

Elasticity of Games

Gilad Bavly*

February 16, 2014

Abstract

We develop an elasticity index of a strategic game. The index measures the robustness of the set of rational outcomes of a game. The elasticity index of a game is the maximal ratio between the change of the rational outcomes and the size of an infinitesimal perturbation. The perturbation is on the players' knowledge of the game.

The elasticity of a strategic game is a nonnegative number. A small elasticity is indicative of the robustness of the rational outcomes (for example, if there is only one player the elasticity is 0), and a large elasticity is indicative of non-robustness. For example, the elasticity of the (normalized) n -stage finitely repeated prisoner's dilemma is at least exponential in n , as is the elasticity of the n -stage centipede game and the n -ranged traveler's dilemma.

The concept of elasticity enables us to look from a different perspective at Neyman's (1999) repeated games when the number of repetitions is not commonly known, and Aumann's (1992) demonstration of the effect of irrationality perturbations.

*Department of Statistics and Operations Research, Tel Aviv University, Israel (<http://giladbavly.weebly.com>).

This work is part of the author's Ph.D. thesis, done under the supervision of Professor Abraham Neyman. I am deeply grateful to Prof. Neyman for his kind and illuminating guidance. The research was supported in part by Israel Science Foundation grants 1123/06 and 1596/10.

1 Introduction

“One can expect agreement between philosophers sooner than between clocks”

Claudius the God: and His Wife Messalina, Robert Graves

This paper concerns the effect that a (small) change in the players’ knowledge about the game has on the rational outcomes of that game. That is, looking at a complete information game as representing an ideal situation, where the game data is common knowledge among the players, we ask, to what extent is the set of equilibria of this game sensitive to changes in that knowledge.

We propose an index, which we call *elasticity*, that measures this sensitivity for every game. The elasticity of a game is defined by the maximal ratio between the change of the equilibrium payoffs and the size of a perturbation of the players’ knowledge. We show that it has some desirable properties: it is always finite for a finite game, and it is determined by the behavior of small perturbations (indeed, we could have alternatively defined elasticity using infinitesimal perturbations).

The elasticity measure can be thought of as the elasticity of a set of outcomes, or “solutions.” Specifically, our definition may be regarded as the elasticity of the correlated equilibrium payoffs of the game. Correlated equilibria represent situations in which the game data are common knowledge, and players may have differential information (possibly correlated) only about payoff-irrelevant events. Therefore, correlated equilibria are the natural benchmark here.

The elasticity is always nonnegative, and a small elasticity is indicative of the robustness of the set of correlated equilibria. An extreme case is a decision problem (namely a one-player game), and indeed the elasticity in this case is 0. As for two-player zero-sum games, if we only allowed perturbations that are themselves zero-sum, then the elasticity would have been 0. As we defined it, the perturbations need not be zero-sum, and the elasticity of zero-sum games is bounded by 2.

Conversely, a large elasticity is indicative of non-robustness. In games such as the repeated prisoner’s dilemma, the centipede game, and the traveler’s dilemma, where the equilibria may be considered non-intuitive or “paradoxical,” elasticity is indeed large.¹

In Section 2 we present Bayesian games, which are used to model uncertainty about the game. Note that in allowing for all Bayesian games, the class of uncertainties is quite general. Section 3 concerns some basic properties of the resulting perturbations of games. Section 4 contains the definition and properties of the elasticity index.

In Section 5, the work of Neyman (1999) on repeated games, where the number of repetitions is not commonly known, is viewed from the perspective of elasticity. We explain how uncertainties of the kind considered there, namely uncertainties about the length of the game, can be transformed and embedded into our uncertainties. This allows us to derive the following corollary from his results: the elasticity of the (normalized) finitely repeated prisoner’s dilemma grows very rapidly (at least exponentially) as the number of repetitions grows.

Section 6 views Aumann (1992), again from the perspective of elasticity. Aumann demonstrates the strong effect that perturbations of rationality may have. We follow this up by defining a parallel notion of elasticity, this time with respect to perturbation of rationality rather than knowledge of the game. It turns out that although conceptually the two parallel notions may be considered quite different, they are essentially equivalent.

2 Uncertainties

Fix a *game form* (N, A) , where $N = \{1, \dots, n\}$ is a finite set of players, and $A = \times_{i \in N} A_i$, where A_i is player i ’s set of actions. An *Interactive Belief*

¹Specifically, elasticity grows rapidly as a function of the size of the exact game (e.g., the number of repetitions in the repeated prisoner’s dilemma, or the number of legs of the centipede).

System (with a common prior) for this game form is a tuple $(p, u, (T_i)_{i \in N})$, where T_i is player i 's set of *types* (or *information states*), and each type $t_i \in T_i$ contains a choice of action $a_i(t_i) \in A_i$; $p \in \Delta(T)$ is the *common prior* probability distribution over $T = \times_{i \in N} T_i$; and $u : T \times A \rightarrow \mathbb{R}^n$ is the payoff. We call a profile of types $t = (t_1, \dots, t_n) \in T$ a *state of the world*.

A payoff function $g : A \rightarrow \mathbb{R}^n$, along with our game form (N, A) , gives a normal form game (N, A, g) . We will also refer to the payoff function g as a *game*, where no ambiguity may arise.

Let $u : T \times A \rightarrow \mathbb{R}^n$ and $v : T \times A \rightarrow \mathbb{R}^n$ be two payoff functions (over the same type space T). A norm $\|\cdot\|$ on \mathbb{R}^n induces a *distance* between u and v , by $E_p(\max_{a \in A} \|u(t, a) - v(t, a)\|)$. Thus, the L_1 -distance is $E_p(\max_{a \in A} (\sum_{i \in N} |u_i(t, a) - v_i(t, a)|))$.

Similarly, the distance between a belief system $S = (p, u, T)$ and a game $g : A \rightarrow \mathbb{R}^n$, is $E_p(\max_{a \in A} \|u(t, a) - g(a)\|)$. Equivalently, it is the distance between u and the type-independent payoff $\tilde{g} : T \times A \rightarrow \mathbb{R}^n$ given by $\tilde{g}(t, a) = g(a)$. When the L_1 -distance² between S and g is δ , we will say that S is an *approximation of order δ* of the game g . We may also write $d(S, g) = \delta$.

In a belief system (p, u, T) , a payoff function $g : A \rightarrow \mathbb{R}^n$ is *common knowledge* among the players, if for every state of the world t and every $a \in A$, $u(t, a) = g(a)$.

Example 2.1. $N = \{1, 2\}$, $A_1 = A_2 = \{D, C\}$. Player 1 has two types, $T_1 = \{c, d\}$, and player 2 has just one type, $T_2 = \{c\}$. The common prior is $p(c, c) = p(d, c) = \frac{1}{2}$. The payoff $u(t, a)$ is given in Figure 1.

What is the L_1 -distance between this belief system and the game g depicted in Figure 2 (“chicken”)? For $t = (d, c)$ the payoffs are the same, i.e., $\forall i \forall a g_i(a) = u_i(t, a)$. For $t = (c, c)$, the maximum difference is 2, attained at $a = (C, D)$, i.e., $\max_{a \in A} (\sum_{i \in N} |g_i(a) - u_i(t, a)|) = |7 - 5| = 2$. Hence the distance is $p(c, c) \cdot 2 = \frac{1}{2} \cdot 2 = 1$.

²Since all norms on a finite space are equivalent, the specific choice of norm is insignificant.

Figure 1:

| | | |
|-----|------|------|
| | D | C |
| D | 0, 0 | 7, 2 |
| C | 2, 5 | 7, 6 |

$t = (c, c)$

| | | |
|-----|------|------|
| | D | C |
| D | 0, 0 | 7, 2 |
| C | 2, 7 | 6, 6 |

$t = (d, c)$

| | | |
|-----|------|------|
| | D | C |
| D | 0, 0 | 7, 2 |
| C | 2, 7 | 6, 6 |

Figure 2: Chicken

We say that player i is *rational* at state t , if his action $a_i(t_i)$ maximizes his expected payoff, given his belief (i.e., his type). That is, $a_i(t_i) \in \arg \max_x [E_p(u_i(t, (x, a^{-i}(t))) | t_i)]$, where $a^{-i}(t) = (a_j(t_j))_{j \neq i}$. We say that a belief system is *rational* if every player is rational at every state (i.e., they play a Bayesian equilibrium).

For a game $\Gamma = (N, A, g)$, define the set of rational distributions, $\text{BE}_\Gamma(\delta) \subseteq \Delta(A)$, to be all distributions over A that are achieved as the action distribution of a rational approximation of Γ , of order $\leq \delta$ (“action distribution” meaning the average over tuples of actions, i.e., $E_p(a(t))$). The set of rational payoffs, $\text{BEP}_\Gamma(\delta)$, is the set of all payoff profiles achieved as the average payoff (i.e., $E_p(u(t, a(t)))$) of such approximations. Obviously, both BE and BEP are monotonically increasing in δ . Note that $\text{BE}(0)$ is just the set of correlated equilibrium distributions, and $\text{BEP}(0)$ the correlated equilibrium payoffs.

3 Properties of Approximations

Similarly to a standard procedure concerning correlated equilibria, we can restrict our attention to systems in which all the types who choose the same action are combined into one.

Proposition 3.1. *For any rational belief system, there exists another rational belief system of the same order or less, with $T_i = A_i$ for every $i \in N$.*

Proof. Let $S = (T, p, u)$ be a rational system. For $a_i \in A_i$, denote $F_i(a_i) = \{t_i \in T_i : a_i(t_i) = a_i\}$ and $F(a) = \{t \in T : a(t) = a\}$. We construct a system in which the types are A_i , by $\hat{S} = (\hat{T}, \hat{p}, \hat{u})$, where $\hat{T}_i = A_i$; player i of type $\hat{t}_i \in A_i$ chooses the action \hat{t}_i (i.e., $a_i(\hat{t}_i) = \hat{t}_i$); $\hat{p}(\hat{t}) = p(F(\hat{t}))$; and for $\hat{t} \in \hat{T}$ and $x \in A$, $\hat{u}(\hat{t}, x) = E_p(u(t, x) \mid t \in F(\hat{t}))$.

S is rational, so for any $t_i \in F_i(a_i)$, the expression $E_p(u_i(t, (x_i, a^{-i}(t))) \mid t_i)$ is maximized by $x_i = a_i$. Therefore $E_p(u_i(t, (x_i, a^{-i}(t))) \mid t_i \in F_i(a_i))$ is also maximized by $x_i = a_i$. To establish the rationality of \hat{S} , we maximize $E_{\hat{p}}(\hat{u}_i(\hat{t}, (x_i, a^{-i}(t))) \mid \hat{t}_i = a_i) = E_{\hat{p}}(E_p(u_i(t, (x_i, a^{-i}(t))) \mid t \in F(\hat{t})) \mid \hat{t}_i = a_i) = E_p(u_i(t, (x_i, a^{-i}(t))) \mid t_i \in F_i(a_i))$. This is the same expression as before, and therefore is maximized by $x_i = a_i$.

$$\begin{aligned} d(\hat{S}, g) &= E_{\hat{p}}(\max_{a \in A} \|\hat{u}(\hat{t}, a) - g(a)\|) = \\ &E_{\hat{p}}(\max_{a \in A} \|E_p(u(t, a) \mid t \in F(\hat{t})) - g(a)\|) \leq \\ &E_{\hat{p}}(\max_{a \in A} E_p(\|u(t, a) - g(a)\| \mid t \in F(\hat{t}))) \leq \\ &E_{\hat{p}}(E_p(\max_{a \in A} \|u(t, a) - g(a)\| \mid t \in F(\hat{t}))) = \\ &E_p(\max_{a \in A} \|u(t, a) - g(a)\|) = d(S, g). \quad \square \end{aligned}$$

Therefore, for any finite game Γ , BE_Γ and BEP_Γ can be defined over belief systems that are finite.

Note that the construction of \hat{S} in the above proof does not depend on g . In fact, the transformation $S \rightarrow \hat{S}$ is a kind of ‘‘coarsening’’ of the system, since it takes the payoffs to be the conditional expectation of payoffs with respect to the partition $\{F(a)\}_{a \in A}$. Such a coarsening is closer than S to any payoff function g .

Corollary 3.2. *If Γ is a finite game, then for any δ , $\text{BE}_\Gamma(\delta)$ and $\text{BEP}_\Gamma(\delta)$ are closed.*

Given two systems $S^1 = (T^1, p^1, u^1)$, $S^2 = (T^2, p^2, u^2)$, we define a convex combination of the two, denoted $\lambda S^1 + (1 - \lambda)S^2$ ($0 \leq \lambda \leq 1$), as the system (T, p, u) , where $T_i = T_i^1 \cup T_i^2$ (assuming that T_i^1 and T_i^2 are disjoint), u is the union of u^1 and u^2 , and $p(t)$ equals $\lambda p^1(t)$ for $t \in T^1$, and $(1 - \lambda)p^2(t)$ for $t \in T^2$.

Lemma 3.3. *If $S = \lambda S^1 + (1 - \lambda)S^2$, then for any game g , $d(S, g) = \lambda d(S^1, g) + (1 - \lambda)d(S^2, g)$.*

Proof. Note that the construction of S is equivalent to the following procedure. We toss a coin with probabilities $(\lambda, 1 - \lambda)$, and accordingly choose either the setting of S^1 , or that of S^2 (and we inform the players of the outcome of the coin toss). Then a direct computation of $d(S, g)$ verifies the result. \square

Corollary 3.4. *For $0 \leq \lambda \leq 1$, and $\delta = \lambda \delta^1 + (1 - \lambda)\delta^2$, $\text{BE}(\delta) \supseteq \lambda \text{BE}(\delta^1) + (1 - \lambda)\text{BE}(\delta^2)$.*

Corollary 3.5. *For any game Γ and any δ , $\text{BE}_\Gamma(\delta)$ and $\text{BEP}_\Gamma(\delta)$ are convex.*

Proof. By Corollary 3.4, $\text{BE}_\Gamma(\delta) \supseteq \lambda \text{BE}(\delta) + (1 - \lambda)\text{BE}(\delta)$. \square

4 Elasticity

Recall that $\text{BEP}_\Gamma(0)$ is just the set of correlated equilibrium payoffs $\text{CEP}(\Gamma)$. Let $d(\text{BEP}_\Gamma(\delta), \text{CEP}(\Gamma))$ be the Hausdorff distance between the two sets, or more simply in our case (since the latter set is a subset of the former) $\max_{x \in \text{BEP}_\Gamma(\delta)} \min_{y \in \text{CEP}(\Gamma)} d(x, y)$.

Definition 4.1. The *elasticity* of a game Γ is

$$\eta(\Gamma) = \sup_{\delta > 0} \frac{d(\text{BEP}_\Gamma(\delta), \text{CEP}(\Gamma))}{\delta} - 1$$

Note that η is nonnegative. For example, if a system S consists of changing the game by adding a constant to every outcome, then this also adds the same constant to the equilibrium payoffs, without affecting the equilibrium strategies. In this case, (equilibrium payoff distance) / (payoff distance) $- 1 = 0$.

Proposition 4.2. (a) If Γ has only one player, then $\eta(\Gamma) = 0$.

(b) If Γ is a two-player zero-sum game,³ then $\eta(\Gamma) \leq 2$.

Proof. (a) Let v be what the player can get in Γ , and let $S = (T, p, u)$ be some system with $d(S, g) = \delta$. Then $|E(u(a(t)) - g(a(t)))| \leq E(|u(a(t)) - g(a(t))|) \leq E(\max_{a \in A} |u(a) - g(a)|) = \delta$. Therefore the expectation in S cannot exceed $v + \delta$, since his current average play would then yield more than v in Γ . And it cannot fall below $v - \delta$, since sticking to his optimal play in Γ (at all types) would then be better in S .

(b) Let v be the value of Γ . Then player 1 has a strategy that guarantees v in Γ . Denote $\delta_i = E(\max_{a \in A} |u_i(a) - g_i(a)|)$ for $i = 1, 2$. Since $\delta = E(\max_{a \in A} |u_1(a) - g_1(a)| + |u_2(a) - g_2(a)|)$, we get that $\delta_1, \delta_2 \leq \delta \leq \delta_1 + \delta_2$. Similarly to (a), the expectation of player 1 in S , $E(u_1)$, is $\geq v - \delta_1$, since that optimal strategy can guarantee that much. Likewise $E(u_2) \geq -v - \delta_2$.

On the other hand, by writing $u = g + (u - g)$, we get that $E(u_1 + u_2) \leq E(g_1 + g_2) + \delta = \delta$. Now let $\alpha_1 = E(u_1) - v$, $\alpha_2 = E(u_2) - (-v)$. So $\alpha_1 \geq -\delta_1$, $\alpha_2 \geq -\delta_2$, and $\alpha_1 + \alpha_2 \leq \delta$. If both numbers have the same sign, then $|\alpha_1| + |\alpha_2| \leq \delta_1 + \delta_2$, and therefore $(|\alpha_1| + |\alpha_2|)/\delta \leq 2$, since $\delta \geq \delta_1, \delta_2$. Otherwise, if, say, $\alpha_1 > 0$ and $\alpha_2 < 0$, then $|\alpha_1| \leq \delta + \delta_2$, and $|\alpha_2| \leq \delta_2$, and therefore $(|\alpha_1| + |\alpha_2|)/\delta \leq (\delta + 2\delta_2)/\delta \leq 3$. \square

The following theorem tells us that in order to establish η , it suffices to consider only small changes in the game data.

Theorem 4.3. For any game Γ ,

$$\eta(\Gamma) = \lim_{\delta \rightarrow 0} \frac{d(\text{BEP}_\Gamma(\delta), \text{CEP}(\Gamma))}{\delta} - 1$$

³There are simple zero-sum games whose elasticity is 2, and this is not a peculiarity of the L_1 norm, as we could also get 2 using the L_∞ norm.

Proof. Denote $\phi(\delta) = d(\text{BEP}_\Gamma(\delta), \text{CEP}(\Gamma))$. First we note that $\phi(\delta)$ is a concave function. Indeed, let $0 \leq \lambda \leq 1$, and let $\delta = \lambda \delta^1 + (1 - \lambda) \delta^2$. Then by Corollary 3.4, $d(\text{BEP}_\Gamma(\delta), \text{CEP}) \geq \lambda d(\text{BEP}_\Gamma(\delta^1), \text{CEP}) + (1 - \lambda) d(\text{BEP}_\Gamma(\delta^2), \text{CEP})$. Now, since $\phi(0) = 0$ and ϕ is concave, the function $\phi(\delta)/\delta$ is decreasing, and the result follows. \square

Theorem 4.4. *For any finite game Γ , $\eta(\Gamma) < \infty$.*

In other words, for any finite Γ there exists a number M , such that for any $\delta \geq 0$ and any $x \in \text{BEP}_\Gamma(\delta)$, $d(x, \text{BEP}_\Gamma(\delta)) \leq M \delta$.

Proof. Let S be a rational belief system of order δ , with average payoff x . Let $z \in \Delta(A)$ be the induced distribution over action profiles. Rationality implies that whenever any type t_i of player i chooses the action $a_i(t_i) \in A_i$, this action is best given t_i . That is, $a_i = a_i(t_i)$ maximizes $E_p(u_i((t_i, t_{-i}), (a_i, a_{-i}(t_{-i}) | t_i)))$.

Let $a_i \in A_i$ be chosen with positive probability. Denote by $F = F(a_i)$ all the types of player i that choose the action a_i . Then $E(u_i(t, a(t)) | F)$ is maximized by the choice of a_i . By writing $u = g + (u - g)$, the maximized expression is $E(g_i(a(t)) | F) + E(u_i(t, a(t)) - g_i(a(t)) | F)$. If we multiply this by $P(F)$, then the first part is equal to $\sum_{a_{-i} \in A_{-i}} z(a_i, a_{-i}) g_i(a_i, a_{-i})$ and the second part is $\leq E(u_i(t, a(t)) - g_i(a(t))) = \delta$. Therefore, for any $b_i \in A_i$, we get that $\sum_{a_{-i} \in A_{-i}} z(a_i, a_{-i}) (g_i(a_i, a_{-i}) - g_i(b_i, a_{-i})) \geq -2\delta$.

Let $\text{CE}(\Gamma)$ be the set of correlated equilibrium distributions. It is the intersection of the halfspaces $H_{a,b}^i = \{\zeta \in \Delta(A) : \sum_{a_{-i} \in A_{-i}} \zeta(a, a_{-i}) (g_i(a, a_{-i}) - g_i(b, a_{-i})) \geq 0\}$, for all $i \in N$ and $a, b \in A_i$.

By Lemma 4.5, there exists M , such that $d(z, \text{CE}) < 2\delta M$. Therefore, $d(\sum_{a \in A} z(a)g(a), \text{CEP}) \leq 2\delta M \cdot \|g(a)_{a \in A}\|_2$. And $d(x, \text{CEP}) \leq d(\sum_{a \in A} z(a)g(a), \text{CEP}) + \delta$. \square

Lemma 4.5. *Let $\emptyset \neq C = \bigcap_{j=1}^k H_j$ be a finite intersection of halfspaces, where $H_j = \{x \in \mathbb{R}^n : \langle e_j, x \rangle \geq 0\}$, $e_j \in \mathbb{R}^n$. Then there exists a number M such that for any $z \in \mathbb{R}^n \setminus C$, $d(z, C) \leq M \cdot \max\{-\langle e_j, z \rangle : 1 \leq j \leq k\}$.*

Proof. Fix $z \in \mathbb{R}^n \setminus C$. Denote $V = \text{span} \{e_j : \langle a_j, z \rangle \leq 0\}$, and choose an independent set of vectors $\{e_j\}$ that spans V ; w.l.o.g. it is $\{e_1, \dots, e_m\}$. Let $\phi(z)$ be the closest point to z in C , and denote $v = z - \phi(z)$. Then $v \in V$, because if $v = v_1 \oplus v_2$, where $v_1 \in V$ and $v_2 \in V^\perp$, then for some $\varepsilon > 0$ $\phi(z) + \varepsilon v_2$ is still in C , and is closer to z .

Since (e_1, \dots, e_m) is a basis for V , it follows that if we define $N : V \rightarrow \mathbb{R}$ by $N(v) = \max\{-\langle e_i, v \rangle : 1 \leq i \leq m\}$, we can verify that N is a norm. Since all norms are equivalent, there exists $K = K(e_1, \dots, e_m)$ with $\|v\| \leq K \cdot N(v) = K \cdot \max\{-\langle e_j, v \rangle : 1 \leq j \leq m\}$. By taking M to be the maximum over all $K(E)$, where E is any linearly independent subset of $\{e_1, \dots, e_k\}$, we get the result. \square

There is, however, no universal bound. Even if we fixed the game form and bound all payoffs to a certain range, we would still find games whose elasticity is as high as we please.

5 Large Elasticity in Repeated Games

Neyman (1999) considers finitely repeated games, where the number of repetitions T is not common knowledge. He shows that with a very small deviation from common knowledge of the length of the repeated game (i.e., common knowledge of the proposition $T = n$), we get equilibrium payoffs that approximate every feasible and strictly individually rational payoff (of the stage game). Thus, for example, in the (normalized) finitely repeated prisoner's dilemma we get an equilibrium whose payoff is close to the full cooperation payoff. There are a few senses in which the required deviation from common knowledge is small. In particular, the expectation of $|T - n|$ is exponentially small in n .

With this class of uncertainties, namely uncertainties about the length of a repeated game, we may define the *L-elasticity* of the game as the maximal ratio between the expected change of the equilibrium payoff and the expected change of the length ($E|T - n|$). In this terminology, Neyman's results say,

in particular, that the scale of growth of the L-elasticity of the repeated prisoner's dilemma is exponential in n .

Proposition 5.1 (Neyman (1999)). *There exist numbers $B \geq A > 1$, such that the L-elasticity of the normalized n -repeated prisoner's dilemma is larger than A^n and smaller than B^n .*

As we will see, the “L-uncertainties,” namely the uncertainties about the length of the game, can be seen as a subclass of the class of uncertainties employed by the elasticity concept.

In order to consider various lengths T of the game, Neyman takes the set of strategies of each player to be the strategies of the infinitely repeated game. Two strategies that completely agree on the first n stages are equivalent in the original n -repeated game.

We can use Neyman's result to prove that the elasticity of the n -repeated prisoner's dilemma grows at least exponentially in n . To do this, we translate the uncertainty about the length of the game T into uncertainty about payoffs. Simply, for every pair of strategies of the infinitely repeated game, a value of T determines the payoffs of the players.

We should also account for the shift from n -stage strategies to infinite strategies.

Definition 5.2. The *reduced form* of a game in normal form (N, A, g) is a game (N, \hat{A}, g) , where for each $i \in N$, \hat{A}_i are the equivalence classes of A_i (i.e., $a_i \sim a'_i \iff g(a_i, a^{-i}) = g(a'_i, a^{-i})$ for every $a^{-i} \in A^{-i}$). Two games are *isomorphic* if they have the same reduced form.

Thus, the original n -stage game is the reduced form of the game with infinite n -equivalent strategies. The following lemma tells us that the two games will have the same elasticity.

Let $\text{BE}(\delta)/\sim$ denote the projection of $\text{BE}(\delta)$ over the equivalence classes of actions.

Lemma 5.3. *Let Γ be a game, and $R(\Gamma)$ its reduced form. Then $\text{BE}_\Gamma(\delta)/\sim = \text{BE}_{R(\Gamma)}(\delta)/\sim$ and $\text{BEP}_\Gamma(\delta) = \text{BEP}_{R(\Gamma)}(\delta)$, for any δ .*

Proof. Let $\Gamma = (N, A, g)$, $R(\Gamma) = (N, \hat{A}, g)$. Let $\varphi_i : A_i \rightarrow \hat{A}_i$ be the function that associates an action $x_i \in A_i$ with its equivalence class $\varphi_i(x_i)$. Let $\psi_i : \hat{A}_i \rightarrow A_i$ be some arbitrary function that associates every equivalence class $\hat{x}_i \in \hat{A}_i$ with a “representative” of that class $\psi_i(\hat{x}_i) \in A_i$ (i.e., $\varphi_i(\psi_i(\hat{x}_i)) = \hat{x}_i$).

Suppose $S = (T, p, u)$ is a rational belief system over (N, A) . For each $t_i \in T_i$, define the function $f_i^{(t_i)} : A_i \rightarrow A_i$ by

$$f_i^{(t_i)}(x) = \begin{cases} \varphi_i(\psi_i(x)) & \text{if } x \neq a_i(t_i) \\ a_i(t_i) & \text{if } x = a_i(t_i) \end{cases}$$

That is, if x is the action chosen by the type t_i , we keep it; otherwise, we replace x with the fixed representative of its equivalence class.

First, we construct the system $S^e = (T, p, u^e)$ where the types and prior are the same T and p as in S , and $u^e(t, (a_i)_{i \in N}) = u(t, (f_i(a_i))_{i \in N})$. Here two actions that are equivalent (i.e., according to g) are also equivalent according to u^e .

If $x \in A_i$ is not equivalent to $a_i(t_i)$, then $E(u_i^e(t, (x, a^{-i}(t))) | t_i) = E(u_i(t, (f_i(x), f^{-i}(a^{-i}(t)))) | t_i) = E(u_i(t, (f_i(x), a^{-i}(t))) | t_i) \leq$ [because S is rational] $E(u_i(t, (a_i(t_i), a^{-i}(t))) | t_i) = E(u_i^e(t, (a_i(t_i), a^{-i}(t))) | t_i)$. So S^e is rational.

For any $t \in T$, $\max_{a \in A} \|u^e(t, a) - g(a)\| = \max_{a \in A} \|u(t, f(a)) - g(a)\| = \max_{x \in f(A)} \|u(t, x) - g(x)\| \leq$ [because $f(A) \subseteq A$] $\max_{x \in A} \|u(t, x) - g(x)\|$. Therefore $d(S^e, g) \leq d(S, g)$.

Now consider the system over (N, \hat{A}) defined by $\hat{S} = (T, p, \hat{u})$, where $\hat{u}(t, \hat{a}) = u^e(t, \psi(\hat{a}))$. Since equivalent actions are equivalent in S^e , it follows that \hat{S} is still rational and that $d(\hat{S}, g) = d(S^e, g) (\leq d(S, g))$.

The other direction is simpler: Any belief system over (N, \hat{A}) can be directly copied into a belief system over (N, A) , simply by taking all actions that are equivalent (according to g) to be equivalent in the belief system (i.e., equivalent according to u). \square

Corollary 5.4. *If Γ_1 and Γ_2 are isomorphic games, then $\text{BE}_{\Gamma_1}(\delta)/\sim = \text{BE}_{\Gamma_2}(\delta)/\sim$ and $\text{BE}_{\Gamma_1}(\delta) = \text{BE}_{\Gamma_2}(\delta)$, for any δ .*

We are thus able to derive the following from Neyman's result:

Proposition 5.5. *There exists a number $A > 1$, such that the elasticity of the normalized n -repeated prisoner's dilemma is greater than A^n .*

6 Uncertainty about Rationality

Aumann (1987) shows that when the rationality of all the players is common knowledge, then the distribution of their action profile is a correlated equilibrium distribution. Indeed, common knowledge of rationality may be considered a strong assumption. As Aumann (1992) writes: "In real interactive situations there is a great deal of uncertainty about what others will do, to what extent they are rational, what they think about what you think and about your rationality, and so on." In the same paper, he sets out to demonstrate what a rational player might do when rationality is not commonly known. He shows how a slight departure from common knowledge of rationality can have a big strategic effect – the players' actions can deviate sharply from the correlated equilibria. Such a state of affairs can account for human behavior in well-known "backward induction paradoxes," such as centipede games and the finitely repeated prisoner's dilemma.

Thus in Aumann's model, in the vast majority of "states of the world" the players of the centipede game are rational. Moreover, in the vast majority of states, rationality is mutually known (i.e., each player knows that the other player is rational). And yet, in the vast majority of states, the players do not "go out" immediately at the start of the game, but rather they "stay in" for a few rounds. So this staying-in behavior, which cannot occur when rationality is common knowledge, is possible in this "almost common knowledge" situation.

In particular, in Aumann's examples each player's "expected irrationality" is small. That is, the irrationality of some type of player i is taken to be the difference between his expected payoff and what a rational player might have expected in his place. Player i 's expected irrationality is then

the weighted average of the irrationality of his types. Note that it does not seem enough to consider only the probability of irrational types; we want to account for the “extent of irrationality” that each type exhibits.

This concept of expected irrationality can be used to define a parallel notion of elasticity with respect to uncertainty about rationality, instead of uncertainty about the game, as follows. In a belief system S , let $L_i(t_i) = \max_{x \in A_i} (E_p(u_i(t, (x, a^{-i}(t)))) - u_i(t, a(t)) \mid t_i)$ be the irrationality of type t_i . If t_i plays optimally given his belief, then $L_i(t_i) = 0$. The irrationality of the system, denoted $\mathcal{I}(S)$, is then $\sum_{i \in N} E_p(L_i(t_i))$, namely the sum of the expected irrationality of the players.⁴ For a game $\Gamma = (N, A, g)$, define $\text{IRR}_\Gamma(\delta)$ to be the action distributions of any system S where the payoff g is common knowledge, and $\mathcal{I}(S) \leq \delta$.

Although uncertainty about rationality is conceptually quite different from uncertainty about the game, the following proposition states that all the distributions achieved with some irrationality bound, can also be achieved by rational systems with a deviation of the same order of magnitude (and vice versa).

Proposition 6.1. *For any game Γ , $\text{IRR}_\Gamma(2\delta) \subseteq \text{BE}_\Gamma(\delta) \subseteq \text{IRR}_\Gamma(2\delta n)$.*

Proof. Let $\Gamma = (N, A, g)$. Let $S = (T, p, u)$ be a rational system, with $d(S, g) = \delta$. Consider the system $\hat{S} = (T, p, g)$, where the payoff g is common knowledge, and the types and prior are the same T and p . By writing $g = u + (g - u)$ we write:

$$g_i(x, a^{-i}(t)) - g_i(a(t)) = [u_i(t, (x, a^{-i}(t))) - u_i(t, (a(t)))] + g_i(x, a^{-i}(t)) - g_i(a(t)) - u_i(t, (x, a^{-i}(t))) + u_i(t, (a(t))).$$

To compute $L_i(t_i)$ we need to take the conditional expectation of this, given t_i . Then for the first parenthesis we get ≤ 0 , because of the rationality of S . Thus $L_i(t_i) = E(g_i(x, a^{-i}(t)) - g_i(a(t)) \mid t_i) \leq E(g_i(x, a^{-i}(t)) - u_i(t, (x, a^{-i}(t))) + u_i(t, (a(t))) - g_i(a(t)) \mid t_i) \leq E(2 \max_{a \in A} |g_i(a) - u_i(t, a)| \mid t_i)$.

⁴A reasonable alternative is to define it as $\max_{i \in N} E_p(L_i(t_i))$.

Hence, $E(L_i(t_i)) \leq 2 E(E(\max_{a \in A} |g_i(a) - u_i(t, a)|) | t_i)$
 $= 2 E(\max_{a \in A} |g_i(a) - u_i(t, a)|) \leq 2 E(\max_{a \in A} \|g(a) - u(t, a)\|) = 2 d(S, g) = 2\delta$. And therefore the irrationality of \hat{S} is $\leq 2\delta n$.

To prove the other inclusion, let $S = (T, p, g)$ with irrationality δ . Denote the function $L_i(t_i)$ for this system by $L_i^S(t_i)$, and define u as follows. If $x_i \neq a_i(t_i)$, then $u_i(t, x) = g_i(x) - L_i^S(t_i)/2$; and if $x_i = a_i(t_i)$, then $u_i(t, x) = g_i(x) + L_i^S(t_i)/2$. Consider the system $\hat{S} = (T, p, u)$. The definition of u ensures that $a_i(t_i)$ maximizes i 's payoff given his belief t_i , namely every type t_i is rational. For any $t \in T$, $a \in A$, and $i \in N$, $|u_i(t, a) - g(a)| = L_i^S(t_i)/2$, therefore also $\|u(t, a) - g(a)\| = L_i^S(t_i)/2$, and so
 $d(\hat{S}, g) = E(\max_{a \in A} \|u(t, a) - g(a)\|) = E(L_i^S(t_i)/2) = \delta/2$. \square

This allows us to easily translate results concerning uncertainty about the game to results concerning uncertainty about rationality, and vice versa.

7 Remarks

7.1 Extending the Distance to Bayesian Games

We have defined elasticity by looking at normal-form games, and in particular the distance between them and their approximations, Bayesian games. It seems natural then to try and extend our space of games to Bayesian games, in which normal-form games are included as a degenerate case, and have some measure of distance between any two points in this space, i.e., two Bayesian games. The definition is simple for two games with the same set of types, as in Section 2; but we want it for any two games.

So consider two Bayesian games $S_1 = (T_1, p_1, u_1)$ and $S_2 = (T_2, p_2, u_2)$. The following definition might seem appealing. Define their distance by $\min E(d(u_1, u_2))$, where the minimum is taken over all worlds that combine S_1 and S_2 (i.e., a probability p over $T_1 \times T_2$, with the appropriate marginals). Note that this amounts to measuring the distance between the two distributions over $(\mathbb{R}^N)^A$, namely the distributions of $u_i(\cdot)$ induced by p_i ($i = 1, 2$).

The reason why such a definition is inappropriate is that it does not reflect the individual information of each player. For example, consider the two Bayesian games in Figure 3. There are two possible states of nature, A

Figure 3:

| | | |
|----------|----------|----------|
| | γ | δ |
| α | A | B |
| β | B | A |

| | | |
|----------|----------|----------|
| | γ | δ |
| α | A | B |
| β | A | B |

and B , and let $d(A, B) = 1$. In each of the games the types of player 1 are α, β and the types of player 2 are γ, δ . Let the prior in both cases be $1/4$ for each state of the world.

Here $\min E(d(u_1, u_2)) = 0$, because we can choose a world that combines the two games where the state of nature, either A or B , is identical in the first and the second games.

We believe that the distance should be defined by the types of the players, and not just by the payoffs, as follows. Take the minimum of $E_p(d(u_1, u_2))$ over all combinations $(p, T_1 \times T_2)$ that embed the two games: namely, the marginal distribution of p over T_1 equals p_1 , and the marginal over T_2 equals p_2 , but also with the following conditional independence properties: for every player i , the type t_2^i is independent of t_1^{-i} given t_1^i , and also independent given t_2^{-i} (and likewise when replacing 1 and 2.)

If we think of a normal-form game as a Bayesian game with only one possible type profile, then we immediately see that this definition of distance between Bayesian games extends our previous definition of distance between a normal-form game and a Bayesian game.

In the above example, we can construct a combination, with $E_p(d(u_1, u_2)) = 1/2$, as follows. Let the states of the world be denoted by the tuple $(t_1^1, t_2^1, t_1^2, t_2^2)$. p is constructed so that the two types of each player are identical (i.e., $t_1^i = t_2^i$) in all possible states of the world, and $p(\alpha, \alpha, \gamma, \gamma) = p(\alpha, \alpha, \delta, \delta) = p(\beta, \beta, \gamma, \gamma) = p(\beta, \beta, \delta, \delta) = 1/4$. Then with probability $1/2$

the states of nature are identical (and then $d(u_1, u_2) = 0$), and with probability $1/2$ they are different, and then $d(u_1, u_2) = 1$, and hence $E_p(d(u_1, u_2)) = 1/2$.

This combination satisfies the stated conditions. First, the marginal distribution over T_1 is $1/4$ for each combination, as required, and similarly for T_2 . As for the conditional independence conditions, take for example player 1's type in the first game and player 2's type in the second game, t_1^1 and t_2^2 , given player 1's type in the second game, t_2^1 . t_1^1 is completely determined given t_2^1 , and therefore, conditional on t_2^1 , t_1^1 and t_2^2 are trivially independent.

Now we show that, in fact, any combination satisfying these conditions has $E_p(d(u_1, u_2)) = 1/2$. Consider any such combination, and first let us confine our attention to the event $E = \{t_1^2 = \delta\}$. t_1^1 and t_2^2 are independent given E . Note that u_2 is always determined by t_2^2 , and that given E , u_1 is determined by t_1^1 . Therefore, u_1 and u_2 are independent given E . Now, the marginal distribution over T_1 has to be the same as in the first game; therefore it follows that given E the distribution of t_1^2 is either γ or δ with equal probability, and hence $u_1 = A$ or B with equal probability. u_2 also gets the values A or B , and is independent of u_1 ; therefore u_1 and u_2 coincide or differ with equal probability. Hence $E_p(d(u_1, u_2) | E) = 1/2$.

We can follow the same argument for the event $\{t_1^2 = \gamma\}$, and thus get that overall $E_p(d(u_1, u_2)) = 1/2$. Therefore, the distance between the two games in our example is $1/2$.

7.2 Individual scales

It is not hard to see that elasticity is invariant to adding a constant to any player's payoff function. It is also invariant to multiplying the payoffs of all players by a constant; but multiplying the payoff of a single player by a constant may change the elasticity.

One way to make it invariant might be to define it so that given a game, we first rescale any player's payoff function in the usual sense (positive linear affine transformation), so that each player's maximal payoff is 1 and the

minimal is 0, and then compute the elasticity. Then if some player's payoff in game A is a multiple of his payoff in game B , the two games will still have the same elasticity. This may seem satisfactory, at least in some cases.

It may be argued that either way the notion involves a comparison of individual utility scales, which some may find problematic. The problem lies not in the concept of utility comparison in general (in fact, this author believes that in many situations utility comparison cannot, or at least should not, be avoided), but rather in the fact that games are often perceived as modelling ordinal utilities and not cardinal utilities.

Yet, if one tries to analyze, through the elasticity concept or through other means, the consequences of an uncertainty of the *modeler* (about the knowledge of the players), then the specific numbers that the modeler ascribes to the different players may not be that meaningless.

Suppose our modeler writes down the prisoner's dilemma, with player 1's payoffs in the four cells being $(0, 1, 4, 5)$ and player 2's payoffs being $(0, 100, 400, 500)$. Why does she choose those particular values? If she is not ready to consider any uncertainty, any kind of perturbation, then perhaps nothing can be deduced from her specific choice.

But if she is taking perturbations into account, we may have to take her numbers seriously. The different scales of 1 and 2 may reflect her ability to discern the preferences of player 2 more clearly than those of player 1. Or they may relate to something like 1's uncertainty about 2's uncertainty about 1's discerning 2's preferences, etc.

In short, at some possible state of the world, player 2's preferences seem to some player (or to the modeler) a lot clearer than player 1's. That degree of clarity is manifested in the difference between 2's payoff in various outcomes, compared to 1's differences. And indeed, this is invariant to adding a constant to a player's payoffs, but not invariant to a positive multiple.

We could, then, say that the numbers do not involve a comparison between the utilities of two agents, but rather they are all in the mind of one "agent," the modeler. And they reflect what she knows about the preferences

of the players, what she knows about their beliefs about the preferences, and so on.

7.2.1 Example

An extreme example is the game in Figure 4, where if we multiplied the payoff of one of the players by some factor, the elasticity becomes arbitrarily small as the factor grows.

Figure 4:

| | |
|------|------|
| 0, 0 | 4, 0 |
| 0, 4 | 3, 3 |

7.3 Growth Rate

Having seen that the elasticity of the repeated prisoner's dilemma grows at least exponentially (as well as that of other games), it will be of interest to show that the order of growth is exponential, i.e., to show an exponential upper bound.

Moreover, perhaps there is room for "classifications" of the elasticity growth rate of repeated games (or maybe some other sorts of parametrized games), with classes such as subexponential, superexponential, etc.

References

- [1] Aumann, R.J. (1987) Correlated equilibrium as an expression of Bayesian rationality, *Econometrica* **55**, 1–18.
- [2] Aumann, R.J. (1992) Irrationality in game theory, in *Economic Analysis of Markets and Games (Essays in Honor of Frank Hahn)*, edited by P.

Dasgupta, D. Gale, O. Hart, and E. Maskin, Cambridge and London: MIT Press, 214–227.

- [3] Basu, K. (1994) The Traveler’s Dilemma: Paradoxes of rationality in game theory, *American Economic Review* **84**, 391–395.
- [4] Neyman, A. (1999) Cooperation in repeated games when the number of stages is not commonly known, *Econometrica* **67**, 45–64.