Int J Game Theory (2017) 46:591–594
DOI 10.1007/s00182-016-0537-7

ERRATUM

Erratum to: A nonstandard characterization of sequential equilibrium, perfect equilibrium, and proper equilibrium

Joseph Y. Halpern[1]

Published online: 1 June 2016

Erratum to: Int J Game Theory (2009) 38:37–49 DOI 10.1007/s00182-008-0139-0

As pointed out by an anonymous referee, Theorem 1.4 in the original article, which gives a characterization of perfect equilibrium in terms of nonstandard probability, is incorrect as stated. Fortunately, the theorem is easily corrected. There are actually two distinct problems, which I outline here.

The first problem involves the definition of best response. Both sequential equilibrium and perfect equilibrium consider best responses at all information sets, even ones off the equilibrium path. The subtlety arises in making clear what it means to make a best response at an information set off the equilibrium path. In the standard definition of what it means for a behavioral strategy profile $\vec{\sigma}$ to be a sequential or perfect equilibrium (see, e.g., Osborne and Rubinstein 1994), a sequence of *completely-mixed behavioral strategy profiles* that converges to $\vec{\sigma}$ is considered (where a behavioral strategy σ_i for player i is *completely mixed* if, at each information set I for player i, it assigns positive probability to each possible action at I). In the original article, this sequence of behavioral strategy profiles converging to $\vec{\sigma}$ is replaced by a single completely-mixed *nonstandard* strategy profile $\vec{\sigma}'$ (i.e., a strategy profile where each individual strategy can assign a nonstandard probability to each action) that differs infinitesimally from $\vec{\sigma}$. (The notion of being completely mixed remains the same: each action at an information set gets positive probability, but now that probability can be a nonstandard real, in particular, an infinitesimal). The idea is that, by using a

The online version of the original article can be found under doi:10.1007/s00182-008-0139-0.

✉ Joseph Y. Halpern
halpern@cs.cornell.edu

1 Computer Science Department, Cornell University, Ithaca, NY, USA

completely-mixed strategy $\vec{\sigma}'$, the notion of σ_i' being a best response to $\vec{\sigma}'_{-i}$ becomes completely unambiguous.

The characterization of sequential equilibrium involves σ_i' being an ϵ-best response to $\vec{\sigma}'_{-i}$ at all information sets for player i, where ϵ is an infinitesimal. The characterization of perfect equilibrium requires something stronger: that σ_i be a best response to $\vec{\sigma}'_{-i}$ at all information sets for player i. The problem is that $(\sigma_i, \vec{\sigma}'_{-i})$ is not a completely mixed strategy, and may not reach an information set I at all. This leads to problems with the definition of best response (Definition 2.1) in the original article.

To solve this problem, we should go back to the original motivation for these definitions. What we really want to say is that, after reaching I using $\vec{\sigma}'$, σ_i is the best continuation for player i from then on. To make this precise, we need to know i's beliefs regarding the relatively likelihood of histories in I; this belief is determined by the completely-mixed strategy profile used to reach I. Recall the notation from the original article; given a behavioral-strategy profile $\vec{\sigma}$, $\Pr_{\vec{\sigma}}$ is the distribution on terminal histories induced by $\vec{\sigma}$. Since we can identify a partial history with the terminal histories that extend it, $\Pr_{\vec{\sigma}}(h)$ and $\Pr_{\vec{\sigma}}(I)$ are well defined for a partial history h and an information set I. As usual, we take a *belief system* μ to be a function that associates with each information set I a probability denoted μ_I on the histories in I. Given a behavioral strategy $\vec{\sigma}$ and a belief system μ in an extensive-form game Γ, let

$$\mathrm{EU}_i((\vec{\sigma}, \mu)|I) = \sum_{h \in I} \sum_{z \in Z} \mu_I(h) \Pr_{\vec{\sigma}}(z|h) u_i(z),$$

where Z is the set of terminal histories in Γ. Here $\mu_I(h)$ is the probability that history $h \in I$ is where in I player i actually is; $\Pr_{\vec{\sigma}}(z|h)$ is the probability that z will be the history played given that play proceeds according to $\vec{\sigma}$ from h on; and $u_i(z)$ is the utility that i gains from z being played. Finally, if $\vec{\sigma}$ is a completely-mixed behavioral strategy profile, let $\mu^{\vec{\sigma}}$ be the belief system determined by $\vec{\sigma}$ in the obvious way:

$$\mu_I^{\vec{\sigma}}(h) = \Pr_{\vec{\sigma}}(h|I).$$

We can now define what it means for a player's strategy to be a best response at an information set.

Definition 2.1$'$ If $\epsilon \geq 0$ and I is an information set for player i that is reached with positive probability by $\vec{\sigma}'$, then σ_i is an *ϵ-best response to $\vec{\sigma}'_{-i}$ for i conditional on having reached I using $\vec{\sigma}'$* if, for every strategy τ for player i, we have $\mathrm{EU}_i(((\sigma_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'})|I) \geq \mathrm{EU}_i(((\tau_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'})|I) - \epsilon$.

By way of contrast, here is Definition 2.1 in the original article:

Definition 2.1 If $\epsilon \geq 0$ and I is an information set for player i, σ_i is an *ϵ-best response to $\vec{\sigma}'_{-i}$ for i conditional on having reached I* if, for every strategy τ for player i that agrees with σ_i except possibly at I and information sets preceded by I, we have $\mathrm{EU}_i(\sigma_i, \vec{\sigma}'_{-i}) \geq \mathrm{EU}_i(\tau, \vec{\sigma}'_{-i}) - \epsilon$.

Note the differences between Definitions 2.1 and 2.1$'$. Definition 2.1$'$ enforces the assumption that I is reached using $\vec{\sigma}'$ by using the belief system $\mu^{\vec{\sigma}'}$. In Definition 2.1,

σ_i' plays no role; rather, σ_i and τ have to agree up to I, so the implicit assumption is that I is reached using $(\sigma_i, \vec{\sigma}'_{-i})$, not $\vec{\sigma}'$. It is this difference that causes problems (see the discussion of the statement of Theorem 4.1 below).[1]

The second major problem with Theorem 1.4 involves the definition of perfect equilibrium. According to Definition 2.2 in the original article, "$\vec{\sigma}$ is a *perfect equilibrium* in an extensive-form game Γ iff there exists a sequence $\vec{\sigma}^n$ of completely mixed behavior strategies such that $\vec{\sigma}^n \to \vec{\sigma}$ and, for all n and each information set I of player i, σ_i is a best response to $\vec{\sigma}^n_{-i}$ conditional on having reached I." Even ignoring the issue of best response discussed above, this is not the standard definition of perfect equilibrium. The standard definition involves games in *agent-normal form*, with a different agent assigned to each information set. When considering a best response for agent i at an information set I, agent-normal form allows for the possibility that i will tremble at information sets below I; Definition 2.2 in the original article does not allow this. In Footnote 2 of the original article, it is claimed that Definition 2.2 is equivalent to the standard definition. This claim is false. In fact, Definition 2.2 in the original article characterizes the notion of *quasi-perfect equilibrium* (van Damme 1984), which is much like perfect equilibrium, except that trembles by agent i are not allowed when considering i's best response. As van Damme (1984) shows, quasi-perfect equilibrium and perfect equilibrium can differ in quite simple games.

To prove Theorem 1.4, it is most convenient to use a characterization of perfect equilibrium due to van Damme (1984, Lemma 1). Given an information set I for player i, let A_I be the set of actions for i at I. As usual, we take $\Delta(A_I)$ to be the set of probability measures on A_I. Note that if σ_i is a behavioral strategy for player i then, by definition, $\sigma_i(I) \in \Delta(A_I)$.

Definition 2.1$''$ If $\epsilon \geq 0$ and I is an information set for player i that is reached with positive probability by $\vec{\sigma}'$, then $b \in \Delta(A_I)$ is a *local ϵ-best response to $\vec{\sigma}'_{-i}$ for i conditional on having reached I using $\vec{\sigma}'$* if, for all $b' \in \Delta(A_I)$, we have that

$$\mathrm{EU}_i(((\sigma_i'[I/b], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'})|I) \geq \mathrm{EU}_i(((\sigma_i'[I/b'], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'})|I) - \epsilon,$$

where $\sigma_i'[I/b']$ is the behavioral strategy that agrees with σ_i' except possibly at information set I, and $\sigma_i'[I/b'](I) = b'$.

Note that Definition 2.1$''$ allows changes only at I; by way of contrast, Definition 2.1$'$ allows changes at I and all information sets below I.

Perfect equilibrium can be characterized as follows:

Definition 2.2$'$ $\vec{\sigma}$ is a *perfect equilibrium* in an extensive-form game Γ iff there exists a sequence $\vec{\sigma}^n$ of completely mixed behavior strategies such that $\vec{\sigma}^n \to \vec{\sigma}$ and, for all n and each information set I of player i, $\sigma_i(I)$ is a local best response to $\vec{\sigma}^n_{-i}$ conditional on having reached I using $\vec{\sigma}^n$.[2]

[1] There is actually a second problem in Definition 2.1: it uses the ex ante probability, rather than the probability conditional on reaching I. This problem is also corrected in Definition 2.1$'$. However, it is the former problem that is the deeper conceptual problem.

[2] Of course, a (*local*) *best response* is a (local) 0-best response.

To understand the impact of these changes, consider the following theorem from the original article, which attempts to characterize perfect equilibrium:

Theorem 1.4 *If Γ is an extensive-form game with perfect recall, then $\vec{\sigma}$ is a perfect equilibrium in Γ iff there exists a nonstandard completely-mixed strategy profile $\vec{\sigma}'$ that differs infinitesimally from $\vec{\sigma}$ such that, for each player i and each information set I for player i, σ_i is a best response to $\vec{\sigma}'_{-i}$, conditional on having reached I.*

When attempting to evaluate whether σ_i is a best response to $\vec{\sigma}'_{-i}$ conditional on having reached I according to Definition 2.1, we must implicitly assume that I is reached using $(\sigma_i, \vec{\sigma}'_{-i})$. But σ_i is not completely mixed and, indeed, might not reach I at all. This causes problems in stating the theorem. If we add the words "using $\vec{\sigma}'$" at the end of the theorem (i.e., I is reached using $\vec{\sigma}'$), where "best response to $\vec{\sigma}'_{-i}$, conditional on having reached I using $\vec{\sigma}'$" is defined by Definition 2.1$'$, then we get a characterization of quasi-perfect equilibrium. The proof of Theorem 1.4 sketched in the original article actually proves this result, since the definition of quasi-perfect equilibrium is used rather than that of perfect equilibrium. We can get a characterization of perfect equilibrium by further replacing "σ_i is a best response to $\vec{\sigma}'_{-i}$" by "$\sigma_i(I)$ is a local best response to $\vec{\sigma}'_{-i}$" (i.e., using Definition 2.1$''$ rather than Definition 2.1$'$). The proof is just a slight variant of that given in the paper, and is left to the reader.

A similar change must be made in Definition 2.2 to get an appropriate definition of perfect equilibrium; this is exactly what was done in Definition 2.2$'$. The key point is that, again, the phrase "using $\vec{\sigma}^n$" needs to be added to Definition 2.2, the definition given in the original article.

References

Osborne MJ, Rubinstein A (1994) A course in game theory. MIT Press, Cambridge

van Damme E (1984) A relationship between perfect equilibria in extensive form games and proper equilibria in normal form games. Int J Game Theory 13:1–13