

The value of two-person zero-sum repeated games with incomplete information and uncertain duration

Abraham Neyman

International Journal of Game Theory

ISSN 0020-7276

Volume 41

Number 1

Int J Game Theory (2012) 41:195-207

DOI 10.1007/s00182-011-0281-y

International Journal of **Game Theory**
Official Journal of the Game Theory Society

Volume 39
2010
Issues 1–2

Founded by
Oskar Morgenstern

Editor
Shmuel Zamir

Co-Editors
B. von Stengel
R. V. Vohra

Past Editor
William Thomson

Advisory Board
R. J. Aumann
W. Lucas
E. Maskin
R. Myerson
R. Selten
L. S. Shapley

Managing Editor
Romina Goldman

Editorial Board
Pierpaolo Battigalli
Adam Brandenburger
In-Koo Cho
Youngguk Chun
Ezra Einy
Ido Erev
Eduardo Faingold
Andrea Galeotti
Peter Hammerstein
Jean-Jacques Herings
Toru Hokari
Tatsuro Ichishi
Philippe Jehiel
Ehud Kalai
Frédéric Koessler
Yishay Mansour
Leslie Marx
Aki Matsui
Andy McLennan
Dov Monderer
Hervé Moulin
John Nachbar
Clara Ponsati
Klaus Ritzberger
Dinah Rosenberg
Abdolkarim Sadrieh
Dov Samet
William H. Sandholm
Alvaro Sandroni
Eilon Solan
Tamas Solymosi
Marilda Sotomayor
Peter Sudhölter
William Thomson
Tristan Tomala
Fernando Vega-Redondo
Nicolas Vieille
Ravi Vohra
William Zwicker

Special Issue in Honor of Michael Maschler
Guest Editor: Salvador Barberà

EDITORIAL

Memorial
S. Barberà · S. Zamir 1

ORIGINAL PAPERS

Some non-superadditive games, and their Shapley values, in the Talmud
R. J. Aumann 3

Flow sharing and bankruptcy games
E. Bjørndal · K. Jörnsten 11

Repeated games with public uncertain duration process
A. Neyman · S. Sorin 29

Explicit formulas for repeated games with absorbing states
R. Laraki 53

Bargaining among groups: an axiomatic viewpoint
S. Chae · H. Moulin 71

The nucleolus of a standard tree game revisited: a study of its monotonicity and computational properties
M. Maschler · J. Potters · H. Reijnders 89

A characterization of the average tree solution for tree games
D. Mishra · A. J. J. Talman 105

The positive core of a cooperative game
G. Orshan · P. Sudhölter 113

The rights egalitarian solution for NTU sharing problems
C. Herrero · A. Villar 137

(Continued on back cover page)

 Springer

Your article is protected by copyright and all rights are held exclusively by Springer-Verlag. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your work, please use the accepted author's version for posting to your own website or your institution's repository. You may further deposit the accepted author's version on a funder's repository at a funder's request, provided it is not made publicly available until 12 months after publication.

The value of two-person zero-sum repeated games with incomplete information and uncertain duration

Abraham Neyman

Accepted: 10 April 2011 / Published online: 23 June 2011
© Springer-Verlag 2011

Abstract Fix a zero-sum repeated game Γ with incomplete information on both sides. It is known that the value of the infinitely repeated game Γ_∞ need not exist (Aumann and Maschler 1995). It is proved that any number between the minmax and the maxmin of Γ_∞ is the value of a long finitely repeated game Γ_n where players' information about the uncertain number of repetitions n is asymmetric.

Keywords Repeated games · Incomplete information · Uncertain duration

1 Introduction

Two-player repeated games with incomplete information (henceforth, RGII), introduced by Aumann and Maschler (1995),¹ model long-term interactions in which players have asymmetric information about the actual one-shot game that is repeatedly played. Modeling the long-term interactions was focused initially on the infinitely repeated game Γ_∞ and the finitely repeated game Γ_n . Studying the repeated game Γ_n assumes that the number of repetitions n is known to both players, and moreover that n is common knowledge. These assumptions are difficult to justify in many applications of long-term interactions. Neyman and Sorin (2010) studies two-player repeated games where the players have symmetric information about the uncertain number

¹ This book is based on reports by Robert J. Aumann and Michael Maschler which appeared in the sixties in *Report of the U.S. Arms Control and Disarmament Agency*. See “Game theoretic aspects of gradual disarmament” (1966, ST–80, Chapter V, pp. V1–V55), “Repeated games with incomplete information: a survey of recent results” (1967, ST–116, Chapter III, pp. 287–403), and “Repeated games with incomplete information: the zero-sum extensive case” (1968, ST–143, Chapter III, pp. 37–116).

A. Neyman (✉)
Institute of Mathematics, and Center for the Study of Rationality, The Hebrew University of Jerusalem,
Givat Ram, 91904 Jerusalem, Israel
e-mail: aneyman@math.huji.ac.il

of repetitions. The present paper studies the model of the zero-sum RGII where the players have asymmetric information about the number of repetitions.

In the zero-sum infinitely repeated game Γ_∞ , Player 1 (henceforth, P1) can guarantee v if for every $\varepsilon > 0$ he has a strategy σ such that for any sufficiently large number of repetitions n , for each strategy τ of Player 2 (henceforth, P2) the expected average per-stage payoff is at least $v - \varepsilon$. Similarly, P2 can guarantee v in Γ_∞ if for every $\varepsilon > 0$ he has a strategy τ such that for any sufficiently large number of repetitions n , for each strategy σ of P1 the expected average per-stage payoff is at most $v + \varepsilon$. The game Γ_∞ has a uniform value if the maximal payoff that P1 can guarantee, \underline{v} , equals the minimal payoff that P2 can guarantee, \bar{v} .

The definition of the uniform value implies that whenever the uniform value exists, e.g., RGII on one side (henceforth, RGII-OS) (Aumann and Maschler 1995) or stochastic games (Mertens and Neyman 1981), the limit of the values of the finitely repeated games (where payoffs are the average per-stage payoffs) converge to the uniform value as the number (or the expected number in the model with uncertain duration) of repetitions goes to infinity. This limiting result is independent of the information players have on the number of repetitions.

The uniform value need not exist in RGII on both sides (henceforth, RGII-BS) (Aumann and Maschler 1995, Sect. 4.3). In this case, $\underline{v} < \bar{v}$. Nonetheless, v_n , the value of the n -stage RGII-BS (with state-independent signaling) converges to a limit as $n \rightarrow \infty$ (Mertens and Zamir 1971; Mertens 1971), and more generally, v_θ , the value of the finitely RGII-BS (with state-independent signaling) with a random number of repetitions θ and where the players have symmetric information about θ , converges to a limit as the expectation of the number θ of repetitions goes to infinity (Neyman and Sorin 2010). The present paper characterizes the limit points (as $E(\theta) \rightarrow \infty$) of v_θ where players' information about the number of repetitions θ is asymmetric.

This characterization answers two natural questions: (i) what happens in repeated games with asymmetric information about the number of repetitions and (ii) how can one relate all the points in between \underline{v} and \bar{v} to long games?

In RGII, one of finitely many one-shot games is repeatedly played and each player has only partial information about the one-shot game that is being repeated. The RGII (denoted Γ) is described as follows. There is a finite set of normal form games G^m , $m \in M$, with finite action sets I for P1 and J for P2. The state $m \in M$ is chosen at random according to a publicly known probability p , and each player receives partial information about m . The partial information of the players is defined by two functions, $c : M \rightarrow C$ and $d : M \rightarrow D$; P1 observes $c = c(m)$ and P2 observes $d = d(m)$. In addition, after each stage the players obtain some further information about the previous choice of moves.² This is represented by a map Q from $I \times J$ to probabilities on $A \times B$. At stage t , given the state m and the moves (i_t, j_t) , a pair (a_t, b_t)

² This is called state-independent signaling. In more general state-dependent signaling, the players obtain further information about the previous choice of moves and the state.

is chosen at random according to the distribution $Q(i_t, j_t)$.³ A play of the game is thus a sequence $m, i_1, j_1, a_1, b_1, \dots, i_t, j_t, a_t, b_t, \dots$, while the information to P1 before his play at stage t is $c(m), i_1, a_1, \dots, i_{t-1}, a_{t-1}$, and the information to P2 before his play at stage t is $d(m), i_1, b_1, \dots, j_{t-1}, b_{t-1}$. The repeated game is thus represented by the tuple $\Gamma = \langle M, p, M^1, M^2, I, J, G, Q, A, B \rangle$, where M^1 is the partition of M defined by the values of c and M^2 is the partition of M defined by the values of d .

The payoff at stage t of the repeated game, $g_t := G_{i_t, j_t}^m$, depends on the chosen state m and the action pair (i_t, j_t) at stage t . A pair of strategies σ of P1 and τ of P2 in the repeated game Γ defines a probability distribution $P_{\sigma, \tau}$ on the space of plays, and thus a probability distribution on the stream of payoffs g_1, g_2, \dots . The value of the n -stage zero-sum game, v_n , where P1 maximizes the (expectation of the) average $\bar{g}_n := (g_1 + \dots + g_n)/n$ of the payoffs in the first n stages, exists and equals $\max_{\sigma} \min_{\tau} E_{\sigma, \tau} \bar{g}_n$ (where the max is over all strategies σ of P1 and the min is over all strategies τ of P2, and $E_{\sigma, \tau}$ stands for the expectation w.r.t. the probability $P_{\sigma, \tau}$), which by the minmax theorem is equal to $\min_{\tau} \max_{\sigma} E_{\sigma, \tau} \bar{g}_n$.

Special subclasses of RGII are defined by the signaling structure and the initial information about the state. The classical case of standard signaling corresponds to $A = J, B = I$, and to $Q(i, j)$ being the Dirac measure on (j, i) , or equivalently, to $A = B = I \times J$ and to $Q(i, j)$ being the Dirac measure on $((i, j), (i, j))$. RGII-OS corresponds to the case where $c(m) = m$ and $d(m)$ is a constant, or equivalently, only P1 receives a signal about m . Deterministic signaling corresponds to $Q(i, j)$ (respectively, $Q(m, i, j)$ in the state-dependent signaling) being a Dirac measure; in this case we can think of the signal to a player as a deterministic function of (i, j) (respectively, (m, i, j)).

The independent case corresponds to an initial probability p such that the probability defined on $C \times D$ by $p(c, d) = p(\{m : c(m) = c \text{ and } d(m) = d\})$ is a product probability. In this case we may assume without loss of generality that $M = C \times D$ and that the initial probability distribution is a product probability $p \otimes q$ where p is a probability on C and q is a probability distribution on D . Aumann and Maschler (1995, Sect. 4.2) shows that each game with incomplete information in the dependent case is equivalent to a game with incomplete information in the independent case. Therefore, it is sufficient for our main result to handle the independent case, where the statement and the proof of the main result are simplified.

In this paper we study the asymptotic behavior of the value of zero-sum repeated games with an uncertain number of repetitions θ . θ is an integer-valued random variable on a probability space $(\Omega, \mathcal{B}, \mu)$ with finite expectation and each player observes partial information about θ . The normalized value is denoted v_{θ} . We prove that any value between the max min (the maximal payoff that P1 can guarantee) and the min max (the minimal payoff that P2 can guarantee) of Γ_{∞} can be obtained as the value v_{θ} for an asymmetric uncertainty about the number of repetitions θ with arbitrarily large expected duration $E(\theta)$. As any limit point of v_{θ} as $E(\theta) \rightarrow \infty$ is in the interval

³ In state-dependent signaling, Q is a map from $I \times J \times M$ to probabilities on $A \times B$, and at stage t , given the state m and the moves (i_t, j_t) , the pair (a_t, b_t) is chosen at random according to the distribution $Q(m, i_t, j_t)$.

$[\maxmin \Gamma_\infty, \minmax \Gamma_\infty]$, the result characterizes the set of limit points of v_θ as $E(\theta) \rightarrow \infty$.

2 The game model

The model of a RGII-BS of uncertain duration is described by two independent components: the classical RGII-BS Γ and an uncertain duration structure Θ . We define each component separately, and then the RGII-BS of uncertain duration. Thereafter, we recall the (vexcav and cavvex) formulas (as a function of the stage game) of the maxmin and the minmax of Γ , and close the section with the statement of the main result.

RGII-BS is defined in the standard signaling and the independent case by the tuple $\langle C, D, p, q, I, J, G \rangle$, where C, D, I, J are finite sets, p and q are probability distributions on C and D respectively, and G is a list of $I \times J$ two-person zero-sum games $G^{c,d}$, $c \in C$ and $d \in D$. The repeated game proceeds in stages. In stage 0, nature chooses a pair (c, d) with probability $p(c)q(d)$. P1 is informed of c and P2 is informed of d . At stage $t \geq 1$, P1 is first informed of j_{t-1} and then chooses $i_t \in I$, and simultaneously P2 is first informed of i_{t-1} and then chooses $j_t \in J$. The payoff (from P2 to P1) in stage t is $g_t = G_{i_t, j_t}^{c, d}$.

The repeated game is denoted Γ for short, or $\Gamma(p, q)$ to emphasize the dependence on the probability distributions p and q and the fixing of the other parameters C, D, I, J, G that define the repeated game.

A behavioral strategy of P1 in Γ is a map $\sigma : C \times (I \times J)^* \rightarrow \Delta(I)$, where $(I \times J)^*$ stands for all finite strings of $I \times J$ elements, namely, $(I \times J)^* = \cup_{t \geq 0} (I \times J)^t$, and $\Delta(X)$ stands for all probability distributions on X , and a behavioral strategy of P2 is a map $\tau : D \times (I \times J)^* \rightarrow \Delta(J)$. A pair of behavioral strategies, σ of P1 and τ of P2, defines a probability distribution $P_{\sigma, \tau}$ on the space of plays $(c, d, i_1, j_1, i_2, j_2, \dots)$ by $P_{\sigma, \tau}(c, d) = p(c)q(d)$, $P_{\sigma, \tau}(c, d, i_1, j_1) = p(c)q(d)\sigma(c)[i_1]\tau(d)[j_1]$, and by induction on t

$$P_{\sigma, \tau}(c, d, h_t, i_t, j_t) = P_{\sigma, \tau}(c, d, h_t) \sigma(c, h_t)[i_t] \tau(d, h_t)[j_t]$$

for $h_t = (i_1, j_1, \dots, i_{t-1}, j_{t-1}) \in (I \times J)^{t-1}$.

The uncertainty of the number of repetitions θ is modeled as follows. The number of repetitions θ is an integer-valued random variable θ defined on a probability space $(\Omega, \mathcal{B}, \mu)$ and with finite expectation. Before the start of the repeated game the players receive partial information about the value of θ ; P1 observes $s^1(\omega) \in S^1$ and P2 observes $s^2(\omega) \in S^2$, where S^1 and S^2 are finite sets. The interpretation is that at stage 0, nature chooses $\omega \in \Omega$ according to the probability μ , and independently of the choices of nature in the repeated game Γ , the number of repetitions is $\theta(\omega)$, and P1 and P2 are informed of $s^1(\omega)$ and $s^2(\omega)$ respectively. The joint distribution of (θ, s^1, s^2) is assumed to be independent⁴ of the state (c, d) . We call the tuple $\Theta = \langle (\Omega, \mathcal{B}, \mu), \theta, s^1, s^2 \rangle$ an *uncertain duration structure*.

⁴ The more general model, where the duration depends on the state, is obviously of interest. However, restrictive assumptions on the uncertain duration make our main result—Theorem 1—stronger.

The repeated game of uncertain duration Γ_θ is the repeated game Γ , where the choice of P1's (respectively, P2's) action at stage t , i_t (respectively, j_t), may depend in addition on $s^1(\omega)$ (respectively, $s^2(\omega)$). Therefore a strategy σ of P1 in Γ_θ is in fact a list of strategies σ^s ($s \in S^1$) in Γ , and a strategy τ of P2 in Γ_θ is in fact a list of strategies τ^s ($s \in S^2$) in Γ .

The un-normalized payoff in Γ_θ is $\sum_{t=1}^\theta g_t$ ($:= \sum_{t \geq 1} g_t I(t \leq \theta)$ where I stands for the indicator function) and the normalized one is $\frac{1}{E(\theta)} \sum_{t=1}^\theta g_t$. The value of Γ_θ (with the normalized payoff) exists, is denoted v_θ , and equals $\max_\sigma \min_\tau E_{\sigma,\tau,\mu} \frac{1}{E(\theta)} \sum_{t=1}^\theta g_t$ (where the max is over all strategies σ of P1 in Γ_θ , the min is over all strategies τ of P1 in Γ , and $E_{\sigma,\tau,\mu}$ stands for the expectation with respect to the probability $P_{\sigma,\tau,\mu}$ induced on the joint probability of the number of repetitions θ and the play by σ, τ, μ). We are interested in the asymptotic behavior of v_θ as the expected duration $E(\theta)$ goes to ∞ .

Given $p \in \Delta(C)$ and $q \in \Delta(D)$ we denote by $G^{p,q}$ the $I \times J$ stage-payoff matrix $\sum_{c,d} p(c)q(d)G^{c,d}$ and by $u(p, q)$ its minmax value. For $x \in \Delta(I)$, $y \in \Delta(J)$, and an $I \times J$ matrix G we denote by xGy the sum $\sum_i \sum_j x(i)G_{i,j}y(j)$. This is the classical notation for matrix multiplication, where x is the I row vector and y is the J column vector.

Given a compact convex set Y and a bounded function $u : Y \rightarrow \mathbb{R}$ we denote by $\text{cav}_y u$ the smallest concave function from Y to \mathbb{R} that is $\geq u$ and by $\text{vex}_y u$ the largest convex function from Y to \mathbb{R} that is $\leq u$. If $u : \Delta(C) \times \Delta(D) \rightarrow \mathbb{R}$ we denote by $\text{cav}_p u$ the smallest function on $\Delta(C) \times \Delta(D)$ that is concave in p and is not smaller than u at each point (p, q) . Similarly, $\text{vex}_q u$ is the largest function on $\Delta(C) \times \Delta(D)$ that is convex in q and is not larger than u at each point (p, q) . Note that cav_p and vex_q are operators on bounded functions on $\Delta(C) \times \Delta(D)$, and thus can be iterated. The value of the function $\text{vex}_q \text{cav}_p u$, respectively $\text{cav}_p \text{vex}_q u$, at the point (p, q) is denoted $\text{vex}_q \text{cav}_p u(p, q)$, respectively $\text{cav}_p \text{vex}_q u(p, q)$.

P2 can guarantee v in $\Gamma_\infty(p, q)$ if for every $\varepsilon > 0$ there is a strategy τ of P2 and a positive integer N such that for every $n \geq N$ and every strategy σ of P1 we have

$$E_{\sigma,\tau} \frac{1}{n} \sum_{t=1}^n g_t \leq v + \varepsilon$$

Similarly, P1 can guarantee v in $\Gamma_\infty(p, q)$ if for every $\varepsilon > 0$ there is a strategy σ of P1 and a positive integer N such that for every $n \geq N$ and every strategy τ of P2 we have

$$E_{\sigma,\tau} \frac{1}{n} \sum_{t=1}^n g_t \geq v - \varepsilon$$

It follows that if P2, respectively P1, can guarantee v in $\Gamma_\infty(p, q)$, then for every $\varepsilon > 0$ there is N such that for every uncertain duration with $E(\theta) > N$ we have $v_\theta \leq v + \varepsilon$, respectively $v_\theta \geq v - \varepsilon$. If each player can guarantee v in $\Gamma_\infty(p, q)$, then v is called the *uniform value*, or for short a *value*, of $\Gamma_\infty(p, q)$, and is denoted $v_\infty(p, q)$.

Aumann and Maschler (1995), respectively Stearns Aumann and Maschler (1995, Theorem 4.11), proved that P1 can guarantee, respectively cannot guarantee more than, $\text{cav}_p \text{vex}_q u(p, q)$ and that P2 can guarantee, respectively cannot guarantee more than, $\text{vex}_q \text{cav}_p u(p, q)$, and therefore $\Gamma_\infty(p, q)$ has a uniform value iff

$$\text{cav}_p \text{vex}_q u(p, q) = \text{vex}_q \text{cav}_p u(p, q)$$

There are games for which

$$\text{cav}_p \text{vex}_q u(p, q) > \text{vex}_q \text{cav}_p u(p, q);$$

see Aumann and Maschler (1995).

Our main result asserts that for every $\varepsilon > 0$ and $\text{cav}_p \text{vex}_q u(p, q) - \varepsilon \geq v \geq \text{vex}_q \text{cav}_p u(p, q) + \varepsilon$ there is an uncertain duration structure Θ such that $v_\theta = v$ and $E(\theta) > 1/\varepsilon$. Together with the above-mentioned result that for every $\varepsilon > 0$ there is N such that $\text{cav}_p \text{vex}_q u(p, q) + \varepsilon \geq v_\theta \geq \text{vex}_q \text{cav}_p u(p, q) - \varepsilon$ whenever $E(\theta) > N$, we have a complete characterization of the limit points of v_θ as $E(\theta)$ goes to ∞ .

3 Preliminary results

3.1 The posteriors and conditional payoffs

In this section we review a few classical tools used in the analysis of RGII. The space of plays of a RGII-BS with standard signaling is the space of sequences $(c, d, i_1, j_1, i_2, j_2, \dots)$ with the minimal σ -algebra for which all functions $(c, d, i_1, j_1, i_2, j_2, \dots) \mapsto (c, d, i_1, j_1, i_2, j_2, \dots, i_t, j_t)$, $t \geq 0$, are measurable. \mathcal{H}_t denotes the minimal σ -algebra (in fact, an algebra) for which the function $(c, d, i_1, j_1, i_2, j_2, \dots) \mapsto h_t := (i_1, j_1, i_2, j_2, \dots, i_{t-1}, j_{t-1})$ is measurable.

In the following notations and observations we assume the independent case. Let τ be a behavioral strategy of P2 in Γ . We define the functions q_t , $t \geq 1$, from plays to $\Delta(D)$ (called posteriors) by induction on t as follows. $q_1 = q$, and

$$q_{t+1}(d) = \frac{q_t(d)\tau(d, h_t)[j_t]}{\sum_d q_t(d)\tau(d, h_t)[j_t]} \tag{1}$$

Note that q_t is \mathcal{H}_t -measurable.

Lemma 1 *For every strategy σ of P1 and every c, h_t such that $P_{\sigma, \tau}(c, h_t) > 0$, the conditional probability*

$$P_{\sigma, \tau}(d | c, h_t) = q_t(h_t)[d] \tag{2}$$

and thus ($= P_{\sigma, \tau}(d | h_t)$ and) is independent of the strategy σ of P1.

The next lemma is a classical tool in the study of games with incomplete information. It is presented here for completeness. Note that if P is the joint distribution

of $(d, j) \in D \times J$, then $\sum_j P(j) \sum_d |P(d | j) - P(d)| = \sum_{d,j} |P(d, j) - P(d)P(j)| = \sum_d P(d) \sum_j |P(j | d) - P(j)|$. Therefore, if we set

$$\begin{aligned} y_t^d(h_t) &= \tau(d, h_t), y_t(j) = \sum_d q_t(d) \tau(d, h_t) \quad \text{and} \quad \|y_t^d - y_t\| \\ &= \sum_j |y_t^d(j) - y_t(j)| \end{aligned}$$

and apply the above equalities to the conditional distribution of (d, j_t) given \mathcal{H}_t^1 —the algebra spanned by (c, h_t) — we have

Lemma 2

$$E_{\sigma, \tau}(\|q_{t+1} - q_t\| | \mathcal{H}_t^1) = E_{\sigma, \tau}(\|q_{t+1} - q_t\| | \mathcal{H}_t) = \sum_d q_t(d) \|y_t^d - y_t\|$$

3.2 The variation of martingales of probabilities

Lemma 3 *Let $q_t, t = 1, \dots, K + 1$ be a martingale with values in $\Delta(D)$ where D is a finite set. Then*

$$E \sum_{t=1}^K \|q_{t+1} - q_t\| \leq \sqrt{K} \min(\sqrt{2 \log |D|}, \sqrt{|D| - 1}) \tag{3}$$

where $\|q_{t+1} - q_t\| = \sum_{d \in D} |q_{t+1}(d) - q_t(d)|$ and $|D|$ stands for the number of elements of D .

Proof The bound $\sqrt{K} \sqrt{|D| - 1}$ is classical (see, e.g., [Aumann and Maschler 1995](#) and [Mertens et al. 1994](#)), and the bound $\sqrt{2K \log |D|}$ is proved in [Neyman \(2009\)](#). \square

3.3 A strategy of the informed player in RGII-OS

We present here a result of [Aumann and Maschler \(1995\)](#) that is used in the proof of our main result.

Lemma 4 ([Aumann and Maschler 1995](#)) *There is a strategy σ in $\Gamma(p)$ such that for every t and every strategy τ we have*

$$E_{\sigma, \tau} \left(G_{i_t, j_t}^m | \mathcal{H}_t \right) \geq \text{cav}_p u(p)$$

The following implication of this result is used in our analysis of RGII-BS of uncertain duration. Fix a sequence $n_1 < n_2 < \dots < n_K$ and a vector of independent random variables $\vec{c} = c_1, \dots, c_K$, each c_k distributed according to p_k (e.g., as in our application $p_k = p$), whose realization is private information of P1, e.g., generated

by a secret lottery performed by P1. Then, for every strategy τ of P2 in $\Gamma(p, q)$ and every sequence $\hat{q}_k \in \Delta(D)$ where \hat{q}_k is measurable w.r.t. $\mathcal{H}_{n_{k-1}+1}$ (e.g., as in our application, the posteriors of d before the play at stage $n_{k-1} + 1$), there is a strategy σ of P1 such that for every $n_{k-1} < t \leq n_k$ we have

$$E_{\sigma, \tau} G_{i, j_t}^{c_k, \hat{q}_k} \geq \text{cav}_p u(p_k, \hat{q}_k)$$

3.4 Mixing uncertain durations

The next lemma is a trivial (but useful) observation on the value of the repeated games with an uncertain duration structure that is a mixture of (two) uncertain duration structures.

Lemma 5 *For every two uncertain duration structures Θ_1 and Θ_2 and $0 \leq \beta \leq 1$ there is an uncertain duration structure Θ such that $E(\theta) \geq \min(E(\theta_1), E(\theta_2))$ and*

$$v_\theta = \beta v_{\theta_1} + (1 - \beta)v_{\theta_2}$$

Proof Let $\Theta_1 = \langle (\Omega_1, \mathcal{B}_1, \mu_1), \theta_1, s_1^1, s_1^2 \rangle$ and $\Theta_2 = \langle (\Omega_2, \mathcal{B}_2, \mu_2), \theta_2, s_2^1, s_2^2 \rangle$ be two uncertain duration structures. W.l.o.g. we can assume that Ω_1 and Ω_2 are disjoint and that $S_1^i = s^i(\Omega_1)$ and $S_2^i = s^i(\Omega_2)$ are disjoint. For every $0 \leq \alpha \leq 1$, we define the uncertain duration structure $\alpha\Theta_1 + (1 - \alpha)\Theta_2$ as the uncertain duration structure $\Theta = \langle (\Omega, \mathcal{B}, \mu), \theta, s^1, s^2 \rangle$, where Ω is the disjoint union of Ω_1 and Ω_2 , the restriction of s^j ($j = 1, 2$), respectively θ , to Ω_i ($i = 1, 2$) is s_i^j , respectively θ_i , \mathcal{B} consists of all unions $B_1 \cup B_2$ where $B_i \in \mathcal{B}_i$, and $\mu(B_1 \cup B_2) = \alpha\mu_1(B_1) + (1 - \alpha)\mu_2(B_2)$. Then

$$E(\theta) = \alpha E_{\mu_1}(\theta_1) + (1 - \alpha)E_{\mu_2}(\theta_2) \geq \min(E_{\mu_1}(\theta_1), E_{\mu_2}(\theta_2))$$

and

$$v_\theta = \frac{\alpha E(\theta_1)v_{\theta_1} + (1 - \alpha)E(\theta_2)v_{\theta_2}}{E(\theta)} = \beta v_{\theta_1} + (1 - \beta)v_{\theta_2}$$

and note that as α ranges over $[0,1]$ so does $\beta = \beta(\alpha)$. □

4 The main result

Theorem 1 *For every repeated game with incomplete information on both sides, every $\varepsilon > 0$, and every $v \geq \text{cav}_p \text{cav}_q u(p, q) - \varepsilon \geq v \geq \text{cav}_p \text{vex}_q u(p, q) + \varepsilon$, there is an uncertain duration structure Θ with $E(\theta) > 1/\varepsilon$ and such that*

$$v_\theta = v$$

Proof It suffices to prove that

$$\forall \varepsilon > 0 \exists \Theta \text{ with } E(\theta) > 1/\varepsilon \text{ and } v_\theta \geq \text{vex}_q \text{cav}_p u(p, q) - \varepsilon \tag{4}$$

Explicitly, for every $\varepsilon > 0$ there is an uncertain duration structure $\Theta = \langle (\Omega, \mathcal{B}, \mu), \theta, s^1, s^2 \rangle$ such that $E(\theta) > 1/\varepsilon$ and $v_\theta \geq \text{vex}_q \text{cav}_p u(p, q) - \varepsilon$. Indeed, (4) implies by duality⁵ that

$$\forall \varepsilon > 0 \exists \Theta \text{ with } E(\theta) > 1/\varepsilon \text{ and } v_\theta \leq \text{cav}_p \text{vex}_q u(p, q) + \varepsilon \tag{5}$$

The conclusion of the theorem follows from Lemma 5 together with (4) and (5).

We now turn to the proof of (4). Without loss of generality assume that $\max_{c,d,i,j} |G_{i,j}^{c,d}| \leq 1$. Fix $\varepsilon > 0$.

Let K be sufficiently large so that $B := \min(\sqrt{2 \log |D|}, \sqrt{|D| - 1}) \leq \varepsilon \sqrt{K}/3$. Fix a sequence $n_0 = 0 < n_1 < n_2 < \dots < n_K$ with $n_{k-1} \leq \varepsilon n_k/2$. Set $\ell_k = n_k - n_{k-1}$. Let $\mu(\theta = n_k) = \frac{1}{n_k \sum_{k=1}^K 1/n_k}$, $k = 1, \dots, K$. P1 is informed of the value of θ ; P2 is not informed of the value of θ . Note that

$$\forall k \leq K, \quad n_k \mu(\theta = n_k) = \frac{1}{\sum_{k=1}^K 1/n_k} = \frac{E_\mu(\theta)}{K} \tag{6}$$

We prove that for every strategy τ in $\Gamma_\theta(p, q)$ ($p \in \Delta(C)$ and $q \in \Delta(D)$) there is a strategy $\sigma = \sigma(\tau)$ such that

$$g_\theta(\sigma, \tau) \geq \text{vex}_q \text{cav}_p u(p, q) - 3\varepsilon$$

Let τ be a strategy of P2 in Γ_θ . As P2 has no information about the realized value of θ , τ is a strategy in Γ . Let q_t be the posterior of d before the play at stage t . Let $\hat{q}_k := q_{n_{k-1}+1}$ (the posterior of d before the play at stage $n_{k-1}+1$). Note that \hat{q}_k is a function of the strategy τ and the sequence of actions $\hat{h}_k := h_{n_{k-1}+1} = (i_1, j_1, \dots, i_{n_{k-1}}, j_{n_{k-1}})$.

We now define a strategy σ of player 1. Let $\vec{c} = c_1, c_2, \dots, c_K$ be a sequence of C -valued random variables such that conditional on the value of θ they are independent, c_k has distribution p , and for k such that $\theta = n_k$ we have $c_k = c$.

The strategy σ will collate a sequence of strategies σ_k , $k = 1, \dots, K$, by following σ_k in stages $n_{k-1} < t \leq n_k$. The strategy σ_k will depend on $\hat{h}_k := h_{n_{k-1}+1}$ by being a function of \hat{q}_k . By Lemma 4, we can select σ_k to be a strategy of P1 in the repeated game $\Gamma^{\hat{q}_k}(p)$ such that for every strategy $\bar{\tau}$ of P2 in $\Gamma^{\hat{q}_k}(p)$, and every $1 \leq t \leq \ell_k$, we have

$$E_{\sigma_k, \bar{\tau}}(g_t | \mathcal{H}_t) \geq \text{cav}_p u(p, \hat{q}_k) \tag{7}$$

In stage $n_{k-1} + t \leq n_k$ of the repeated game $\Gamma(p, q)$, the behavioral strategy σ of P1 plays the mixed action

⁵ Namely, by reversing the roles of P1 and P2 so that P2 is the maximizer and P1 the minimizer with stage payoff $-g$.

$$\sigma_k (c_k, i_{n_{k-1}+1}, j_{n_{k-1}+1}, \dots, i_{n_{k-1}+t-1}, j_{n_{k-1}+t-1})$$

We define the auxiliary stage payoffs g_t^* as follows. For $n_{k-1} < t \leq n_k$ we set

$$g_t^* = G_{i_t, j_t}^{c_k, d}$$

Recall that $2n_{k-1} \leq \varepsilon n_k$, and note that on $\theta = n_k$ we have $g_t^* = g_t := G_{i_t, j_t}^{c, d}$ for $n_{k-1} < t \leq n_k$. Therefore,

$$\sum_t g_t I(t \leq \theta) \geq \sum_t g_t^* I(t \leq \theta) - \varepsilon \theta$$

and thus

$$E_{\sigma, \tau, \mu} \sum_t g_t I(t \leq \theta) \geq E_{\sigma, \tau, \mu} \sum_t g_t^* I(t \leq \theta) - \varepsilon E(\theta) \tag{8}$$

The definition of σ implies that the conditional distribution of g_1^*, g_2^*, \dots , given θ , is independent of θ .

The definition of σ implies that for every $n_{k-1} < t \leq n_k$ we have

$$E_{\sigma, \tau} (G_{i_t, j_t}^{c_k, \hat{q}_k} | \mathcal{H}_t) \geq \text{cav}_p \mu (p, \hat{q}_k) \tag{9}$$

For every $1 \leq t$ we set

$$y_t^d = \tau(d, h_t) \quad \text{and} \quad y_t = \sum_d q_t(d) y_t^d$$

Recall that $y_t^d = \tau(d, h_t)$ and that $y_t = \sum_d q_t(d) y_t^d$ is measurable w.r.t. \mathcal{H}_t . The play of the strategy σ depends on the realization of \vec{c} . Its play in stages $n_{k-1} < t \leq n_k$ depends only on the value of c_k (which need not be equal to the actual value of c) and therefore (by abuse of notation) we denote

$$x_t^{c_k} = \sigma(c_k, h_t)$$

and for every $n_{k-1} < t \leq n_k$ we denote by p_t the posterior given h_t of c_k .

The definitions of σ , p_t , q_t , y_t^d , and y_t (all as a function of the given strategy τ of P2), together with property (7), imply that for every $n_{k-1} < t \leq n_k$ we have

$$\begin{aligned} E_{\sigma, \tau} (g_t^* | \mathcal{H}_t) &= \sum_c p_t(c) \sum_d q_t(d) x_t^c G^{c, d} (y_t + y_t^d - y_t) \\ &\geq \sum_c p_t(c) \sum_d (\hat{q}_k(d) + q_t(d) - \hat{q}_k(d)) x_t^c G^{c, d} y_t - \sum_d q_t(d) \|y_t^d - y_t\| \\ &\geq \sum_c p_t(c) \sum_d \hat{q}_k(d) x_t^c G^{c, d} y_t - \|q_t - \hat{q}_k\| - E_{\sigma, \tau} (\|q_{t+1} - q_t\| | \mathcal{H}_t) \end{aligned}$$

$$\begin{aligned} &\geq \sum_c p_t(c) x_t^c G^{c, \hat{q}_k} y_t - \|q_t - \hat{q}_k\| - E_{\sigma, \tau}(\|q_{t+1} - q_t\| \mid \mathcal{H}_t) \\ &\geq \text{cav}_p u(p, \hat{q}_k) - \|q_t - \hat{q}_k\| - E_{\sigma, \tau}(\|q_{t+1} - q_t\| \mid \mathcal{H}_t) \end{aligned}$$

where the second inequality uses Lemma 2 and the last inequality uses inequality (9). Therefore, as $E_{\sigma, \tau} \hat{q}_k = q$ and $\text{vex}_q \text{cav}_p u$ is convex in q and $\leq \text{cav}_p u$,

$$E_{\sigma, \tau}(g_t^*) \geq \text{vex}_q \text{cav}_p u(p, q) - E_{\sigma, \tau} \|q_{t+1} - q_t\| - E_{\sigma, \tau} \|q_t - \hat{q}_k\|$$

By the triangle inequality (or equivalently, the convexity of the norm) we have $E_{\sigma, \tau} \|q_t - \hat{q}_k\| \leq E_{\sigma, \tau} \|\hat{q}_{k+1} - \hat{q}_k\|$ and $\|q_{t+1} - q_t\| \leq \|q_{t+1} - \hat{q}_k\| + \|q_t - \hat{q}_k\|$, and therefore, by setting $\eta_k = E_{\sigma, \tau} \|\hat{q}_{k+1} - \hat{q}_k\|$, we have

$$\sum_{n_{k-1} < t \leq n_k} E_{\sigma, \tau}(g_t^*) \geq \ell_k \text{vex}_q \text{cav}_p u(p, q) - 3\eta_k \ell_k$$

and therefore

$$\sum_{1 \leq t \leq n_k} E_{\sigma, \tau}(g_t^*) \geq n_k \text{vex}_q \text{cav}_p u(p, q) - \varepsilon n_k - 3\eta_k n_k$$

Recall (6) and that the distribution of g_t^* is independent of θ . Therefore

$$E_{\sigma, \tau, \mu} \sum_t g_t^* I(t \leq \theta) \geq E(\theta) \text{vex}_q \text{cav}_p u(p, q) - \varepsilon E(\theta) - \frac{3E(\theta)}{K} \sum_k \eta_k \quad (10)$$

By Lemma 3, $\sum_k \eta_k \leq B\sqrt{K}$, where $B = \min(\sqrt{2 \log |D|}, \sqrt{|D| - 1})$. Therefore, as K is sufficiently large so that $3B\sqrt{K} \leq \varepsilon K$, we have

$$E_{\sigma, \tau, \mu} \frac{1}{E(\theta)} \sum_t g_t^* I(t \leq \theta) \geq \text{vex}_q \text{cav}_p u(p, q) - 2\varepsilon$$

which together with (8) completes the proof of (4). □

The next result is a simple corollary of Theorem 1 and the definition of the maxmin and minmax.

Corollary 1 *For every repeated game with incomplete information on both sides, the set of values v_θ , where the uncertainty structure $((\Omega, \mathcal{B}, \mu), \theta, s^1, s^2)$ ranges over all uncertainty structures with $E\theta > N$, converges, as N goes to ∞ , to $[\text{cav}_p \text{vex}_q u(p, q), \text{vex}_q \text{cav}_p u(p, q)]$.*

5 Remarks

There are obviously alternative and/or more general models of asymmetric uncertain duration. For example, the signal sets S^1 and S^2 need not be finite, the signals can be probabilistic, and incremental information can be transmitted to the players in the course of the game and as a function of players' actions. However, each such class of models will include our class of uncertain durations. As our main result asserts that for any v in a given range there is an uncertain duration with a given property, it follows that the more restricted is the class of models of asymmetric uncertain duration, the stronger is the conclusion of Theorem 1. We conclude that Theorem 1 holds for any reasonable class of models of asymmetric uncertain duration.

The other (easy part) of the result, that for every $\varepsilon > 0$ there is N sufficiently large such that $v_\theta \geq \underline{v} - \varepsilon$ (and $v_\theta \leq \bar{v} + \varepsilon$) whenever $E\theta \geq N$, follows from the fact that P1 can guarantee \underline{v} and P2 can guarantee \bar{v} in Γ_∞ . Therefore, this conclusion holds for all possible models of an asymmetric duration uncertainty (where the number of repetitions is independent⁶ of the state). In short, the conclusion of Theorem 1 is independent of the particular choice of modeling asymmetric duration uncertainty.

We turn now to questions motivated by the results of the present paper. A natural question that arises is whether we can characterize the asymptotic conditions on the distribution of θ (with finite expectation) so that independently of players' signals about θ we will have $v_\theta \rightarrow \lim v_n$.

A simple sufficient condition is that $E(\theta) \rightarrow \infty$ and $E(|\theta - E(\theta)| + 1)/E(\theta) \rightarrow 0$. Indeed, if $n(\theta)$ is the integer part of $E(\theta)$ then $|\sum_t g_t I(t \leq \theta) - \sum_{t \leq n(\theta)} g_t| \leq |\theta - n(\theta)| \leq |\theta - E(\theta)| + 1$. Therefore, if $\|G\| := 2 \max_{c,d,i,j} |G_{i,j}^{c,d}|$, an optimal strategy of P1 in $\Gamma_{n(\theta)}$ guarantees in Γ_θ a payoff of at least $v_{n(\theta)} - \frac{\|G\|E(|\theta - E(\theta)| + 1)}{E(\theta)} \rightarrow \lim v_n$ as $\frac{E(|\theta - E(\theta)| + 1)}{E(\theta)} \rightarrow 0$.

Another natural question that arises is the asymptotic characterization of the distributions μ of the number θ of repetitions that when P1 is informed of θ and P2 is not, then the value v_θ is close to the minmax ($\text{vex}_q \text{cav}_p u(p, q)$) of the repeated game $\Gamma(p, q)$. A close look at the proof of the main result reveals a sufficient condition. Given a distribution μ of the uncertain number of repetitions θ and $0 \leq \beta \leq 1$ we define $\theta(\beta)$ to be

$$\inf\{\beta : E_\mu(\theta I(\theta \leq \beta)) \geq \beta E_\mu(\theta)\}$$

Note that $\theta(\beta)$ is monotonic nondecreasing in β and that the distribution μ constructed in our proof obeys $\theta(k/K) = n_k$. We have the following result: for every $\varepsilon > 0$ there is $\delta > 0$ such that for an uncertainty structure where P1 is informed of the value of θ and P2 is not, if $E(\theta) > 1/\delta$ and for every $\beta < 1 - \delta$ we have $\theta(\beta + \delta) > \theta(\beta)/\delta$, then

$$v_\theta \geq \text{vex}_q \text{cav}_p u(p, q) - \varepsilon$$

⁶ The more general model, where the number of repetitions may be correlated to the state, is obviously of interest.

It is also of interest to find out the limit behavior of v_θ for specific classes of asymmetric uncertain durations. Two suggestive examples are when θ is uniformly distributed on $\{1, 2, \dots, n\}$ and when θ has the distribution $P(\theta = n) = (1 - \lambda)\lambda^{n-1}$, and P1 is informed and P2 is not informed of the value of θ . Denote the normalized values by v_{n^*} and v_{λ^*} . What are the limits, if they exist, of v_{n^*} as $n \rightarrow \infty$ and of v_{λ^*} as $\lambda \rightarrow 1$ —?

It is also of interest to study the payoff outcomes of repeated games with incomplete information and uncertain duration where the number of repetitions is known to both players, but not commonly known. A study of such non-zero-sum repeated games with complete information is presented in [Neyman \(1999\)](#).

Acknowledgments This research was supported in part by Israel Science Foundation grants 382/98, 263/03, and 1123/06.

References

- Aumann RJ, Maschler M (1995) Repeated games with incomplete information, with the collaboration of R. Stearns, MIT Press, Cambridge
- Mertens J-F (1971) The value of two-person zero-sum repeated games: the extensive case. *Int J Game Theory* 1:217–227
- Mertens J-F, Neyman A (1981) Stochastic games. *Int J Game Theory* 10:53–66
- Mertens J-F, Zamir S (1971) The value of two-person zero-sum repeated games with lack of information on both sides. *Int J Game Theory* 1:39–64
- Mertens J-F, Sorin S, Zamir S (1994) Repeated games. CORE D.P. 9420, 9421, 9422. CORE, LouvainLa-Neuve
- Neyman A (1999) Cooperation in repeated games when the number of stages is not commonly known. *Econometrica* 67:45–64
- Neyman A (2009) The maximal variation of martingales of probabilities and repeated games with incomplete information, DP 510. Center for the Study of Rationality, Hebrew University, Jerusalem
- Neyman A, Sorin S (2010) Repeated games with public uncertain duration processes. *Int J Game Theory* 39:29–52