University of Pretoria
Department of Economics Working Paper Series

Optimal Information Transmission

Wei Ma
University of Pretoria
Working Paper: 2015-30
May 2015

Department of Economics
University of Pretoria
0002, Pretoria
South Africa
Tel: +27 12 420 2413

Optimal Information Transmission

Wei Ma[*]

Abstract

This paper addresses the issue of how a given piece of information should be transmitted from a better-informed doctor to an ill-informed patient. The information to be transmitted is expressed as a probability distribution on a space of the patient's possible health states. For a formal analysis of the issue we develop a two-person dynamic game, in which the doctor sends a sequence of messages to the patient to inform him of his health state, and the patient, after receiving each message, chooses an action in an attempt to improve upon his current health status. We study some standard properties of the equilibria of this game; in particular, we show that it has a subgame perfect equilibrium.

Keywords: Information transmission; Dynamic game theory; Subgame perfect equilibrium.
JEL classification: D73; D83

1 Introduction

Crawford and Sobel (1982) studied the question *how much* information a better-informed agent should send to the other agent in a two-person game, and this question was then further analyzed by Caplin and Leahy (2004) in the context of behavioral medicine. As a continuation of these two papers, we shall consider here, also in the context of behavioral medicine, a different but related question: Supposing the amount of information to be sent from one agent to the other has been determined, or more simply, given a piece of information, *how* should it be sent from the better-informed agent (henceforth call *her* the doctor) to the other agent (henceforth call *him* the patient).

Before making this question more precise let us first see an example, partly as an illustration and partly as a motivation. Suppose that a patient is diagnosed as having tumour with probability 0.8, and this tumour has an even chance of being either benign or malignant. Then the patient's health condition can formally be described by a probability distribution which yields benign tumor with probability 0.4, malignant tumor with probability 0.4, and health with probability 0.2. Suppose further that the doctor has two chances to communicate with the patient on his health condition. So one way for her to do so is this: First tell him that he has tumour with probability 0.8, and then, in the second chance, tell him that the tumour is with an even chance of being benign or malignant. A second way is to tell him, in a one-shot fashion, that he has benign tumor with probability 0.4, etc.

[*]Department of Economics, University of Pretoria, Pretoria 0002, South Africa.

The first way is usually called gradual, and the second one-shot, resolution of uncertainty. This situation is connected with the framing effect of decision making under risk (Tversky and Kahneman (1981, Problems 6 and 7)). Given two risky situations S_1 and S_2, Tversky and Kahneman's experiment showed that an individual may prefer S_1 to S_2 when their uncertainties are resolved gradually, but he may reverse his preference when the uncertainties are instead resolved in a one-shot fashion. Additionally, a number of experimental studies in financial economics, too, have shown that people are not indifferent between those two kinds of uncertainty resolution. For instance, Gneezy and Potters (1997) demonstrated that a subject will invest less in a financial instrument, the more frequently he evaluates its returns (i.e. he prefers one-shot resolution of uncertainty).

Returning to our doctor-patient example, there now arises the question: For a doctor not knowing the patient's preference between the two kinds of uncertainty resolution, how should she send to him a piece of information on his health condition such that both of their utility levels are maximized in a certain sense; and even when she knows his preference but he prefers gradual resolution of uncertainty, which one is the most preferred among a number of different ways for gradually resolving the uncertainty involved?

To answer this question we have first to provide a quantitative description of an individual's preference between one-shot and gradual resolution of uncertainty. For this we shall invoke the theories of Gul (1991) and Dillenberger (2010). According to them an individual preference is completely characterized by a parameter $\beta \in (-1,\infty)$; more precisely, when $\beta \in [0,\infty)$ an individual displays a preference for one-shot resolution of uncertainty, and when $\beta \in (-1,0]$ he displays a preference for gradual resolution of uncertainty. These two preferences and their foundation—a theory of multi-stage lotteries—are discussed in detail in Section 2.

In Section 3, we first describe the game we will study, and then present its extensive form and define the notion of subgame perfect equilibrium. Specifically, there are two players: a doctor and a patient. The doctor has in hand a piece of information on the patient's health condition, and suppose she has K chances to send him this information. Let us call each chance a period. In each period, the doctor sends the patient a message; after receiving the message, he will take an action, which might in turn alter the original information the doctor wants to send. Then in the next period, the doctor, based on this updated information, will send another message, and based on this message the patient will take another action. To illustrate, consider again the above doctor-patient example, which has $K = 2$. In the case of gradual resolution of uncertainty, the message sent by the doctor in period one is that the patient has tumour with probability 0.8, and the one in period two is that the tumour is with an even chance of being benign or malignant. After receiving the message in the first period, the patient may choose to undergo an operation, and this operation is most likely to change his health condition.

To get some intuition of the game and some feeling of how the way of uncertainty resolution will affect players' payoffs, we proceed in Section 4 to examine a special example of the game in which the patient reveals his preference to the doctor and the actions he takes are assumed not to alter the information that the doctor wants to send him at the very beginning. This example, albeit rather simple, is studied also for two other reasons: in the first place it is related to the framing effect mentioned above, and, in the second, it serves as a more specific motivation for the present paper.

Section 5 studies some standard properties of the equilibria of the game; in particular, it is shown that the game has a subgame perfect equilibrium. Finally, Section 6 concludes the paper with a brief remark.

2 Theory of Multi-stage Lotteries

To our knowledge Segal (1990) is the first systematic study of multi-stage lotteries. Here we first give the definition of a multi-stage lottery, and then make a brief review of Gul's theory of disappointment aversion (Gul (1991)) and Dillenberger's preference for one-shot and gradual resolution of uncertainty (Dillenberger (2010)), and finally discuss the relationship between multi-stage lotteries and partitions of a set.

2.1 Definition

Let X be a finite set of a patient's states (such as health, indisposition, or serious disease), and let $\mathbb{R}_+$ be the set of nonnegative real numbers. We assume given on X a utility function $u : X \to \mathbb{R}_+$, mapping each state to its utility level which, for technical reasons and without loss of generality, is supposed to be nonnegative. For the purpose of this paper, only the utility level of a state matters, and so we shall identify X with $u(X)$, and assume henceforth that X be a finite subset of $\mathbb{R}_+$. A one-stage lottery, or more simply a lottery, is a probability distribution on $(X, 2^X)$, where 2^X stands for the power set of X. To simplify the exposition we shall identify each $x \in X$ with the Dirac measure δ_x at x (i.e. the degenerate lottery that yields x with certainty). Let $\mathcal{L}^1$ be the set of all such lotteries. We then define a two-stage lottery to be a probability distribution on $\mathcal{L}^1$ with finite support, and let $\mathcal{L}^2$ be the set of all two-stage lotteries. That is,

$$\mathcal{L}^2 = \{Q : \mathcal{L}^1 \to [0,1] | \; Q(p) > 0 \text{ for finitely many } p\text{'s in } \mathcal{L}^1 \text{and} \sum_{p \in \mathcal{L}^1} Q(p) = 1\}.$$

Inductively, we can define a $(k+1)$-stage lottery to be a probability distribution on $\mathcal{L}^k$ with finite support, where $\mathcal{L}^k$ is the set of all k-stage lotteries, $k = 1, 2, \ldots$.

Throughout this paper we make the following convection: We use small Latin letters such as p, q to designate a generic k-stage lottery, and when we discuss the relationship between a $(k+1)$-stage lottery and a k-stage lottery, we use capital Latin letters such as P, Q to designate the former, and small Latin letters such as p, q to designate the latter. Given $P, Q \in \mathcal{L}^k$ and $\alpha \in (0,1)$ we define $R = \alpha P + (1-\alpha)Q$ to be a lottery in $\mathcal{L}^k$ such that $R(p) = \alpha P(p) + (1-\alpha)Q(p)$ for all $p \in \mathcal{L}^{k-1}$. With this algebraic operation we can define two worthwhile notions: reduction and extension of a k-stage lottery. Let us begin with the simpler case of reduction. It is easily seen that to each $Q \in \mathcal{L}^k$, $k \geq 2$, there corresponds in $\mathcal{L}^{k-1}$ a lottery $\rho_r(Q) = \sum_{p \in \mathcal{L}^{k-1}} Q(p)p$. This $\rho_r(Q)$ we shall call a reduction of Q, and we say that $\rho_r(Q)$ and Q are algebraically equivalent.

To define extension of a k-stage lottery we take $Q \in \mathcal{L}^k$. Let $\mathrm{supp}(Q) = \{p_1, \ldots, p_n\}$ be the support of Q. For any partition, $\Pi = \{\Pi_1, \ldots, \Pi_l\}$, of $\mathrm{supp}(Q)$, we construct $Q_i \in \mathcal{L}^k$, $i = 1, \ldots, l$, such that

$$Q_i(p_j) = \begin{cases} \frac{1}{\tau_i} Q(p_j), & \text{for } p_j \in \Pi_i, \\ 0, & \text{otherwise,} \end{cases}$$

where $\tau_i = \sum_{p_j \in \Pi_i} Q(p_j)$. Let $\rho_e(Q)$ be the lottery that yields Q_i with probability τ_i. It is obvious that $\rho_e(Q) \in \mathcal{L}^{k+1}$, and that $\rho_e(Q)$ and Q are algebraically equivalent. This $\rho_e(Q)$ we shall call an extension of Q. From this construction we see that the cardinality of the support of $\rho_e(Q)$ is equal to that of Π, a fact that will be used in a little while.

2.2 Preferences for One-shot and Gradual Resolution of Uncertainty

We now discuss evaluation of a k-stage lottery, $k = 1, 2, \ldots$. Let us begin with the basic case of $\mathcal{L}^1$. To evaluate a lottery we shall invoke Gul's theory of disappointment aversion, and below we make a brief review of the computational aspect of this theory (see Gul (1991, pp. 684–685)). First, we set the utility level of each Dirac measure $\delta_x \in \mathcal{L}^1$ to be x. Second, for an arbitrary non-degenerate lottery $p \in \mathcal{L}^1$, let $\{x_1, \dots, x_n\}$ be its support. Without loss of generality we assume $x_{j-1} \le x_j$ for $j = 2, \dots, n$.[1] For each $j \in \{1, \dots, n-1\}$, we construct a real number $\alpha_j = \sum_{i=j+1}^{n} p(x_i)$, and a pair of lotteries, denoted by (q_j, r_j):

$$q_j(x) = 0 \text{ for } x \notin \{x_{j+1}, \dots, x_n\},\ q_j(x) = p(x)/\alpha_j \text{ for } x \in \{x_{j+1}, \dots, x_n\}$$
$$r_j(x) = 0 \text{ for } x \notin \{x_1, \dots, x_j\},\ r_j(x) = p(x)/(1-\alpha_j) \text{ for } x \in \{x_1, \dots, x_j\}.$$

Given $\beta \in (-1, \infty)$ let

$$\gamma_\beta(\alpha) = \frac{\alpha}{1 + (1-\alpha)\beta} \text{ for all } \alpha \in [0,1]; \tag{2.1}$$

and let $\hat{V}(\alpha_j, q_j, r_j) = \gamma(\alpha_j) E q_j + (1 - \gamma(\alpha_j)) E r_j$, where E denotes the expectation operator. Then we have the following two cases:

(I) $x_{j^*} < \hat{V}(\alpha_{j^*}, q_{j^*}, r_{j^*}) < x_{j^*+1}$ for some $j^* \in \{1, \dots, n-1\}$;
(II) $\hat{V}(\alpha_{j^*-1}, q_{j^*-1}, r_{j^*-1}) = \hat{V}(\alpha_{j^*}, q_{j^*}, r_{j^*}) = x_{j^*}$ for some $j^* \in \{2, \dots, n\}$.

According to Gul (1991), one and only one of them will occur; moreover, the j^* that satisfies the inequalities in case (I) and the x_{j^*} that satisfies the equalities in case (II) are unique.[2] If case (I) (resp. (II)) occurs then we call p of type (I) (resp. (II)). Let $u_\beta(p)$ be the utility level of p; we define $u_\beta(p) = \hat{V}(\alpha_{j^*}, q_{j^*}, r_{j^*})$. In case (I) let $D(p) = \{(\alpha_{j^*}, q_{j^*}, r_{j^*})\}$ and in case (II) let

$$D(p) = \{(\alpha, q, r) \mid \alpha q + (1-\alpha) r = p \text{ and } q(x) = 0 \text{ if } x \notin \{x_{j^*}, \dots, x_n\},$$
$$r(x) = 0 \text{ if } x \notin \{x_1, \dots, x_{j^*}\}\}.$$

We compile here, but leave their proofs to the appendix, a number of properties of u_β that will be useful later in the following lemma: Let j^* be defined as above; for p^i with support $\{x_1^i, \dots, x_n^i\}$ and p with support $\{x_1, \dots, x_n\}$, let $p^i \rightrightarrows p$ stand for $x_j^i \to x_j$ and $p(x_j^i) \to p(x_j)$ as $i \to \infty$ for all $j = 1, \dots, n$.

LEMMA 2.1. (i) u_β is continuous under the topology generated by the L^1 metric.[3]
(ii) $u_\beta(p) = u_\beta(q)$ implies $u_\beta(p) = u_\beta(\alpha p + (1-\alpha) q)$ for all $\alpha \in (0,1)$ and all $p, q \in \mathcal{L}^1$.
(iii) $u_\beta(p) = \gamma(\alpha) E q + (1 - \gamma(\alpha)) E r$ for any $(\alpha, q, r) \in D(p)$.
(iv) For any p of type (II), $\hat{V}(\alpha_j, q_j, r_j) > x_j$ when $x_j < x_{j^*-1}$ and $\hat{V}(\alpha_j, q_j, r_j) < x_j$ when $x_j > x_{j^*}$.
(v) If $\beta^i \to \beta$ and $p^i \rightrightarrows p$, then $u_{\beta^i}(p^i) \to u_\beta(p)$.

Inductively, we can use Gul's theory to evaluate any k-stage lottery, $k = 2, 3, \dots$. Take a two-stage lottery for example; let $Q \in \mathcal{L}^2$ with its support given by $\{p_1, \dots, p_n\}$. Evaluate each p_i by means of Gul's theory, and let $v_i = u_\beta(p_i)$, $i = 1, \dots, n$. Then we form a (one-stage) lottery q which yields v_i with probability $Q(p_i)$. Let $u_\beta(Q)$ be the utility level of Q; we define $u_\beta(Q) = u_\beta(q)$.

Therefore, we obtain a utility function $u_\beta : \cup_{k \ge 1} \mathcal{L}^k \to \mathbb{R}$ which maps each k-stage lottery to its utility level. With these preliminaries we can introduce Dillenberger's preference for one-shot and gradual resolution of uncertainty.

DEFINITION 1. A patient displays preference for one-shot resolution of uncertainty (PORU) if for any $Q \in \mathcal{L}^2$, $u_\beta(\rho_r(Q)) \ge u_\beta(Q)$. He displays preference for gradual resolution of uncertainty (PGRU) if for any $Q \in \mathcal{L}^2$, $u_\beta(\rho_r(Q)) \le u_\beta(Q)$.

[1] With a view to later application we, as opposed to Gul (1991), allow $x_{j-1} = x_j$ for some j. But this is not essential and all results of Gul (1991) are still valid.

[2] Since it is possible that $x_{j-1} = x_j$, the j^* that satisfies the equalities in case (II) may not be unique.

[3] For the definition of this topology see Gul (1991, Footnote 5, p. 671).

According to this definition and Dillenberger (2010, p. 1989), a patient displays PORU if $\beta \geq 0$ in Eq. (2.1) and PGRU if $-1 < \beta \leq 0$. From this we can see that a patient is completely characterized by the value of β, and henceforth for easier reference, we shall call a patient with $\beta = \beta_0$ of β_0-type.

2.3 Multi-stage Lottery and Partition of X

This subsection is preliminary to Section 3. Let us first recall the definition of a partition (see for example Billingsley (1995, A3, p. 536)).

DEFINITION 2. (i) A partition of X is a finite collection of subsets of X, say $\{X_1, \ldots, X_n\}$, such that $X_i \cap X_j = \emptyset$ for $i \neq j$ and $\cup_{i=1}^n X_i = X$.
(ii) For any two partitions Π_1, Π_2 of X, Π_2 is finer that Π_1, written $\Pi_1 \geq \Pi_2$, if every element of Π_2 is a subset of some element of Π_1.

Remark. Suppose $\Pi_1 = \{X_1, \ldots, X_{n_1}\}$, $\Pi_2 = \{Y_1, \ldots, Y_{n_2}\}$, and $\Pi_1 \geq \Pi_2$. Let $I_i = \{j | \, Y_j \subset X_i\}$; then it is easily checked that $X_i = \cup_{j \in I_i} Y_j$. Moreover, assume given a sequence $\Pi_1 \geq \Pi_2 \geq \cdots \geq \Pi_k$ and let $\mathcal{F}_i$ be the algebra generated by Π_i; then it is easily verified that $\{\mathcal{F}_1, \ldots, \mathcal{F}_k\}$ constitutes a filtration (see Billingsley (1995, p. 458) for its definition). All the development that follows could equivalently be stated in term of filtrations, but as will become clear later, it is more convenient to adopt the language of partitions.

The following construction, depending on the notion of extension of a lottery, is fundamental to later development. Given $p \in \mathcal{L}^1$ and a sequence $\Pi_1 \geq \Pi_2 \geq \cdots \geq \Pi_k$ of partitions of X, we can use this sequence to construct a $(k+1)$-stage lottery that is algebraically equivalent to p in the following fashion. Let $\Pi_i = \{X_1^i, \ldots, X_{n_i}^i\}$, $i = 1, \ldots, k$. Without loss of generality we may assume $\mathrm{supp}(p) = X$, so that Π_i constitutes a partition of $\mathrm{supp}(p)$ (otherwise let $Y_j^i = X_j^i \cap \mathrm{supp}(p)$ and then $\Pi_i' = \{Y_1^i, \ldots, Y_{n_i}^i\}$ is a partition of $\mathrm{supp}(p)$). Let ρ_e be as defined in Section 2.1; using Π_k we can construct $P^2 = \rho_e(p)$, a two-stage lottery that is algebraically equivalent to p. Now consider Π_{k-1}. Let $I_i^{k-1} = \{j | \, X_j^k \subset X_i^{k-1}\}$; by means of the above remark, the set, $\{I_i^{k-1} | \, i = 1, \ldots, n_{k-1}\}$, is a partition of $\{1, \ldots, n_k\}$. Also as remarked at the end of Section 2.1, Π_k and $\mathrm{supp}(P^2)$ have the same cardinality, so that we may assume $\mathrm{supp}(P^2) = \{p_1, \ldots, p_{n_k}\}$. Then to Π_{k-1} there corresponds a partition of $\mathrm{supp}(P^2)$: $\{\{p_j | \, j \in I_i^{k-1}\} | \, i = 1, \ldots, n_{k-1}\}$. Using this partition we can construct $P^3 = \rho_e(P^2)$, a three-stage lottery that is algebraically equivalent to P^2, hence to p. Iterating this process we will obtain a $(k+1)$-stage lottery that is algebraically equivalent to p; let us denote this lottery by $\rho(\Pi_1, \ldots, \Pi_k, p)$.

3 The Model

3.1 The Game

We consider the following game. Suppose that there is a patient of β-type who comes to see a doctor. After diagnosis the doctor gets some knowledge about the patient's health condition, which is represented by a probability distribution, say π, on X. This distribution is then the information that the doctor wants to convey to the patient. We assume that she is endowed with a prior probability distribution, say μ, on the possible types of the patient, and has K chances to communicate with him. Each chance we shall from now on call a period. In each period, the doctor sends to the patient a message about his health condition. After receiving the message, the patient chooses an action from $\{0, 1\}$ (which may refer, for example, to refusing or accepting a treatment), and this action will in turn update π to, say, π' (which means, for example, that after receiving a treatment, his health condition may be improved upon). Let us designate this game by Γ, and assume that all aspects of Γ except π and its update π' are common knowledge.

To make this game more precise we have to address the following questions: (i) what the word 'message' means; (ii) how an action taken by the patient will change the information that the doctor wants to pass on to him; (iii) the payoff structure of Γ. Let us deal with these questions one by one.

For (i), by a message we shall mean a partition of X, and a sequence of messages, a sequence of partitions of X that become increasingly finer. To illustrate and justify this let us consider again the doctor-patient example that opened the introduction. Formally the patient's health condition can be described by the following distribution, denoted by p:

health status	benign tumour	malignant tumour	healthy
probability	0.4	0.4	0.2

Remember that $K = 2$ and that one way for the doctor to inform the patient of the distribution p is this: First tell him that he has tumour with probability 0.8, and then, in the second period, tell him that the tumour is with an even chance of being benign or malignant. Formally this constitutes a two-stage lottery, which can be represented diagrammatically as follows:

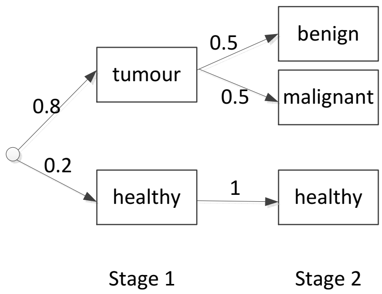

Let tumour $= \{$benign tumour, malignant tumour$\}$. Then the set, $\Pi = \{\{$tumour$\}$, $\{$healthy$\}\}$, is a partition of $\{$benign tumour, malignant tumour, healthy$\}$. Let ρ be as defined at the end of Section 2.3; then the message sent by the doctor in period one is given by $\rho(\Pi, p)$. With regard to a sequence of messages, since the message in the $(k+1)$-st period is to be more informative than the one in the k-th period, it is natural to require messages in the sequence to become ever finer. This solves question (i).

For question (ii), let $\mathcal{T}^K = \times_{i=1}^{K}\{0,1\}$. We can formally represent how a sequence of actions in $\mathcal{T}^K$ updates $\pi \in \mathcal{L}_1$ by a function $f_\pi : \mathcal{T}^K \to \mathcal{L}_1$. For question (iii), we shall discuss it in detail in the next subsection.

3.2 Extensive Form and Equilibrium

We now define the payoff structure of Γ, the notion of mixed-strategy, and the corresponding equilibrium. To this end we have to formulate the game in extensive form. To be concrete and for ease of understanding let us first do this for the doctor-patient game.

For notational convenience, let $x_1 =$ benign tumor, $x_2 =$ malignant tumor, and $x_3 =$ health. Then the set $X = \{x_1, x_2, x_3\}$ has four partitions (excluding X itself): $m_1 = \{\{x_1\}, \{x_2\}, \{x_3\}\}$,

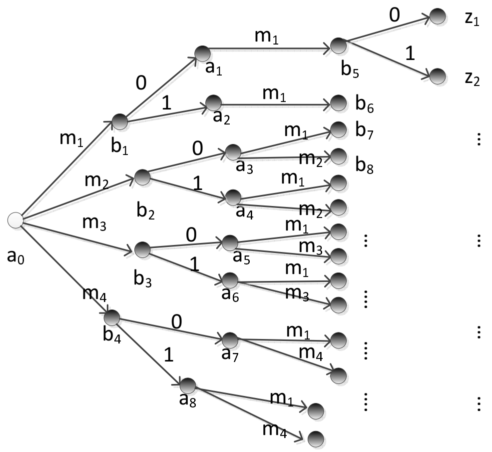

Figure 3.1: Extensive form of the doctor-patient game

$m_2 = \{\{x_1,x_2\},\{x_3\}\}$, $m_3 = \{\{x_1,x_3\},\{x_2\}\}$, $m_4 = \{\{x_1\},$ $\{x_2,x_3\}\}$. Let $F(m) = \{m' |\ m \geq m'\}$; i.e., $F(m)$ is the set of partitions that are finer than m.

Recall that $K = 2$ in the doctor-patient game; its extensive form is given in Figure 3.1. In this figure symbols as $a_0, a_1, \ldots$, represent the doctor's decision nodes, and symbols as $b_1, b_2, \ldots$, represent the patient's decision nodes. The game starts with a_0, at which the doctor can take four actions $m_1, \ldots, m_4$, each of them denoting a message that she can send to the patient. These four actions then lead to four decision nodes for the patient: $b_1, \ldots, b_4$; at every b_i the patient can take either of the two actions: 0 or 1. This in turn leads to eight decision nodes for the doctor: $a_1, \ldots, a_8$. We define a function ϕ_1 on $\{a_1, \ldots, a_8\}$ such that $\phi_1(a_i)$ gives the penultimate action taken to reach a_i. So the actions available at a_i must be members of the set $F(\phi_1(a_i))$, $i = 1, \ldots, 8$. Take a_1 for example; note that $\phi_1(a_1) = m_1$ and $F(m_1) = \{m_1\}$, so the action available at a_1 is m_1 only. Take a_3 for another example; we now have $\phi_1(a_3) = m_2$ and $F(m_2) = \{m_1, m_2\}$, so the actions available at a_3 are m_1 and m_2. Therefore, each of $a_i, i = 1, \ldots, 8$ will yield one or two decision nodes for the patient, which all together are denoted by $b_5, b_6, \ldots$. At each b_i, $i \geq 5$, the patient again takes either of the two actions: 0 or 1; and this in turn leads to the endpoints of the game, given by $z_1, z_2, \ldots$.

It is not hard to extend the above description to any game Γ. Abstractly put and in the terminology of Selten (1975), let T be the game tree for Γ, Z the set of all endpoints, D the set of all vertices of T that are not endpoints, and a_0 the origin of T. In the notation of last paragraph we let $D = A \cup B$, where $A = \{a_0, a_1, \ldots\}$ with each a_i being a decision node for the doctor and $B = \{b_1, b_2, \ldots\}$ with each b_i being a decision node for the patient. For notational convenience we shall call the doctor player one, the patient player two, and write D^1 for A and D^2 for B. Every member of D^i is an information set for player i. Recall that ϕ_1 is a function on D^1 that labels each of its member with the penultimate action taken to reach it. Let C_d be the set of all choices at $d \in D$. It is readily checked that $C_a = F(\phi_1(a))$ for every $a \in D^1$, and $C_b = \{0, 1\}$ for every $b \in D^2$. Moreover, it is easily seen that both players have perfect recall.

Again following Selten (1975), we define a local strategy s_{id} at the information set $d \in D^i$ to be a probability distribution over C_d for $i = 1, 2$. A local strategy s_{id} is called pure if it assigns 1 to one choice at d and 0 to the other choices. A behavior strategy s_i for player i is a function which assigns a local strategy s_{id} to every $d \in D^i$; s_i is called pure if every local strategy s_{id} is. According to Kuhn (1953) and remembering that the game is of perfect recall, mixed-strategies are equivalent to behavior strategies, and therefore we shall restrict our following consideration to the latter, hereafter simply called strategies.

Let S_{id} denote the set of local strategies at $d \in D^i$, S_i the set of strategies for player i, and $S = S_1 \times S_2$ the set of strategies for the game. To complete description of the game there remains to specify the payoffs of the endpoints and of the strategies. For this let $\xi_0(d) = d$, $\xi_1(d)$ be the immediate predecessor of d for $d \neq a_0$, $\xi_n(d)$ the n-th predecessor of d; let $\mathcal{P}$ be the set of partitions of X and $\phi_2 : D \backslash \{a_0\} \to \mathcal{P} \cup \{0,1\}$ a function that maps each non-initial node to the last action taken to reach it (see Kreps and Wilson (1982)). Let $v_1 : Z \to \mathbb{R}_+$ be a function that maps each endpoint to its payoff for the doctor, and v_2 for the patient.

With this notation and for any $z \in Z$ let $m(z) = (\xi_1(z), \xi_3(z), \ldots, \xi_{2K-1}(z))$ and $t(z) = (\xi_2(z), \xi_4(z), \ldots, \xi_{2K}(z$ Then the sequence $t(z)$ of patient's actions will update π into $f_\pi(t(z))$, and the uncertainty of this updated information will be resolved according to the sequence $m(z)$ of patient's messages. This way of uncertainty resolution, by means of the construction at the end of Section 2.3, will lead to a $(K+1)$-stage lottery $\sigma(z) = \rho(m(z), f_\pi(t(z)))$. Under this interpretation, it is reasonable to set

$$v_2(z) = u_\beta(\sigma(z)). \tag{3.1}$$

For the doctor, we assume her, as in Caplin and Leahy (2004), to be entirely empathetic by setting

$$v_1(z) = \int_{-1}^{\infty} u_{\beta'}(\sigma(z)) d\mu(\beta'). \tag{3.2}$$

In particular, when μ is the Dirac measure at β, we have $v_1(z) = v_2(z)$, so that the doctor shares the same utility function with the patient.

To specify payoffs of the strategies note that each $s = (s_1, s_2) \in S$ induces a probability distribution $\mathbf{P}^s$ on Z according to the formula

$$\mathbf{P}^s(z) = \prod_{k=1}^{2K} s_{\iota(k), \xi_k(z)}(\phi_2(\xi_{k-1}(z))), \tag{3.3}$$

where $\iota : \{1, 2, \ldots, 2K\} \to \{1, 2\}$ is defined by $\iota(k) = 2$ for k odd and $\iota(k) = 1$ for k even. Let $V_1 : S \to \mathbb{R}_+$ be a function that maps each strategy to its payoff for the doctor, and V_2 for the patient. Since the patient is assumed in this paper to follow the theory of Gul (1991) (instead of the standard expected utility theory), we set

$$V_2(s) = u_\beta(\mathbf{P}^s),$$

where $u_\beta(\mathbf{P}^s)$ is the utility level of $\mathbf{P}^s$ evaluated according to Gul's theory. And for the doctor, we set as above

$$V_1(s) = \int_{-1}^{\infty} u_{\beta'}(\mathbf{P}^s) d\mu(\beta').$$

With these preparations we can now define the notion of equilibrium. We shall study subgame perfect equilibrium. It is interesting to observe that, as every information set of the game is a singleton, the notion of subgame perfection coincides with that of sequential rationality of Kreps and Wilson (1982). To define it we need the following standard notation: for $s \in S$ let (s_{-id}, s'_{id}) be a strategy in S which is derived from s by replacing s_{id} with s'_{id}.

DEFINITION 3. A strategy $s^* = (s_1^*, s_2^*) \in S$ is a subgame perfect equilibrium for Γ if for every $d \in D^i$, $i = 1, 2$,

$$V_i(s^*) \geq V_i(s^*_{-id}, s_{id}) \text{ for every } s_{id} \in S_{id}.$$

Remark. Again as every information set of the game is a singleton, this definition is equivalent to the usual one, for example, the one of Selten (1975).

4 An Example

To get some intuition of the game and some feeling of the phenomenon that the way of uncertainty resolution would affect the players' payoff, we consider a specific example. Recall that π denotes the information that the doctor wants to convey to the patient. For the sake of simplicity we assume $f_\pi(t) = \pi$ for all $t \in \mathcal{T}^K$, and that the patient reveals his type to the doctor, so that they share the same utility function. We take $K = 2$ and $\pi = [(1,2,3,4);(0.1,0.2,0.3,0.4)]$, i.e. a lottery that yields 1 with probability 0.1 and 2 with probability 0.2, etc.

We examine two cases: $\beta = 0.5$ and $\beta = -0.5$. Let $m_1 = \{\{1\},\{2\},\{3\},\{4\}\}$, $m_2 = \{\{1,2\},\{3,4\}\}$, and $m_3 = \{\{1,2,3\},\{4\}\}$. Since $f_\pi(t) = \pi$ for all $t \in \mathcal{T}^2$, the behavior of the patient has no effect on the payoffs of both players, and so we shall restrict our attention to the doctor. Take as an example the following three ways of uncertainty resolution: (m_1,m_1), (m_2,m_1), and (m_3,m_1), of which the first means that the doctor sends the message m_1 in the both periods, and the other two have a similar interpretation.

When $\beta = 0.5$, it is obvious that (m_1,m_1) is preferred to both (m_2,m_1) and (m_3,m_1), as a nonnegative β implies a PORU. More concretely, a direct calculation shows that the utility level of (m_1,m_1) is 2.826, that of (m_2,m_1) is 2.727, and that of (m_3,m_1) is 2.754. When $\beta = -0.5$, a direct calculation shows that the utility level of (m_1,m_1) is 3.286, that of (m_2,m_1) is 3.387, and that of (m_3,m_1) is 3.381, so that (m_2,m_1) is the most preferred. Additionally, it is interesting to note that (m_3,m_1) is preferred to (m_2,m_1) when $\beta = 0.5$, but this preference relation is reserved when β changes to -0.5.

5 Main Result

Let Δ be the set of Borel probability distributions on $(-1,\infty)$, and let Γ_μ be the game as defined in Section 3 when the doctor's belief on the patient's types is given by $\mu \in \Delta$. The object of this section is to study some standard properties of the equilibria for Γ_μ. As a preliminary, it is easily seen that all properties of Lemma 2.1 continue to be valid when the domain is changed to Z, and therefore we shall refer to the lemma without explicit mention of the underlying domain.

Proposition 5.1. *The game Γ_μ has a subgame perfect equilibrium for every μ in Δ.*

Proof. Before turning to the proof we must first discuss the continuity of V_i, $i = 1,2$. For this note that the set D and every C_d, $d \in D$, are all nonempty and finite, so that we can embed every S_{id}, and hence S, into some finite dimensional Euclidean spaces and impose on them the corresponding Euclidean metrics. We now show that both V_1 and V_2 are continuous with respect to the Euclidean metric. Take $s^k \in S$ with $s^k \to s^0$. Referring to Eq. (3.3), we have $\mathbf{P}^{s^k}(z) \to \mathbf{P}^{s^0}(z)$ for every $z \in Z$. This means that $\mathbf{P}^{s^k} \to \mathbf{P}^{s^0}$ under the L_1 metric, hence that V_2 is continuous, by property (i) of Lemma 2.1.

For the continuity of V_1 let $\bar{x} = \max_{x \in X} x$. From the definition of $\sigma(z)$ (see the line immediately above Eq. (3.1)), there follows $0 \le u_\beta(\sigma(z)) \le \bar{x}$ for all β and all $z \in Z$, hence $0 \le u_\beta(\mathbf{P}^s) \le \bar{x}$ for all β and all $s \in S$. By property (v) of Lemma 2.1 we know that $u_\beta(\mathbf{P}^s)$ is continuous in β for every $s \in S$, hence that $u_\beta(\mathbf{P}^s)$, as a function of β, is Borel measurable. This together with the continuity of V_2 and Lebesgue dominated convergence theorem implies that V_1 is continuous.

We turn now to the proof of the proposition. It is somewhat standard and based on Kakutani's fixed point theorem (see for instance Debreu (1959)). Recall that S_{id} is the set of local strategies at

$d \in D^i$; we define

$$r_{id}(s) = \arg\max_{s'_{id} \in S_{id}} V_i(s_{-id}, s'_{id}), \text{ for all } d \in D^i, i = 1,2,$$

define $r_i : S \to S_i$ by $r_i(s) = \times_{d \in D^i} r_{id}(s)$, the Cartesian product of all r_{id}, and define $r : S \to S$ by $r(s) = r_1(s) \times r_2(s)$. To prove the proposition it is sufficient to show that the correspondence r has a fixed point. For this, it is then, according to Kakutani's theorem, sufficient to check the following four conditions:

(1) The set S is nonempty, compact, and convex.
(2) $r(s)$ is nonempty for all s.
(3) $r(s)$ is convex for all s.
(4) The correspondence r is upper hemi-continuous.

According to the discussion in the first paragraph of this proof, every S_{id} is a nonempty, compact, and convex subset of some finite dimensional Euclidean space, so that Condition (1) is true by Debreu (1959, Sections 1.6, (7) and 1.9, (11)). Condition (2) follows immediately from the compactness of S_{id} and the continuity of V_i. For Condition (3), let $s^k = (s_{-id}, s^k_{id}) \in S$ for any $s \in S$, $s^k_{id} \in S_{id}$, and $d \in D^i$, $i,k = 1,2$; it is readily checked that $\mathbf{P}^{\lambda s^1 + (1-\lambda)s^2} = \lambda \mathbf{P}^{s^1} + (1-\lambda)\mathbf{P}^{s^2}$ for all $\lambda \in [0,1]$. This together with property (ii) of Lemma 2.1 implies that $r_{id}(s)$ is convex, hence Condition (3) is true, again by Debreu (1959, Section 1.9, (11)). As regards Condition (4), let $\varphi_i(s_{-id}) = S_{id}$ for all $s \in S$, so that φ_i is continuous at s_{-id}. Since V_i is continuous on S it follows from Debreu (1959, Section 1.8, (4)) that each r_{id} is upper hemi-continuous, hence that r is upper hemi-continuous, by Debreu (1959, Section 1.8, (3)). This verifies Condition (4). *Q.E.D*

It is easy to observe that the equilibria of Γ_μ depend on μ. To study this dependence, we define $E : \Delta \to S$ to be such that $E(\mu)$ is the set of equilibria of Γ_μ for every $\mu \in \Delta$, and endow Δ with the weak topology.

PROPOSITION 5.2. *The correspondence $E : \Delta \to S$ is upper hemi-continuous.*

PROOF. Let $s^k \in E(\mu^k)$ where $\mu^k \in \Delta$, $k = 1,2,\ldots$, and assume $s^k \to s^0$, $\mu^k \to \mu^0$. It suffices to show that $s^0 \in E(\mu^0)$. First we consider player 2 (the patient); this case is simpler as V_2 is independent of μ. Note that $s^k \in E(\mu^k)$ implies $V_2(s^k) \geq V_2(s^k_{-2b}, t_{2b})$ for every $t_{2b} \in S_{2b}$ and every $b \in D^2$. As argued in the proof of Proposition 5.1, V_2 is continuous; passing to the limit we have

$$V_2(s^0) \geq V_2(s^0_{-2b}, t_{2b}) \text{ for every } t_{2b} \in S_{2b} \text{ and every } b \in D^2. \tag{5.1}$$

Now consider player 1 (the doctor). Let

$$V_1^k(s) = \int_{-1}^{\infty} u_\beta(\mathbf{P}^s) d\mu^k(\beta),\ V_1^0(s) = \int_{-1}^{\infty} u_\beta(\mathbf{P}^s) d\mu^0(\beta);$$

and we have to show that

$$V_1^0(s^0) \geq V_1^0(s^0_{-1a}, t_{1a}) \text{ for every } t_{1a} \in S_{1a} \text{ and every } a \in D^1. \tag{5.2}$$

For this note that $s^k \in E(\mu^k)$ implies $V_1^k(s^k) \geq V_1^k(s^k_{-1a}, t_{1a})$ for every $t_{1a} \in S_{1a}$ and every $a \in D^1$. Therefore it suffices to show that $V_1^k(s^k) \to V_1^0(s^0)$ and $V_1^k(s^k_{-1a}, t_{1a}) \to V_1^0(s^0_{-1a}, t_{1a})$. We shall prove the former only; a similar argument holds also for the latter.

For notational convenience let $g^k(\beta) = u_\beta(\mathbf{P}^{s^k})$ and $g^0(\beta) = u_\beta(\mathbf{P}^{s^0})$. The idea is to use Theorem 3.5 of Serfozo (1982). To reduce the need for frequent cross-referencing we first reproduce some of Serfozo's definitions and a relevant part of his Theorem 3.5 (see Serfozo (1982, pp.383, 388, 390)): The sequence $\{g^k\}$ is uniformly $\{\mu^k\}$-integrable if

$$\lim_{\alpha \to \infty} \sup_k \int_{|g^k| \geq \alpha} |g^k| d\mu^k = 0;$$

this sequence converges continuously to g^0, denoted by $g^k \xrightarrow{c} g^0$, if $g^k(\beta^k) \to g^0(\beta^0)$ for any $\beta^k \to \beta^0$ with $\beta^0 \in (-1,\infty)$. Then Theorem 3.5 of Serfozo (1982) states, in terms of our notation, that if $\mu^k \to \mu^0$, $g^k \xrightarrow{c} g^0$, $V_1^k(s^k) < \infty$ for all k, and $\{g^k\}$ is uniformly $\{\mu^k\}$-integrable, then $V_1^k(s^k) \to V_1^0(s^0)$. Since $\mu^k \to \mu^0$ is assumed *a priori*, to complete the proof, it remains to check the other three conditions.

To see $g^k \xrightarrow{c} g^0$, we take $\beta^k \to \beta^0$. By means of property (v) of Lemma 2.1 and the definition of $\sigma(z)$, we have $u_{\beta^k}(\sigma(z)) \to u_\beta(\sigma(z))$ for all $z \in Z$. Since $s^k \to s^0$ it follows that $\mathbf{P}^{s^k}(z) \to \mathbf{P}^{s^0}(z)$ for every $z \in Z$. Referring again to property (v) of Lemma 2.1 we obtain $g^k(\beta^k) \to g^0(\beta^0)$. This shows $g^k \xrightarrow{c} g^0$.

For the other two conditions, recall that $\bar{x} = \max_{x \in X} x$; again as argued in the proof of Proposition 5.1, we have $0 \le u_\beta(\mathbf{P}^s) \le \bar{x}$ for all β and all $s \in S$, from which follows $0 \le g^k(\beta) \le \bar{x}$ for all β. This then implies that $V_1^k(s^k) < \infty$ for all k and that $\{g^k\}$ is uniformly $\{\mu^k\}$-integrable. So the proof of inequality (5.2) is completed, and therefore the whole proof is completed by combining (5.1) and (5.2). *Q.E.D*

6 Conclusion

Like commodities, information is also a kind of scarce resources. Two of the central issues revolving around it are how much and how information should be transmitted from a better-informed agent to an ill-informed one. While the former issue has been well studied in the literature, this paper concentrates on the latter. We begin the analysis of the issue by formalizing it into a doctor-patient dynamic game, and then show that the game has an equilibrium, and the equilibrium depends, in an upper hemi-continuous manner, upon the doctor's belief μ of the patient's type.

This belief has been assumed throughout the paper to be fixed during the information transmission process. But it is not difficult to notice that the actions taken by the patient must contain some information about his type, and therefore the doctor may well take advantage of this information to update her belief μ. As a future work one may study properties of the game (as for instance existence of equilibrium) in this new situation.

As another future work note that the doctor's belief μ is assumed to be a probability distribution. A great deal of experimental evidence falling under the category of decision making under ambiguity (see for example Ellsberg (1961)), however, suggests that the doctor's uncertainty about the patient's type may not be able to totally captured by a single probability distribution, but instead has to be captured by a set of probability distributions (Gilboa and Schmeidler (1989)) or a capacity (Schmeidler (1989)). So another direction for future work is to study properties of the game in which the doctor has an ambiguous belief on the patient's type.

A Appendix: Proof of Lemma 2.1

The first four properties follow easily from Gul (1991): More specifically, property (i) is a direct consequence of his Axiom 2; property (ii) follows from the remark immediately after his Axiom 3; properties (iii) and (iv) are easily seen from Gul (1991, pp. 684–685).

We turn now to the proof of property (v). Although somewhat lengthy, the basic idea of the proof is fairly simple, namely, to make full use of statements (I) and (II) of Section 2.2. Without loss of generality we may assume that $x_1^i \le x_2^i \le \cdots \le x_n^i$ for all i.[4] Let $(\alpha_j^i, q_j^i, r_j^i)$ and (α_j, q_j, r_j),

[4]This is easily seen to cause no problem for $x_1 < x_2 < \cdots < x_n$, because $x_j^i \to x^i$ for all j. When $x_{j-1} = x_j$ but

$j = 1, \ldots, n-1$, be a series of triples constructed respectively from p^i and p according to the procedure described in Section 2.2. From their construction it follows easily that $\alpha_j^i \to \alpha_j$, $q_j^i \rightrightarrows q_j$, and $r_j^i \rightrightarrows r_j$ for all $j = 1, \ldots, n-1$. Recall that $\hat{V}(\alpha_j, q_j, r_j) = \gamma(\alpha_j)Eq_j + (1-\gamma(\alpha_j))Er_j$ and $\hat{V}(\alpha_j^i, q_j^i, r_j^i) = \gamma(\alpha_j^i)Eq_j^i + (1-\gamma(\alpha_j^i))Er_j^i$. By Eq. (2.1), γ_β is continuous in α and β. This combined with $\beta^i \to \beta$ gives that

$$\hat{V}(\alpha_j^i, q_j^i, r_j^i) \to \hat{V}(\alpha_j, q_j, r_j) \text{ for all } j. \tag{A.1}$$

When p is of type (I), there exists a j_0 satisfying $x_{j_0} < u_\beta(p) < x_{j_0+1}$ and $u_\beta(p) = \hat{V}(\alpha_{j_0}, q_{j_0}, r_{j_0})$. Referring to Eq. (A.1) this means $x_{j_0}^i < \hat{V}(\alpha_{j_0}^i, q_{j_0}^i, r_{j_0}^i) < x_{j_0+1}^i$, hence $u_{\beta^i}(p^i) = \hat{V}(\alpha_{j_0}^i, q_{j_0}^i, r_{j_0}^i)$, for i sufficiently large. Referring again to Eq. (A.1) we have $u_{\beta^i}(p^i) \to u_\beta(p)$.

When p is of type (II), there exists a j_0 such that $u_\beta(p) = \hat{V}(\alpha_{j_0-1}, q_{j_0-1}, r_{j_0-1}) = \hat{V}(\alpha_{j_0}, q_{j_0}, r_{j_0}) = x_{j_0}$. If there exists a j such that $x_j < x_{j_0}$, let $j_1 = \max\{j \mid x_j < x_{j_0}\}$; otherwise let $j_1 = 0$. Similarly, if there exists a j such that $x_j > x_{j_0}$, let $j_2 = \min\{j \mid x_j > x_{j_0}\}$; otherwise let $j_2 = n+1$. Then we have four cases:

(1) $j_1 = 0$ and $j_2 = n+1$,
(2) $j_1 = 0$ and $j_2 < n+1$,
(3) $j_1 > 0$ and $j_2 = n+1$,
(4) $j_1 > 0$ and $j_2 < n+1$.

We shall show that property (v) holds in each case. Let us begin with case (1). In this case we have obviously $u_\beta(p) = x_1 = x_n$. Note that $x_1^i \le u_{\beta^i}(p^i) \le x_n^i$ for every i; passing to the limit we get $u_{\beta^i}(p^i) \to u_\beta(p)$.

For the remaining three cases, their proofs are similar, so let us take case (4) as an example and leave cases (2) and (3) to the interested reader. In case (4) we claim that

CLAIM A.1. $x_{j_1}^i < u_{\beta^i}(p^i) < x_{j_2}^i$ for i sufficiently large.

PROOF. We shall prove $u_{\beta^i}(p^i) < x_{j_2}^i$ only; a similar argument holds also for $x_{j_1}^i < u_{\beta^i}(p^i)$. Assume by contradiction that there exists a subsequence of $\{1, 2, \ldots\}$, which we might as well assume is the sequence itself, such that $x_{j_2}^i \le u_{\beta^i}(p^i)$, $i = 1, 2, \ldots$. We first show that $j_2 < n$. To see this suppose by contradiction that $j_2 = n$. Let $J_i = \{j \mid x_j^i = x_n^i\}$; then

$$\sum_{j \in J_i} p_j^i = 1. \tag{A.2}$$

Let J be the set-theoretic limit superior of $\{J_i\}$, and let $\tau = \min_{j \in J} j$. By the very definition of limit superior, there exists a subsequence $\{J_{i_t} \mid t = 1, 2, \ldots\}$ such that $\tau = \min_{j \in J_{i_t}} j$. This implies that $J_{i_t} = \{\tau, \tau+1, \ldots, n\}$ for all t, hence that $J = J_{i_t}$ and $x_j = x_n$ for every $j \in J$. Using Eq. (A.2) we have $\sum_{j \in J} p_j^{i_t} = 1$. Passing to the limit we get $\sum_{j \in J} p_j = 1$, hence $u_\beta(p) = x_n$. But this contradicts $u_\beta(p) = x_{j_0}$, as $j_2 < n+1$ implies $x_{j_0} < x_n$. So we have $j_2 < n$.

As a consequence, there exists an integer $l \ge 0$ and a subsequence of $\{1, 2, \ldots\}$, which we might again assume is the sequence itself, such that $x_{j_2+l}^i \le u_{\beta^i}(p^i) < x_{j_2+l+1}^i$. From the construction of u_{β^i} it follows that $u_{\beta^i}(p^i) = \hat{V}(\alpha_{j_2+l}^i, q_{j_2+l}^i, r_{j_2+l}^i)$. Passing to the limit and referring to Eq. (A.1) we have $\hat{V}(\alpha_{j_2+l}, q_{j_2+l}, r_{j_2+l}) \ge x_{j_2+l}$. But note that $x_{j_2+l} > x_{j_0}$, and therefore, according to property (iv), $\hat{V}(\alpha_{j_2+l}, q_{j_2+l}, r_{j_2+l}) < x_{j_2+l}$, a contradiction. This shows that $u_{\beta^i}(p^i) < x_{j_2}^i$. *Q.E.D*

From Claim A.1 it follows that every p^i, $i = 1, 2, \ldots$, must belong to one of the following sets:

$x_{j-1}^i > x_j^i$ for some j and some i, we can simply interchange x_{j-1}^i and x_j^i.

- $I^1_{j_3} = \{p^i | \, u_{\beta^i}(p^i) = x^i_{j_3}\}$, where $j_1 < j_3 < j_2$;
- $I^2_{j_4} = \{p^i | \, x^i_{j_4} < u_{\beta^i}(p^i) < x^i_{j_4+1}\}$, where $j_1 \le j_4 < j_2$.

Those sets that are either empty or finite we can simply disregard and restrict our consideration to those that are infinite. From the definitions of j_1 and j_2 it follows that $x_j = x_{j_0}$, hence $(\alpha_j, q_j, r_j) \in D(p)$, for $j_1 < j < j_2$. This implies that $\lim_{i\to\infty, i\in I^1_{j_3}} u_{\beta^i}(p^i) = x_{j_3} = u_\beta(p)$, and, using $u_{\beta^i}(p^i) = \hat{V}(\alpha^i_{j_4}, q^i_{j_4}, r^i_{j_4})$ and Eq. (A.1), that $\lim_{i\to\infty, i\in I^2_{j_4}} u_{\beta^i}(p^i) = \hat{V}(\alpha_{j_4}, q_{j_4}, r_{j_4}) = u_\beta(p)$.

References

Billingsley, P. (1995). *Probability and measure*. Wiley.

Caplin, A. and Leahy, J. V. (2004). The supply of information by a concerned expert. *The Economic Journal*, 114(497):487–505.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.

Debreu, G. (1959). *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*. Yale University Press.

Dillenberger, D. (2010). Preferences for one-shot resolution of uncertainty and allais-type behavior. *Econometrica*, 78(6):1973–2004.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics*, 75(4):643–669.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153.

Gneezy, U. and Potters, J. (1997). An experiment on risk taking and evaluation periods. *The Quarterly Journal of Economics*, 112(2):631–645.

Gul, F. (1991). A theory of disappointment aversion. *Econometrica*, 59(3):667–686.

Kreps, D. M. and Wilson, R. (1982). Sequential equilibria. *Econometrica*, 50(4):863–894.

Kuhn, H. (1953). Extensive games and the problem of information. In andA. Tucker, H. K., editor, *Contributions to the Theory of Games*, volume 2, pages 193–216, Princetion: Princeton University Press.

Schmeidler, D. (1989). Subjective-probability and expected utility without additivity. *Econometrica*, 57(3):571–587.

Segal, U. (1990). Two-stage lotteries without the reduction axiom. *Econometrica*, 58(2):349–77.

Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1):25–55.

Serfozo, R. (1982). Convergence of lebesgue integrals with varying measures. *SankhyÄĄ: The Indian Journal of Statistics, Series A (1961-2002)*, 44(3):380–402.

Tversky, A. and Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458.