Elasticity in Extensive Form Games

Gilad Bavly

Preliminary Draft

Abstract

Some simple examples show that the elasticity of "perfect" outcomes of extensive-form games is unbounded in a strong sense, as the smallest perturbations can transform an outcome that was far removed from the set of perfect outcomes into a perfect one. This also highlights the role of differential information in the concept of elasticity, as we do get a bound for the case of perturbations with only symmetric information.

The detailed structure of extensive-form games brings about various equilibrium refinement notions, many of them related to subgame-perfectness. It is natural, at least in some well-behaved classes of extensive games, to consider the elasticity (Bavly(2015)) of some such refined solution concept.

we will see, however, simple examples of well-behaved games, where the sensitivity to our perturbations is unbounded in a strong sense. The smallest perturbations can transform an outcome that was far away from the set of perfect outcomes into a perfect outcome.[1]

Of course, this does tell us something about the volatility of these solution concepts. But perhaps some alteration of the definition of elasticity, designed

[1] Following the perturbation, a game may have no proper subgames, so a refinement other than simple subgame-perfectness is needed.

for extensive games, may prove to be more appropriate for these games, or to capture more.

These examples also emphasize the role that differential information has in the concept of elasticity: we do get a bound, when perturbations are restricted to those with symmetric information.

Finally, we do find some special class of extensive games, where the effects of perturbations are bounded. In fact, the bound we get is not only for some refined equilibria concept, but for all equilibria of the perturbed game. Thus, it is in fact a bound on the elasticity in our regular sense, namely elasticity of correlated equilibria.

1 Examples of Unboundedness

Figure 1:

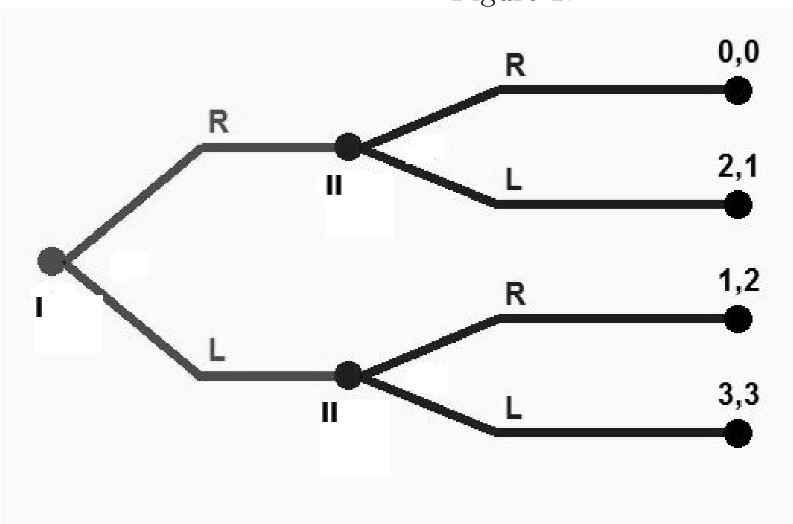

Let Γ be the perfect information game in Figure 1 (any player in Γ gets a different payoff at any two distinct terminal nodes, to make sure there are no pathologies related to ties). Player 1 chooses L or R, and then player 2, informed of 1's choice, chooses L or R. This game has a unique subgame-perfect equilibrium, with payoffs $(3,3)$. Upon this Γ we construct a belief system S with the following properties:

a. The distance of S from Γ can be arbitrarily small.

b. The play even constitutes an extensive-form-perfect equilibrium.

c. The expected payoff from this play is far away from $(3,3)$. Player 1 expects ~ 2, and player 2 expects ~ 1.

d. A finite set of types (in fact $|T| = 2$).

The description of S: player 2 has only one type. Player 1 has two types, α and β. $p(\alpha) = 1 - \varepsilon$, $p(\beta) = \varepsilon > 0$. In the state of nature α, the payoff is the same as g, i.e., $\forall a \in A$, $u(\alpha, a) = g(a)$. In state β, if player 1 chooses L and 2 chooses R, they get $(4,4)$. Otherwise, they get the same as in g.

Type α of player 1 chooses R, and type β chooses L. Player 2 chooses L after the history R, and chooses R after the history L (intuitively, we can explain this play by saying that player 2 believes that if he saw L then it's an indication that player 1 observed the state of world β, and R indicates α). Therefore in state α the play goes (R, L), with payoffs $(2,1)$, and in state β the play goes (L, R), with payoffs $(4,4)$. At every information set of every player, there is a strict incentive not to switch to the other action, and all these information sets are reached with positive probability. It follows that this is an extensive-form perfect equilibrium. This applies to any $\varepsilon > 0$, and therefore S can be arbitrarily close to Γ.

Note that this equilibrium does not contradict Theorem **??**, because it is close to an equilibrium of the original game, albeit a non-subgame-perfect one: the equilibrium where 1 chooses R, and 2 chooses L after R, and R after L.

In this system player 2 deviated from the original backward induction play, because he could deduce something about his own payoff from the previous choice made by player 1. We can, however, construct a similar example with all the above properties, where the payoff of each player depends only on his own type.

In the game of Figure 2, player 1 chooses R and then the game ends, or he chooses L, and then player 2 chooses L and the game ends, or he chooses R and then player 1 chooses last. The backward induction solution is always play L, resulting in the payoff $(3, 1)$.

Figure 2:

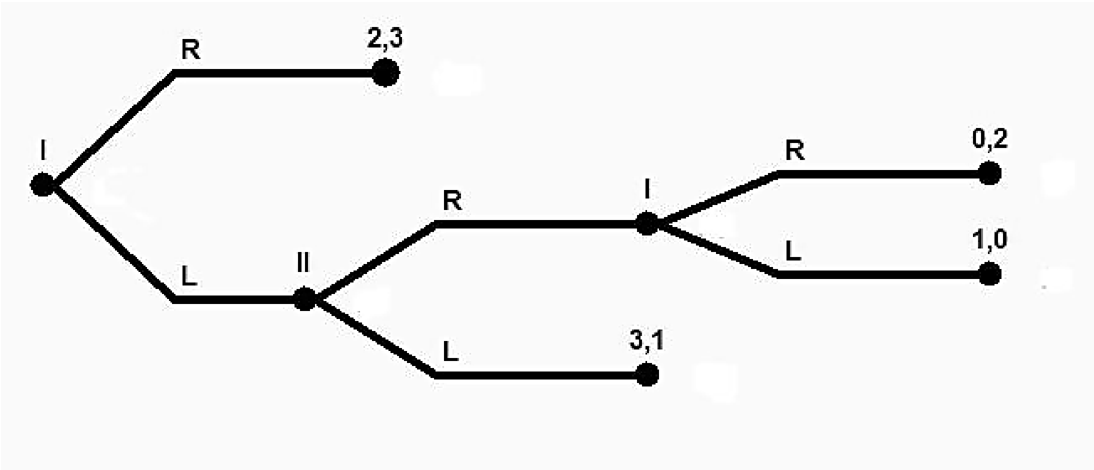

The belief system: player 2 again has only one type, and player 1 has two types, α and β. $p(\alpha) = 1 - \varepsilon$, $p(\beta) = \varepsilon > 0$. In the state of nature α, the payoff is the same as g. In state β, upon the play (L, R, R) they get $(2, 2)$ instead of $(0, 2)$. Otherwise, they get the same as in g. So player 2's payoff is independent of player 1's type.

The play: Type α of player 1 chooses R at his first decision node and L at his second (unreached) node. Type β chooses L at the first and R at the second. Player 2 chooses R. So player 2 acts as if the fact that he gets to choose is an indication that player 1's type is β. With probability $1 - \varepsilon$ the state is α, in which the payoff is $(2, 3)$.

Again, the achieved payoff is indeed close to a (non-subgame-perfect) equilibrium payoff of the original game, where 1 plays R in the first node (and whatever in the second), and 2 plays R.

2 Some Special Cases

Such phenomena as in the above examples cannot, however, occur without differences in the players' information. A belief system (T, p, u) has *symmet-*

ric information if for any type $s_i \in T_i$ of any player i, such that $p(s_i) > 0$, there is just one state of the world $t = (t_1, \ldots, t_n) \in T$ such that $p(t) > 0$ and $t_i = s_i$. Therefore, at every state all players have the same information, as they all know each other's type.

Proposition 2.1. *Let Γ be a generic*[2] *extensive-form game with perfect information. If the belief system S has symmetric information, then the payoff in any subgame-perfect equilibrium payoff of S is within a distance of $2d(S,\Gamma)(1+(n-1)M/m) \leq 2d(S,\Gamma)nM/m$ of the payoff of the unique backward induction solution of Γ, where $M = \max\{g_i(a)-g_i(b) : i \in N; a,b \in A\}$, $m = \min\{|g_i(a) - g_i(b)| : i \in N; a \neq b \in A\}$, and n is the number of players.*[3]

Proof. Let a^* be the unique backward induction play of Γ, and s^* the terminal node reached in this play.

First we note that since information is symmetric, it is meaningful to talk about subgame-perfectness, since every state of the world defines a subgame. Next we note that w.l.o.g. we can prove the result for a game without any types at all. This is because each such subgame is an image of Γ with perhaps altered payoffs, and the overall distance is the average over these subgames.

Now consider a game $\grave{\Gamma}$ with the same structure as Γ, and a payoff function u, and consider a subgame-perfect equilibrium of $\grave{\Gamma}$. Let s be a terminal node reached with positive probability in this equilibrium.

Denote $\delta = d(\Gamma, \grave{\Gamma})$. If $s = s^*$ then $\|u(s^*) - g(s)\| = \|u(s) - g(s)\| \leq \delta$ and it is certainly $\leq 2\delta(1 + (n-1)M/m)$.

If $s \neq s^*$, then let i be the player that last deviated from playing according to a_i^*. Then [i] $g_i(s^*) > g_i(s)$, and therefore $g_i(s^*) \geq g_i(s) + m$. Equilibrium

[2] A statement holds in a *generic* game if, for a given game structure, the statement holds in an open and dense set within the set of payoffs for that structure. Here we can take that set to be all the payoffs in which any player gets a different payoff at any two distinct terminal nodes.

[3] We could have also defined m to be generally larger, by taking the maximum, over all subtrees, of the minimal difference between any two payoffs of any player who has an actual choice within that subtree. The result would still hold.

in $\grave{\Gamma}$ implies that [ii] $u_i(s) \geq u_i(s^*)$, and therefore it must be that either $|u_i(s^*) - g_i(s^*)| \geq m/2$, or $|u_i(s) - g_i(s)| \geq m/2$. This implies that $\delta \geq m/2$.

[i] and [ii] also imply that $|u_i(s) - g_i(s^*)| \leq \delta$, since $|u_i(s^*) - g_i(s^*)|, |u_i(s) - g_i(s)| \leq \delta$. Now, $\|u(s) - g(s^*)\| = \sum_{j=1}^{n} |u_j(s) - g_j(s^*)| \leq \delta + \sum_{j \neq i} |u_j(s) - g_j(s^*)| \leq \delta + \sum_{j \neq i} |u_j(s) - g_j(s)| + \sum_{j \neq i} |g_j(s) - g_j(s^*)| \leq \delta + \delta + (n-1)M = 2\delta + (n-1)M$, and since $\delta \geq m/2$, it is $\leq 2\delta + (2\delta/m)(n-1)M = 2\delta(1 + (n-1)M/m)$.

Since the required inequality holds for every reached terminal node, it also holds for an average thereof. □

The following proposition identifies a special class of extensive-form games, characterized by a dominance property, in which the distance is similarly bounded, but without any restriction on the information structure. Moreover, the result holds not just for subgame-perfect equilibria, but for any equilibria, and therefore can be stated in elasticity terms.

We will say that Γ has the *extensive dominance* property, if for every player i there exists a strategy d_i (namely a map from i's decision nodes to his choices at those nodes), such that given that i plays according to d_i, then at any of his nodes s, $d_i(s)$ is strictly better for i than any other choice, no matter what the other players do afterwards.

Proposition 2.2. *If Γ is a game with extensive dominance, then $\eta(\Gamma) \leq 2n^2M/m$ (with M, m, and n defined as in the previous proposition).*

Proof. Let s^* be the terminal node reached by the strategies $d = (d_1, \ldots, d_n)$. s^* is the unique (correlated) equilibrium outcome of Γ. Let $S = (T, p, u)$ be some rational belief system for Γ (with strategies a). For $i \in N$, let $F_i \subseteq T$ be all states of the world in which i will actually get to play differently than d_i at some point. Choose i that maximizes $p(F_i)$; then $p(F_i) \geq p(F)/n$, where $F = \cup_{j=1}^{n} F_j$.

For any $t \in F_i$ let $\phi(t) \subseteq F_i$ be the information that i has at the point where he makes his first deviation. It is contained in F_i because i knows that he is going to deviate. Also, $t \in \phi(t)$, and therefore the different values of $\phi(t)$ induce a partition of F_i. Let ϕ_0 be some possible value. Since i is rationally deviating at that point, his expectation of the difference between u_i and g_i must be enough to compensate for a loss of at least m (had g_i been his payoff). Therefore $E\left(d(u,g) \mid \phi_0\right) \geq m/2$. As this holds for every possible value of $\phi(t)$, it also holds on the union, namely $E\left(d(u,g) \mid F_i\right) \geq m/2$. And therefore $E\left(d(u,g) \mid F\right) \geq m/2n$.

For any $t \in T \setminus F$, $\|g(a(t)) - g(s^*)\| = 0$, because $a(t) = s^*$. For $t \in F$, $\|g(a(t)) - g(s^*)\| \leq nM$. Therefore $E\left(\|g(a(t)) - g(s^*)\|\right) \leq nMp(F)$. On the other hand, $\delta = E\left(d(u,g)\right) \geq p(F)m/2n$; so $E\left(\|g(a(t)) - g(s^*)\|\right) \leq \delta 2n^2M/m$.

As $E\left(\|u(a(t)) - g(s^*)\|\right) \leq E\left(\|u(a(t)) - u(s^*)\|\right) + E\left(\|u(s^*) - g(s^*)\|\right) \leq \delta 2n^2M/m + \delta$, we get that $\eta(\Gamma) \leq 2n^2M/m$. $\square$

References

[1] Bavly, G. (2015) Elasticity of games, mimeo.