 in ν_n^2 and ρ_n , respectively, $s^1 \in \arg \max_{\tilde{s}^1 \in S^1} P^1(\tilde{s}^1 | t, \nu_n^2)$ for all $(s^1, t) \in \text{supp } \nu_n^1$ and $s^2(s^1) \in \arg \max_{\tilde{a} \in A^2} P^2(\tilde{a} | s^1, \rho_n(\cdot | s^1))$ for all $s^2 \in \text{supp } \nu_n^2$. Hence, $((\nu_n^1, \nu_n^2), \{\rho_n(\cdot | s^1)\}_{s^1 \in S^1})$ is a DSE.

Part (ii). Consider a sequence of strategically ambiguous DSE $(\nu_n^1, \nu_n^2) \rightarrow (\bar{\nu}^1, \bar{\nu}^2)$ and, for all $s_1 \in S_1$, $\rho_n(\cdot | s^1)$ converges to an additive update $\bar{\rho}(\cdot | s^1)$. Since $\bar{\nu}^1$ is additive, $(s^1, t) \in \text{supp } \bar{\nu}^1$ implies $\bar{\nu}^1(s^1, t) > 0$. Suppose there is a $\tilde{s}^1 \in \text{supp } \bar{\nu}^1$ such that $P^1(\tilde{s}^1 | t, \nu^2) > P^1(s^1 | t, \nu^2)$. By continuity of P^1 in ν^2 , $P^1(\tilde{s}^1 | t, \nu_n^2) > P^1(s^1 | t, \nu_n^2)$ for large n . Moreover, for large n , $\nu_n^1(\tilde{s}^1, t) > 0$ and, therefore, $\tilde{s}^1 \in \text{supp } \nu_n^1$. This contradicts the assumption that (ν_n^1, ν_n^2) is a DSE. Hence, $((\bar{\nu}^1, \bar{\nu}^2), \{\bar{\rho}(\cdot | s^1)\}_{s^1 \in S^1})$ satisfies Condition PBE-a.

Since $\bar{\rho}(\cdot | s^1)$ is additive, an analogous argument shows that Condition PBE-b must hold for $(\bar{\nu}^1, \bar{\nu}^2)$,

$\{\bar{\rho}(\cdot | s^1)\}_{s^1 \in S^1}$. Finally, if $\bar{\nu}^1(\{s^1\} \times T) = 1 - \bar{\nu}^1((S^1 \setminus \{s^1\}) \times T) > 0$, then $\nu_n^1((S^1 \setminus \{s^1\}) \times T) < 1$ for n large enough. Hence, all DS-updates are well-defined for large n and $\rho_n(T' | s^1) = \nu_n^1(T' | s^1) \rightarrow \bar{\rho}(T' | s^1) = \bar{\nu}^1(T' | s^1)$. Thus, Condition PBE-c is satisfied. Given the assumption that the DSEL agrees with an additive prior distribution over types it is therefore a perfect Bayesian equilibrium. ■

References

- Aumann, R., "Backward Induction and Common Knowledge of Rationality," *Games and Economic Behaviour* 8 (1994), 6–19.
- Battigalli, P., "On Rationalizability in Extensive Form Games," *Journal of Economic Theory* 74 (1997), 40–61.
- Banks, J., C. Camerer and D. Porter, "An Experimental Analysis of Nash Refinements in Signaling Games," *Games and Economic Behavior* 6 (1994), 1-31.
- Bebchuk, L. A., "Suing Solely to Extract a Settlement Offer," *Journal of Legal Studies* 17 (1988), 437-449.
- Brandts J. and C. A. Holt, "Adjustment Patterns and Equilibrium Selection in Experimental Signaling Games," *International Journal of Game Theory* 22 (1993), 279-302.
- Brandts J. and C. A. Holt, "An Experimental Test of Equilibrium Dominance in Signaling Games," *American Economic Review* 82 (1992), 1350-1365.
- Camerer, C. and M. Weber, "Recent Developments in Modelling Preferences: Uncertainty and Ambiguity," *Journal of Risk and Uncertainty* 5 (1992), 325-370.
- Choquet, G., "Theory of Capacities," *Annales Institut Fourier* 5 (1953-4), 131-295.
- Dow, J. and S. R. C. Werlang, "Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction," *Journal of Economic Theory* 64 (1994), 305-324.
- Dow, J. and S. R. C. Werlang, "Risk Aversion, Uncertainty Aversion and the Optimal Choice of Portfolio," *Econometrica* 60 (1992), 197-204.
- Eichberger, J. and D. Kelsey, "Ambiguity and Social Interaction," Discussion Paper No. 03-30, Sonderforschungsbereich 504, Universität Mannheim, 2003.

- Eichberger, J. and D. Kelsey, "Strategic Complements, Substitutes and Ambiguity: The Implications for Public Goods," *Journal of Economic Theory* 106 (2002), 436-466.
- Eichberger, J. and D. Kelsey, "A Note on the Support Notion for Non-Additive Beliefs," mimeo, University of Heidelberg, 2001.
- Eichberger, J. and D. Kelsey, "Non-Additive Beliefs and Strategic Equilibria," *Games and Economic Behaviour* 30 (2000), 182-215.
- Eichberger, J. and D. Kelsey, "E-Capacities and the Ellsberg Paradox," *Theory and Decision* 46 (1999), 107-138.
- Eichberger, J. and D. Kelsey, "Education-Signalling and Uncertainty," in M. J. Machina and B. R. Munier, eds., *Beliefs, Interactions and Preferences in Decision Making* (Amsterdam: Kluwer, 1999a), 135-157.
- Eichberger, J. and D. Kelsey, "Uncertainty Aversion and Preference for Randomisation," *Journal of Economic Theory* 71 (1996), 31-43.
- Gilboa, I., "Duality in Non-Additive Expected Utility Theory," *Annals of Operations Research* 19 (1989), 405-414.
- Gilboa, I., "Expected Utility Theory with Purely Subjective Probabilities," *Journal of Mathematical Economics* 16 (1987), 65-88.
- Gilboa, I. and D. Schmeidler, "Updating Ambiguous Beliefs," *Journal of Economic Theory* 59 (1993), 33-49.
- Haller, H., "Non-Additive Beliefs in Solvable Games," *Theory and Decision* 49 (2000), 313-338.
- Hendon, E., H. Jacobsen, B. Sloth and T. Tranaes, "Nash Equilibrium with Lower Proba-

- bilities,” mimeo, University of Copenhagen, 1994.
- Klibanoff, P., “Uncertainty, Decision, and Normal Form Games,” mimeo, Northwestern University, 1996.
- Kohlberg, E. and J.-F. Mertens, “On the Strategic Stability of Equilibria,” *Econometrica* 54 (1986), 1003-1037.
- Ledyard, J. O., “Public Goods: A Survey of Experimental Research,” in J. H. Kagel and A. E. Roth, eds., *The Handbook of Experimental Economics* (Princeton, N. J.: Princeton University Press, 1995), 111-194.
- Lo, K. Ch., “Extensive Form Games with Uncertainty Averse Players,” *Games and Economic Behaviour* 28 (1999), 256-270.
- Lo, K. Ch., “Equilibrium in Beliefs under Uncertainty,” *Journal of Economic Theory* 71 (1996), 443-484.
- Mailath, G.J., “Signalling Games,” in J. Creedy, J. Borland and J. Eichberger, eds., *Recent Developments in Game Theory* (Aldershot: Edgar Elgar, 1992), 65-93.
- Marinacci, M., “Ambiguous Games,” *Games and Economic Behaviour* 31 (2000), 191-219.
- Milgrom, P. and K. Roberts, “Distributional Strategies for Games with Incomplete Information,” *Mathematics of Operations Research* 10 (1986), 619-631.
- Mukerji, S., “Understanding the Non-Additive Probability Decision Model,” *Economic Theory* 9 (1997), 23-46.
- Nehring, K., “On the Interpretation of Sarin and Wakker’s A Simple Axiomatisation of Non-Additive Expected Utility,” *Econometrica* 62 (1994), 935-938.

- Ochs, J., "Coordination Problems," in J. H. Kagel and A. E. Roth, eds., *The Handbook of Experimental Economics* (Princeton, N. J.: Princeton University Press, 1995), 195-252.
- Roth, A., "Bargaining Experiments," in J. H. Kagel and A. E. Roth, eds, *The Handbook of Experimental Economics* (Princeton, N. J.: Princeton University Press, 1995), 253-348.
- Ryan, M.J., "Violations of Belief Persistence in Dempster-Shafer Equilibrium," *Games and Economic Behavior* 39 (2002), 167-174.
- Ryan, M.J., "Rational Belief Revision by Uncertainty-Averse Decision-Makers," mimeo, Australian National University, 1998.
- Sarin, R. and P. Wakker, "A Simple Axiomatization of Non-Additive Expected Utility," *Econometrica* 60 (1992), 1255-1272.
- Savage, L. J., *The Foundations of Statistics* (New York: John Wiley & Sons, 1954).
- Schmeidler, D., "Subjective Probability and Expected Utility without Additivity," *Econometrica* 57 (1989), 571-587.
- Shafer, G., *A Mathematical Theory of Evidence* (Princeton, N. J.: Princeton University Press, 1976).
- Spence, M., "Job Market Signalling," *Quarterly Journal of Economics* 87 (1973), 355-374.
- Takayama, A., *Mathematical Economics*, 2nd edition (Cambridge, UK: Cambridge University Press, 1985).




Figure 1: Frivolous suit




Figure 2: Ryan (2002)




Figure 3: Non-additive beliefs off the equilibrium path




Figure 4: Signalling game