BEST-RESPONSE EQUILIBRIUM: AN EQUILIBRIUM IN FINITELY ADDITIVE MIXED STRATEGIES

Igal Milchtaich, Department of Economics, Bar-Ilan University*

July 2020

A generalization of mixed strategy equilibrium is proposed, where mixed strategies need only be finitely additive and payoff functions need not be integrable or bounded. It is based on an extension of the idea that an equilibrium strategy is supported in the player's set of best-response actions, but is applicable also when no best-response actions exist. The proposed notion of *best-response equilibrium* yields simple, natural mixed equilibria in a number of well-known games where other kinds of mixed equilibrium are complicated or not compelling or they do not exist.

1 Introduction

The simplest interpretation of mixed strategy, which is also the original one (von Neumann and Morgenstern 1953), is that it reflects a player's deliberate assignment of probabilities to his possible actions, or pure strategies. Randomization protects the player from his action being found out by an opponent, since the player does not know it himself. Finding out the *probabilities* would not help the other players if the profile of mixed strategies is an equilibrium, as the latter is defined by the condition that each mixed strategy is a best response in the sense that no unilateral deviation to an alternative mixed strategy can increase the deviating player's expected payoff. Checking whether this condition holds requires examining only pure strategies, because a mixed strategy is best response if and only if it is supported in the set of best-response actions. This fact means that from the player's point of view, the probabilities assigned to the actions in the support are unimportant, which suggests an alternative interpretation of mixed strategy as a commonly held external belief about the player's choice of action rather than a deliberate choice of randomized strategy by the player. In addition, since the best-response condition can be stated in terms of actions, alternative mixed strategies play no essential role, which suggests that it may be unnecessary to even consider them.

This paper presents a notion of mixed strategy equilibrium that makes no reference to alternative mixed strategies. For each player i, only one mixed strategy, the equilibrium strategy, is considered. This strategy σ_i is a finitely additive set function defined on some algebra $\mathcal{A}_i$ of subsets of the player's action set S_i. The algebra is not a priory given but is part of the strategy's specification.[1] Importantly, it is not required to be a *sigma*-algebra and

* igal.milchtaich@biu.ac.il http://faculty.biu.ac.il/~milchti/

[1] This contrasts with the usual definition of mixed strategy, where the domain is some pre-specified measurable structure on S_i like the collection of all Borel sets. A conceptual problem with the latter approach is that, unless S_i is finite, the choice of measurable structure is arguably arbitrary, as it is not indicated by the game itself. Yet choosing it is necessary for defining the *mixed extension* of the game, where players use mixed strategies rather than actions. In the model presented here, there is no mixed extension.

σ_i is not required to be sigma-additive. (For a short review of these and related terms, see Section 2.) This aligns with the interpretation of mixed strategy as a (possibly, incomplete) probabilistic description of the player's choice of action rather than a recipe for actually choosing that action at random. The essential element in the definition of mixed-strategy equilibrium, which is that it excludes the choice of actions yielding low payoff, is retained. However, this idea requires a somewhat more elaborate formulation than with sigma-additive strategies. The formulation constitutes the core of the formal definition of *best-response equilibrium* in Section 3.

As Theorem 1 in that section shows, every single-player game has a best-response equilibrium. The same is not true for "real" games, as Examples 1 (for $n = 3$) and 2 (for $n = 2$) show. Yet, as Section 4 demonstrates, the concept has a number of interesting applications also in multiplayer games. In particular, it may be used for describing an optimal choice of some figure (a price, say) that is very close to a specific value (zero, say), but not quite equal to it.

Section 5 considers two-player zero-sum games. Such games may not have a value in the usual sense yet admit a best-response equilibrium.

In Section 6, best-response equilibrium is compared with other solution concepts that also employ finitely additive probabilities, in particular, optimistic equilibrium (Vasquez 2017) and justifiable equilibrium (Flesch et al. 2018). These solution concepts are not compatible with the principles underlying best-response equilibrium, as described above, and may produce different equilibrium predictions. Specifically, Theorem 2 shows that, with bounded payoff functions, every best-response equilibrium is a justifiable equilibrium but not conversely. Thus, the former is essentially the stronger, more demanding solution concept.

2 Preliminaries

An *algebra*, or field, $\mathcal{A}$ on a set S is any collection of subsets of S that includes the empty set and, for all $A, B \in \mathcal{A}$, also includes the complement A^{C} and the union $A \cup B$. If moreover the union $\bigcup_{k=1}^{\infty} A_k$ is in $\mathcal{A}$ for for every sequence $A_1, A_2, \dots \in \mathcal{A}$, then $\mathcal{A}$ is a *sigma-algebra*. A real-valued (set) function μ defined on an algebra $\mathcal{A}$ is *finitely additive* if $\mu(A) + \mu(B) = \mu(A \cup B)$ for all disjoint $A, B \in \mathcal{A}$, and *sigma-additive* if $\sum_{k=1}^{\infty} \mu(A_k) = \mu(\bigcup_{k=1}^{\infty} A_k)$ for all disjoint $A_1, A_2, \dots \in \mathcal{A}$ with $\bigcup_{k=1}^{\infty} A_k \in \mathcal{A}$. If in addition μ only takes values in $[0,1]$ and $\mu(S) = 1$, then it is called a *finitely additive probability* or a *probability* (measure), respectively. The elements of $\mathcal{A}$ are referred to in this context as the *measurable sets*. A finitely additive probability μ' is an *extension* of another one μ if the corresponding algebras satisfy $\mathcal{A} \subseteq \mathcal{A}'$ and $\mu = \mu'|_{\mathcal{A}}$, and it is a *total* extension if $\mathcal{A}' = \mathcal{P}(S)$, the power set of S. The Carathéodory extension theorem states that every (sigma-additive) probability defined on an algebra $\mathcal{A}$ has a unique extension to a probability defined on the smallest sigma-algebra containing $\mathcal{A}$.

The *outer measure* of a finitely additive probability μ is the function $\mu^*: \mathcal{P}(S) \longrightarrow [0,1]$ defined by

$$\mu^*(C) = \inf \{\mu(A) \mid A \supseteq C, A \in \mathcal{A}\}.$$

A set C with $\mu^*(C) = 0$ is said to be μ-*null*. A property of elements of S is said to hold μ–*almost surely* if it holds outside some μ-null set. If all μ-null sets are measurable (therefore, $\mu^*(C) = 0 \Leftrightarrow \mu(C) = 0$), then μ is said to be *complete*.

For a finitely additive probability μ defined on an algebra $\mathcal{A}$ of subsets of a set S, a *simple measurable* function is any function $f: S \longrightarrow \mathbb{R}$ that takes only finitely many values and satisfies $f^{-1}(\{x\}) \in \mathcal{A}$ for all $x \in \mathbb{R}$. The *integral* of f is defined by

$$\int_S f(s)\, d\mu(s) = \sum_{x\in\mathbb{R}} x\, \mu(f^{-1}(\{x\}))\,.$$

More generally, a function f is μ-*integrable* if there is a sequence $(f_n)_{n\in\mathbb{N}}$ of simple measurable functions such that

$$\lim_{n\to\infty} \mu^*(\{s \in S \mid |f(s) - f_n(s)| > \epsilon\}) = 0$$

for every $\epsilon > 0$ (meaning that $f_n \to f$ in μ-probability) and

$$\lim_{m,n\to\infty} \int_S |f_m(s) - f_n(s)|\, d\mu(s) = 0.$$

(If f is bounded, the second condition is redundant, as it is implied by the first one.) The integral of f is then (well) defined by

$$\int_S f(s)\, d\mu(s) = \lim_{n\to\infty} \int_S f_n(s)\, d\mu(s)\,,$$

and the limit is necessarily finite (Dunford and Schwartz 1958). An alternative way of stating that a function is μ-integrable is saying that its integral with respect to μ exists. It is easy to see that, in this case, the integral of f with respect to any extension of μ also exists and the two integrals are equal.

For a bounded function f, the *upper integral* with respect to μ is defined by

$$\overline{\int_S} f(s)\, d\mu(s) \coloneqq \inf\left\{ \int_S g(s)\, d\mu(s) \;\middle|\; g \text{ a simple measurable function}, g \geq f \right\}$$

and the *lower integral* by

$$\underline{\int_S} f(s)\, d\mu(s) \coloneqq \sup\left\{ \int_S g(s)\, d\mu(s) \;\middle|\; g \text{ a simple measurable function}, g \leq f \right\}.$$

The former is always greater than or equal to the latter, and equality holds if and only if f is μ-integrable, in which case the common value is the integral of f.

In the linear space of all bounded functions $f: S \longrightarrow \mathbb{R}$, the subset of μ-integrable functions is easily seen to be a subspace. The integral is a linear functional on this subspace and satisfies $\left|\int_S f(s)\, d\mu(s)\right| \leq \sup_{s\in S}|f(s)|$. By the Hahn–Banach theorem, there is a (generally, non-unique) extension of this linear functional to a linear functional ψ that is defined on the whole space and satisfies a similar inequality, $|\psi(f)| \leq \sup|f|$. It may be viewed as an extension of integration with respect to μ; for any bounded function f, $\psi(f)$ is the integral of f. In particular, the function $\mu^\psi: \mathcal{P}(S) \longrightarrow [0,1]$ defined by $\mu^\psi(A) = \psi(1_A)$ is an extension of μ and, by the linearity of ψ and the above inequality, is also a finitely additive probability. Thus, μ^ψ is a *total* extension of μ. This proves that every finitely additive probability has a total extension. Note that, with respect to a total finitely additive probability, every bounded function is integrable because it is the uniform limit of a sequence of simple measurable functions (as all sets are measurable.) The linear functional ψ considered above is actually integration with respect to μ^ψ.

For an integer $n \geq 2$ and a finitely additive probability μ_i on an algebra $\mathcal{A}_i$ of subsets of a set S_i for each $1 \leq i \leq n$, the *product* $\mu = \prod_i \mu_i$ is a finitely additive probability defined on the *product algebra* $\mathcal{A} = \prod_i \mathcal{A}_i$, whose elements are all sets in the Cartesian product $S = \prod_i S_i$ that are finite unions of measurable rectangles, that is, sets $A \subseteq S$ of the form $A = \prod_i A_i$ with $A_i \in \mathcal{A}_i$ for each i. For such a measurable rectangle, the product probability is given by $\mu(A) = \prod_i \mu_i(A_i)$. Note that the individual μ_i's and $\mathcal{A}_i$'s can be recovered from the product μ and its domain $\mathcal{A}$: the former coincide with the marginals of μ and the latter satisfy $\mathcal{A}_i = \{ A_i \subseteq S_i \mid S_1 \times \cdots \times A_i \times \cdots \times S_n \in \mathcal{A} \}$).

Lemma 1 For a bounded function $f: S \longrightarrow \mathbb{R}$,

$$\overline{\int_S} f(s)\, d\mu(s) \geq \overline{\int_{S_n}} \cdots \overline{\int_{S_2}} \overline{\int_{S_1}} f(s_1, s_2, \dots, s_n)\, d\mu_1(s_1)\, d\mu_2(s_2) \cdots d\mu_n(s_n)$$
$$\geq \underline{\int_{S_n}} \cdots \underline{\int_{S_2}} \underline{\int_{S_1}} f(s_1, s_2, \dots, s_n)\, d\mu_1(s_1)\, d\mu_2(s_2) \cdots d\mu_n(s_n) \geq \underline{\int_S} f(s)\, d\mu(s).$$

Proof. The middle inequality is based on iterated use of the inequality between the upper and lower integrals of bounded functions. The first and last inequalities are analogous. To prove the former, observe first that a similar inequality holds (as equality between integrals) with f replaced by the indicator function of any measurable rectangle, hence also with f replaced by any simple measurable function g. For $g \geq f$, the last conclusion trivially implies that a similar inequality holds with g replacing f only on the left-hand side. Taking the infimum over all $g \geq f$ completes the proof. ∎

An immediate corollary of Lemma 1 is that, if the iterated integral

$$\int_{S_n} \cdots \int_{S_2} \int_{S_1} f(s_1, s_2, \dots, s_n)\, d\mu_1(s_1)\, d\mu_2(s_2) \cdots d\mu_n(s_n)$$

(by implication, also each of the inner integrals) exists, then it lies between the upper and lower integrals of f, and is therefore equal to the ("multiple") integral $\int_S f(s)\, d\mu(s)$ if the latter also exists. Note that the existence of the iterated integral of a bounded function may depend on the order of integration, and it neither implies nor is implied by the existence of the multiple integral. Thus, Fubini's theorem does not hold here. However, if the iterated integral and the multiple integral both exist, then by Lemma 1 they must be equal, and so the *value* of the former cannot depend on the order of integration.

3 Mixed strategies and best-response equilibrium

In an n-player game ($n \geq 1$), each player i has a set S_i of actions, or pure strategies, and a payoff function $u_i: S \longrightarrow \mathbb{R}$, where $S = \prod_j S_j$ is the set of all action profiles. (It is sometimes convenient to view the function u_i as defined on the product set $S_i \times S_{-i}$, where $S_{-i} = \prod_{j \neq i} S_j$.) A (*mixed*) *strategy* for player i is any finitely additive probability σ_i defined on an algebra $\mathcal{A}_i$ of subsets of S_i. A special case is any pure strategy: an action s_i is identifiable with the (total) probability δ_{s_i}, the Dirac measure at s_i. A (mixed-) strategy profile $(\sigma_1, \sigma_2, \dots, \sigma_n)$, which specifies a strategy σ_i for each player i, may be identified with the product $\sigma = \prod_i \sigma_i$ (see the comment immediately preceding Lemma 1), and it may also be written as (σ_i, σ_{-i}), where i is any player and $\sigma_{-i} = \prod_{j \neq i} \sigma_j$.

Mixed strategies can be mixed. That is, for any $L \geq 2$ and nonnegative weights $\lambda_1, \lambda_2, \dots, \lambda_L$ that sum up to 1, if $\sigma_i^1, \sigma_i^2, \dots, \sigma_i^L$ are strategies for player i, defined on algebras $\mathcal{A}_i^1, \mathcal{A}_i^2, \dots, \mathcal{A}_i^L$, then the weighted average $\sum_{l=1}^{L} \lambda_l \sigma_i^l$ is also a strategy, defined on the algebra $\bigcap_l \mathcal{A}_l$.

Definition 1 A strategy profile σ is a *best-response equilibrium* if, for every player i, (i) the integral

$$v_i(s_i) \coloneqq \int_{S_{-i}} u_i(s_i, s_{-i})\, d\sigma_{-i}(s_{-i}) \tag{1}$$

exists for every $s_i \in S_i$, and (ii) the function $v_i: S_i \longrightarrow \mathbb{R}$ thus defined satisfies

$$\sigma_i^*(\{\, s_i \in S_i \mid v_i(s_i) < a \,\}) = 0 \tag{2}$$

for every $a < \sup_{s_i \in S_i} v_i(s_i)$.

Requirement (i) in the definition concerns only the other players' strategies. These need to be such that every action s_i yields player i a well-defined expected payoff $v_i(s_i)$. Requirement (ii) may be interpreted as the condition that strategy σ_i is a *best response* for player i to the other players' strategies. It says that every number smaller than $\sup v_i$ is a σ_i-essential lower bound of the function v_i, put differently, that the supremum of v_i (which may be finite or ∞) coincides with the *σ_i-essential infimum*. If $\sup v_i < \infty$, the requirement can also be stated as the condition that $v_i - \sup v_i$ is a *σ_i-null function*. If σ_i is a probability (thus, sigma-additive), this is equivalent to the condition that the equality $v_i = \sup v_i$ holds σ_i-almost surely. However, if σ_i is only finitely additive, then the equivalence does not hold: the last condition is stronger. Thus, a profile of mixed strategies that are probabilities is a best-response equilibrium if and only if each player's mixed strategy assigns probability 1 to some set of payoff-maximizing actions. But in general, this condition is not necessary but is only a sufficient condition for best-response equilibrium. As the next proposition shows, another familiar equilibrium condition is both necessary and sufficient.

Proposition 1 In Definition 1, if $\sup v_i < \infty$, then the best-response requirement (ii) holds if and only if v_i is σ_i-integrable and

$$\int_{S_i} v_i(s_i)\, d\sigma_i(s_i) = \sup v_i \,.$$

Proof. A nonpositive function is σ_i-null if and only if it is σ_i-integrable and the integral is zero (Dunford and Schwartz 1958, Theorem II.2.20). Apply this to the function $v_i - \sup v_i$. ■

The equilibrium condition identified in Proposition 1 is that the mixed strategy of each player i yields maximal expected payoff. However, this payoff, $\sup v_i$, cannot generally be interpreted as player i's *equilibrium payoff* in the best-response equilibrium σ. That payoff is given by $\int_S u_i(s)\, d\sigma(s)$ – if the integral exists. If in addition the payoff function u_i is bounded, then it follows from Proposition 1 and Lemma 1 that the equilibrium payoff is equal to $\sup v_i$. However, if u_i is not σ-integrable, then the equilibrium payoff is not well defined – it does not exist. The interpretation is that, in this case, the information provided by the strategy profile is not sufficient even for a probabilistic determination of the player's payoff (although it would become so, according to requirement (i) in the definition, if the player's own action were known with certainly).

Even in a finite game, a best-response equilibrium σ does not necessarily assign a probability to every single action. An atom A of $\mathcal{A}_i$ may include several of player i's actions, in particular, equivalent actions. It is, however, always possible to assign probabilities to these actions by arbitrarily dividing the probability $\sigma_i(A)$ among them. Doing so for one or more players i yields an equilibrium that *extends* σ in the sense that its components are extensions of the latter's components. The next proposition generalizes this observation.

Proposition 2 For every best-response equilibrium σ, every strategy profile $\tilde{\sigma}$ that extends σ is also a best-response equilibrium, and there is at least one such $\tilde{\sigma}$ that is total (in the sense that all its components are so).

Proof. As already remarked, if $\tilde{\sigma}$ extends σ, then every function that is σ-integrable is also $\tilde{\sigma}$-integrable and the two integrals are equal. In addition, for every player i, $\tilde{\sigma}_i^*(C) \leq \sigma_i^*(C)$ for every $C \subseteq S_i$. It follows that $\tilde{\sigma}$ is a best-response equilibrium if σ is so. As proved in Section 2, every strategy, hence every strategy profile, has a total extension. ■

A strategy that is total is in particular complete. Therefore, a corollary of Proposition 2 is that there would be essentially no loss of generality in replacing the outer measure σ_i^* in the definition of best-response equilibrium with σ_i itself and requiring the set in Eq. (2) to be measurable.

With respect to a total strategy, every bounded function is integrable (see Section 2). However, this fact does not take the bite out of requirement (i) in Definition 1, because with $n \geq 3$, the requirement refers to integrability with respect to the *product* of strategies. This makes it a substantial, rather than technical, requirement, as the following example demonstrates.

Example 1 *Three-player game without best-response equilibrium.* For three players, the action set is the open interval $(0,1)$. The payoff functions are $u_1(s) = -s_1$, $u_2(s) = -s_2$ and $u_3(s) = \min(s_2/s_1, 1)$ (where $s = (s_1, s_2, s_3)$). For $i = 1,2$, requirement (2) in the definition reads $\int (-s_i)\, d\sigma_i(s_i) = 0$. For player 3, requirement (i) and Lemma 1 imply that, if the iterated integrals $\int \int u_3\, d\sigma_2\, d\sigma_1$ and $\int \int u_3\, d\sigma_1\, d\sigma_2$ exist, they must be equal. However, this condition does not hold. For every s_1, $0 \leq \min(s_2/s_1, 1) \leq s_2/s_1$ for all s_2, which, since $\int 0\, d\sigma_2(s_2) = 0$ and $\int s_2/s_1\, d\sigma_2(s_2) = (-1/s_1) \int (-s_2)\, d\sigma_2(s_2) = 0$, implies that the "sandwiched" integral $\int \min(s_2/s_1, 1)\, d\sigma_2(s_2)$ exists and is also 0 (because the upper and lower integrals both have this value). For every s_2, $1 \geq \min(s_2/s_1, 1) \geq 1 - s_1/s_2$ for all s_1, which similarly implies that $\int \min(s_2/s_1, 1)\, d\sigma_1(s_1)$ exists and is equal to 1. It follows that the two iterated integrals above are 0 and 1, respectively, and so they are not equal. This proves that no strategy profile is an equilibrium.

With $n = 2$, requirement (i) in the definition is not an issue. However, (ii) may well be so.

Example 2 *Two-player game without best-response equilibrium.* For two players, the action set is the set $\mathbb{N}$ of natural numbers. The payoff functions are $u_1(s) = s_1 1_{s_1 \leq s_2}$ and $u_2(s) = s_2 1_{s_1 \leq s_2 \text{ or } s_2 = 1}$, so that the (infinite) payoff matrix is

$$\begin{array}{c} \\ 1 \\ 2 \\ 3 \\ \vdots \\ n \\ \vdots \end{array} \begin{array}{c} \begin{array}{cccccc} 1 & 2 & 3 & \cdots & n & \cdots \end{array} \\ \begin{pmatrix} 1,1 & 1,2 & 1,3 & \cdots & 1,n & \cdots \\ 0,1 & 2,2 & 2,3 & \cdots & 2,n & \cdots \\ 0,1 & 0,0 & 3,3 & \cdots & 3,n & \cdots \\ \vdots & \vdots & & \ddots & \vdots & \\ 0,1 & 0,0 & 0,0 & \cdots & n,n & \\ \vdots & \vdots & \vdots & & & \ddots \end{pmatrix} \end{array}.$$

In a total best-response equilibrium σ, if strategy σ_1 is "diffuse" in sense that $\sigma_1(\{n\}) = 0$ for all n, then $v_2(1) = 1$ and $v_2(n) = 0$ for all $n \geq 2$, which implies that σ_2 must be concentrated at 1, that is, $\sigma_2(\{1\}) = 1$. The conclusion and the best-response requirement imply that σ_1 must also be concentrated at 1, and so it is actually not diffuse. On the other hand, if σ_1 is not diffuse, then $\sum_{n\geq 1} \sigma_1(\{n\}) > 0$ and therefore the sequence $v_2(2), v_2(3), \ldots$ increases to infinity, which implies that σ_2 must be diffuse. The conclusion means that $v_1(n) = n$ for all n, which implies that σ_1 must also be diffuse. These contradictions prove that a total best-response equilibrium does not exist. In view of Proposition 2, the same is true with 'total' omitted.

Existence of best-response equilibrium is guaranteed in the special case $n = 1$.

Theorem 1 Every one-player game has a best-response equilibrium.

Proof. In the player's action set S, let $(s^n)_{n\in\mathbb{N}}$ be some sequence such that $\lim_{n\to\infty} u(s^n) = \sup u$, the supremum (finite or otherwise) of the payoff function. Define a strategy σ by $\sigma(A) = 0$ or $= 1$ if A or its complement, respectively, includes only finitely many points in $(s^n)_{n\in\mathbb{N}}$. If neither condition holds, A is not measurable.[2] By definition of limit, $\sigma(A) = 0$ holds for the set $A = \{ s \in S \mid u(s) < a \}$ for every $a < \sup u$. Thus, σ is a best-response equilibrium. ■

The construction in the proof of Theorem 1 does not use, or assume, any structure on S. However, action sets often do have one or more natural structures – a measurable structure, a topology or an order relation – in which case other, possibly more natural, equilibrium strategies may exit.

Example 3 In a one-player game, the payoff is any real number s the player chooses. While $s = \infty$ is not a legitimate choice, the following strategy δ_∞ may be viewed as coming close: $\delta_\infty(A) = 0$ or $= 1$ if A or A^{C}, respectively, is bounded from above. This strategy is a best-response equilibrium. In an n-player version of the game, player i's request of s_i is granted only if it is higher than all the other players' requests. For a strategy profile σ to be a best-response equilibrium, it suffices that the strategy of some player i is δ_∞. The other strategies do not matter, except that they have to satisfy the technical condition that the ray $(-\infty, x)$ is measureable for all x. This is because, for every player $j \neq i$ and action s_j, the payoff $u_j(s_j, s_{-j})$ is nonzero only if $s_j > s_i$. This implies that $v_j = 0$ identically, so that any strategy is a best response for j. For player i, since the payoff function $u_i(s_i, s_{-i})$ is obviously nondecreasing in s_i, the function v_i is nondecreasing and therefore the set in Eq. (2) is bounded from above for every $a < \sup v_i$, which by definition of δ_∞ means that (2) holds.

The intuitive meaning of the strategy δ_∞ is that it describes the choice of a very large number, indeed, one exceeding any specified x. (This is clearly an impossibility for usual mixed strategies, which are probabilities, as taking $x = 1,2, \ldots$ and using sigma-additivity leads to a contradiction.) Similar constructs, for $x \in \mathbb{R}$, are δ_{x^+}, which is defined by $\delta_{x^+}(A) = 1$ or $= 0$ if A or A^{C}, respectively, includes a right neighborhood of x, and δ_{x^-}, which is defined similarly using left neighborhoods. These strategies, strategy δ_∞, the similarly defined $\delta_{-\infty}$ and the Dirac measures δ_x, $x \in \mathbb{R}$, can all be restricted to a common

[2] However, by Proposition 2, there are extensions of σ that render all sets measurable. Such an extension is the function $A \mapsto \lim_{n\to\infty} \delta_{s^n}(A)$, where $\lim$ refers to some fixed Banach limit (so that it exists for every bounded sequence).

subalgebra, namely, the algebra $\mathcal{I}$ that consists of all finite unions of intervals in the real line (where 'intervals' refers to all convex sets, including $\mathbb{R}$, $\emptyset$, singletons and rays).[3] They are moreover the only finitely additive probabilities defined on $\mathcal{I}$ that take only the values 0 and 1.

4 Applications

Best-response equilibria have a number of notable applications, where they seem quite natural.

Example 4 *Bilateral trade.* A buyer has to offer a price p to the owner of an item whose worth is 1 to the buyer and 0 to the seller. The seller has to decide what prices are acceptable, with the natural proviso that if a price p is acceptable, then so is any higher price. The seller's sensible strategy of accepting any price greater than zero is weakly dominant, yet it is not an equilibrium strategy because no action of the buyer is a best response to it. Offering any $p > 0$ is less profitable than offering, say, half that price. There is moreover no "normal" mixed strategy that is a best response. However, the intuitive idea that the buyer should offer as little as possible, or "an ϵ", is captured by the strategy δ_{0^+}, which together with the seller's (pure) strategy of accepting any positive price constitutes a best-response equilibrium.[4] The traders' payoff functions are integrable with respect to this equilibrium. The integrals, which give the expected profits, are 0 for the seller and 1 for the buyer.

Example 5 *Price competition.* Price competition among identical firms may be expected to drive profits to zero. However, as indicated by Vasquez (2017), considerably higher profits are supported by equilibria involving finitely additive probabilities. This makes these equilibria qualitatively different also from (regular) ϵ-equilibria.

Consider a good that is produced by n identical firms. Each firm i sets a price $p_i \geq 0$. The demand $D(p)$, which is determined by the demand function D and the lowest price $p = \min_i p_i$, is equally divided among the k (≥ 1) firms tied for the lowest price. The profit for firm i is therefore

$$u_i(p_1, p_2, \dots, p_n) = \begin{cases} p_i \dfrac{D(p_i)}{k} - C\left(\dfrac{D(p_i)}{k}\right), & p_i = \min_j p_j \\ 0, & \text{otherwise} \end{cases},$$

where C is the firms' identical cost function. If the monopoly profit function $\pi_M(p) = pD(p) - C(D(p))$ is continuous, unimodal and positive at its maximum point p_M, then $\delta_{p_M^-}$ is the equilibrium strategy in a symmetric best-response equilibrium in which the expected price is the monopoly price p_M. This is because, if a single firm i sets a price $p_i < p_M$, it would be the sole seller, while $p_i \geq p_M$ would mean no sells, and so the expected profit $v_i(p_i)$ is given by

$$v_i(p_i) = \begin{cases} \pi_M(p_i), & 0 \leq p_i < p_M \\ 0, & p_i \geq p_M \end{cases}.$$

[3] More generally, for any $S \subseteq \mathbb{R}$, the collection $\{A \cap S \mid A \in \mathcal{I}\}$ is an algebra on S, which may also be denoted by $\mathcal{I}$ if the meaning is clear from the context.

[4] For an alternative solution to the problem of nonexistence of equilibrium, which employs a set-valued solution concept, see Milchtaich (2019).

The supremum of v_i is the monopoly profit $\pi_M(p_M)$. For every $\epsilon > 0$, the probability that the strategy $\delta_{p_M^-}$ assigns to the set of prices $\{ p_i \mid v_i(p_i) < \pi_M(p_M) - \epsilon \}$ is zero, which shows that it is indeed a best response. Note that the expected equilibrium profit of an individual firm is not well defined, as u_i is not integrable with respect to the equilibrium (but only becomes so after fixing p_i). However, $p = \min_i p_i$ is integrable, and its integral, which is p_M, gives the expected equilibrium price.

There are additional, lower equilibrium prices, and the continuity and unimodality assumptions above are made for illustrative purposes only. In general, a sufficient condition for a price p to be the expected price in a symmetric best-response equilibrium in which the equilibrium strategy is δ_{p^-} is that π_M is nondecreasing in the interval $(0, p)$ and its supremum there is nonnegative. A rather similar result holds for non-identical firms, which differ in their cost functions.

Price competition may have no mixed-strategy equilibrium in the usual sense, that is, with mixed strategies that are (sigma-additive) probabilities (Hoernig 2007, Dastidar 2011). This is so, for example, for $n = 2$, $D(p) = 1 - p$ and quasi-fixed cost, $C(q) = F$ for $q > 0$ and $= 0$ for $q = 0$, with $0 < F < 1/4$. For finitely additive probabilities, by contrast, this case poses no difficulty. By the result in the previous paragraph, $(\delta_{p^-}, \delta_{p^-})$ is a best-response equilibrium for every $1/2 - \sqrt{1/4 - C} \le p \le 1/2$. The upper and lower bounds on the equilibrium price p correspond to the monopoly profit and zero profit, respectively.

Example 6 *Spatial competition with three firms.* With consumers uniformly distributed on the unit interval $[0,1]$, it is well known that this model has no equilibrium in pure strategies (Eaton and Lipsey 1975). It does have a symmetric equilibrium in mixed strategies, where all three firms (independently) choose a location in $[1/4, 3/4]$ according to the uniform distribution on this subinterval (Shaked 1982). There is also a unique (up to permutations of firms) equilibrium with a mixture of pure and mixed strategies, in which one firm chooses $1/2$ and the other two use an identical mixed strategy that specifies a particular continuous distribution on the interval $[5/24, 19/24]$ that is symmetric with respect to $1/2$ and puts most of the weight around $1/4$ and $3/4$ (Osborne and Pitchik 1986). This mixed strategy cannot be replaced by the strategy that simply randomizes fifty-fifty between $1/4$ and $3/4$, as the replacement would make a deviation to $1/2$ profitable for the two randomizing firms. However, it can be replaced with $1/2\, \delta_{1/4^-} + 1/2\, \delta_{3/4^+}$, and more generally by $1/2\, \delta_{x^-} + 1/2\, \delta_{(1-x)^+}$ for any $1/4 \le x \le 1/3$. This is because, if player 2 uses the last strategy and player 3 chooses $1/2$, then the expected profit for player 1 from choosing location $0 \le s_1 \le 1$ is given by $v_1(s_1) = f(\min\{s_1, 1 - s_1\})$, where

$$f(t) = \begin{cases} t/2 + x/4 + 1/8, & 0 \le t < x \\ t/4 - x/4 + 1/4, & x \le t \le 1/2 \end{cases}.$$

If $x \ge 1/4$, then $\sup v_1 = 3/4\, x + 1/8$, and therefore δ_{x^-} is a best response because $\delta_{x^-}(\{ s_1 \mid v_i(s_i) < 3/4\, x + 1/8 - \epsilon \}) \le \delta_{x^-}([x - 2\epsilon, x)^{\mathrm{C}}) = 0$ for every $\epsilon > 0$. Similarly, $\delta_{(1-x)^+}$ is a best response, and therefore also the average of the two is so. The additional requirement $x \le 1/3$ comes from consideration of player 3's alternatives. Thus, with both inequalities holding, the symmetric strategy profile is a best-response equilibrium. With respect to this equilibrium, only player 3's payoff is well defined. That payoff lies in $(1/3, 1/2)$.

5 Zero-sum games

In a finite two-player zero-sum game, an equilibrium in mixed strategies can be found by solving two uncoupled optimization problems, one for each player. The problem is to find for the player an optimal, that is, a maxmin or equivalently minmax, strategy. Thus, a strategy profile $\sigma = (\sigma_1, \sigma_2)$ is an equilibrium if and only if, for each player i, the value of $\sup v_i$ (where v_i is defined by (1)) would not decrease if the strategy σ_j of the other player j were replaced by any other mixed strategy. In this case, $\sup v_1$ is equal to $-\sup v_2$, and it is the value of the game. For best-response equilibrium, characterization in terms of maxmin or minmax is not applicable, because there is no notion of alternative mixed strategies.
A characterization that is applicable is the following one.

Proposition 3 In a two-player zero-sum game with a bounded payoff function u_1, consider a strategy profile σ for which the integral of u_1 with respect to σ and the two corresponding iterated integrals exist. The strategy profile is a best-response equilibrium if and only if $\sup v_1 + \sup v_2 = 0$, and in this case, the players' equilibrium payoffs are given by the respective suprema.

Proof. It follows from Lemma 1 that both iterated integrals must be equal to the multiple integral $\int_S u_1(s)\, d\sigma(s)$. Since $u_2 = -u_1$, this means that

$$\int_{S_1} v_1(s_1)\, d\sigma_1(s_1) + \int_{S_2} v_2(s_2)\, d\sigma_2(s_2) = 0.$$

The first and second integral in this equation are clearly less than or equal to $\sup v_1$ and $\sup v_2$, respectively, and by Proposition 1, both inequalities hold as equalities if and only if σ is a best-response equilibrium. As remarked, in this case, the (well-defined) equilibrium payoff of each player i is $\sup v_i$. ■

Example 7 *Game without a value.* In a two-player zero-sum game, both players' action set is $[0,1]$ and the payoff function is $u_1(s_1, s_2) = g(s_1 - s_2) + 1$, where $g(t) = \operatorname{sign}(t) - \operatorname{sign}(t + 1/2)$. With usual mixed strategies, that is, (sigma-additive) probabilities on the Borel sets, this game does not have an equilibrium or even an ϵ-equilibrium for sufficiently small ϵ, as the maxmin and minmax values are different, $1/3$ and $3/7$ respectively (Sion and Wolfe 1957; see also Dasgupta and Maskin 1986). However, player 2 has a finitely additive mixed strategy that lowers player 1's maximum payoff to $1/3$, namely, $\sigma_2 = 1/3\, \delta_{1/2^-} + 2/3\, \delta_1$ (Vasquez 2017). It follows from Proposition 3 that together with $\sigma_1 = 1/3\, \delta_0 + 2/3\, \delta_1$, for example, against which player 2's maximum payoff is $-1/3$, this strategy constitutes a best-response equilibrium, with equilibrium payoffs $1/3$ and $-1/3$.

The assumption in Proposition 3 that the payoff function is σ-integrable cannot be dropped. Without it, the equality $\sup v_1 + \sup v_2 = 0$ is neither sufficient nor necessary for best-response equilibrium, as the following examples show.

In the game in Example 7, the above equality holds for $\sigma_1 = (3\sqrt{2} - 4)\delta_0 + (3 - 2\sqrt{2})\delta_{1/2^-} + (2 - \sqrt{2})\delta_1$ and $\sigma_2 = (3\sqrt{2} - 4)\delta_{1/2} + (3 - 2\sqrt{2})\delta_{1/2^-} + (2 - \sqrt{2})\delta_1$. Specifically, player 1's strategy makes player 2's maximum payoff equal to $1 - \sqrt{2}$, thus guaranteeing player 1 a minimum of $\sqrt{2} - 1$, and player 2's strategy makes this figure player 1's maximum payoff (Yanovskaya 1970). However, (σ_1, σ_2) is not a best-response equilibrium because, for both players, actions just below $1/2$ (which are picked up by $\delta_{1/2^-}$) yield a significantly lower payoff than the maximum.

In a similar game with $u_1(s_1, s_2) = \text{sign}(s_1 - s_2)\,(-1)^{1_{s_1=1}}(-1)^{1_{s_2=1}}$ (Ville 1938), the strategy profile $(\delta_{1^-}, \delta_{1^-})$ is a best-response equilibrium because $v_1 = v_2 = -1$ identically: all actions yield a player a payoff of -1 if the opponent's strategy is δ_{1^-} (Yanovskaya 1970). But $\sup v_1$ and $\sup v_2$ sum up to -2 rather than zero, which reflects the fact that they are not equilibrium payoffs; the payoff function is not integrable. Note that, with usual mixed strategies, an equilibrium does not exist. For every (sigma-additive) strategy of the opponent, there are for each player actions yielding payoffs arbitrarily close to 1, which means that the infsup and supinf values of u_1 are different: 1 and -1 respectively.

6 Similar solution concepts

The idea of relaxing the sigma-additivity requirement in the definition of mixed strategy to finite additivity is not new (Yanovskaya 1970 credits Karlin 1950 for it). Neither is the realization that integrability with respect to a product algebra, rather than product sigma-algebra, is a strong condition, which is not satisfied by a number of games of interest with non-continuous payoff functions. Non-integrability of a payoff function means that the expected payoff is not well defined, which creates a difficulty for defining, let alone identifying, best response. One solution to this problem is to apply the mixed equilibrium concept only when the payoff functions are integrable (Marinacci 1997). However, such a restriction means that some simple and natural equilibria, or even all equilibria in a game, may be excluded, as demonstrated above. A different approach to dealing with the ambiguity inherent in non-integrality of the payoff functions is to assume that the players' perception of their current payoffs is different from their perception of the payoffs they would receive by deviating to alternative strategies. In particular, a player may be optimistic about the former and pessimistic about the latter. This approach underlies the solution concept of *optimistic equilibrium* proposed by Vasquez (2017). The best-response equilibrium described in Example 5 is viewed by Vasquez as reflecting optimism. All firms are aiming at a price just below the monopoly price p_M, and each of them effectively believes its price will be the lowest. Note, however, that while this equilibrium is similar in spirit to that of the n-person game in Example 3, the latter would have to be interpreted as expressing pessimism. The other players effectively believe they will be "outbid" by player i, even if their strategy is also δ_∞. Rather than reflecting optimism or pessimism, the idea underlying the best-response equilibrium concept is that players evaluate each of their possible actions against the other players' uncertain actions, with the uncertainty specified by the respective mixed strategies. Theirs is therefore a different perspective than that of an outside observer, who is uncertain about everyone's actions. The integral with respect to the product probability represents the latter point of view, and is therefore irrelevant to any of the individual players.

Motivated by the work of Vasquez (2017), Flesch et al. (2018) proposed replacing the integral with the upper integral for the current payoff and with the lower integral for the alternatives. For a game with bounded payoff functions, and for a specified algebra $\mathcal{A}_i$ on the action set S_i of each player i, a strategy profile σ is a *justifiable equilibrium* if for every player i and strategy τ_i (that is also defined on $\mathcal{A}_i$)

$$\overline{\int_S} u_i(s)\, d\sigma(s) \geq \underline{\int_S} u_i(s)\, d(\tau_i, \sigma_{-i})(s). \tag{3}$$

Flesch et al. (2018) prove that a justifiable equilibrium exists for any choice of the players' algebras.[5] They illustrate this concept with an example (Wald's game) that is similar to the following one.

Example 3 (continued) In the two-player case, for $i = 1$ inequality (3) reads

$$\overline{\int_{\mathbb{R}^2}} 1_{s_1>s_2}\, d\sigma(s_1, s_2) \geq \underline{\int_{\mathbb{R}^2}} 1_{s_1>s_2}\, d(\tau_1, \sigma_2)(s_1, s_2). \tag{4}$$

A sufficient condition for this inequality to hold for all τ_1 is that (i) $\sigma_1(A_1) = 0$ for every set $A_1 \in \mathcal{A}_1$ that is bounded from above or (ii) a similar condition holds for σ_2. To see this, consider any simple measurable function $g: \mathbb{R}^2 \longrightarrow \mathbb{R}$, which can be written as $\sum_{l,m} \lambda_{lm} 1_{A_1^l \times A_2^m}$, where $\{A_1^l\}_l \subseteq \mathcal{A}_1$ is a finite partition of $\mathbb{R}$ and similarly for player 2. If $g \geq 1_{s_1>s_2}$, then A_1^l must be bounded from above for all l and m with $\lambda_{lm} < 1$, and so there is some $A_1 \in \mathcal{A}_1$ that is bounded from above such that $g \geq 1_{A_1^\complement \times \mathbb{R}}$. If (i) holds, then the last inequality implies $\int_{\mathbb{R}^2} g\, d\sigma \geq 1$, which proves that the left-hand side of inequality (4) is 1, so that the inequality necessarily holds. Similarly, if $g \leq 1_{s_1>s_2}$, then A_2^m must be bounded from above for all l and m with $\lambda_{lm} > 0$, and so there is some $A_2 \in \mathcal{A}_2$ that is bounded from above such that $g \leq 1_{\mathbb{R}\times A_2}$. If (ii) holds, then the last inequality implies $\int_{\mathbb{R}^2} g\, d(\tau_1, \sigma_2) \leq 0$ for any τ_1, which proves that the right-hand side of inequality (4) is 0, so that the inequality necessarily holds. It follows, by symmetry, that either condition implies that σ is a justifiable equilibrium.

As shown, a sufficient condition for (σ_1, σ_2) to be a best-response equilibrium in the game in Example 3 is that at least one of the two strategies is δ_∞. This is essentially the same sufficient condition established for justifiable equilibrium. However, the two solution concepts are in general different both substantially and conceptually. The differences are illustrated by the following example.

Example 8 In a one-player game, the action set is $[0,1]$ and the payoff is 1 for a choice of a rational number and 0 for an irrational number. Consider the algebra $\mathcal{I}$ of all finite unions of subintervals of $[0,1]$. A simple measurable function $0 \leq g \leq 1$ satisfies $g \leq 1_\mathbb{Q}$ if and only if it is 0 outside some finite set of rational points, and satisfies $g \geq 1_\mathbb{Q}$ if and only if it is 1 outside some finite set of irrational points. The first fact gives that $\underline{\int} 1_\mathbb{Q}\, d\tau = 1$ for $\tau = \delta_0$, which together with the second fact proves that a strategy $\sigma: \mathcal{I} \longrightarrow [0,1]$ is a justifiable equilibrium if an only if $\sigma(\{s\}) = 0$ for all $s \notin \mathbb{Q}$. In particular, the restriction to $\mathcal{I}$ of Lebesgue measure is a justifiable equilibrium, even though it amounts to choosing a number at random and all but a countable number of choices are actually suboptimal in that they give 0 rather than 1. The necessary and sufficient condition for σ to be a best-response equilibrium is more in tune with the payoff function. This condition is $\sigma^*([0,1] \setminus \mathbb{Q}) = 0$, which is equivalent to $\sum_{s\in\mathbb{Q}\cap[0,1]} \sigma(\{s\}) = 1$. It holds if and only if σ is the restriction to $\mathcal{I}$ of some (sigma-additive) probability on the Borel sets in $[0,1]$ that is supported in $\mathbb{Q}$.

The necessary and sufficient condition for justifiable equilibrium in Example 8 would coincide with that for best-response equilibrium if the definition of the former were

[5] It is easy to see that a justifiable equilibrium σ remains so if one (or more) of the algebras $\mathcal{A}_i$ is replaced by a subalgebra, to which the strategy σ_i is restricted. Note that this is the opposite of the situation for best-response equilibria, which are preserved by extensions rather than restrictions.

strengthened by replacing the upper integral on the left-hand side of (3) with a lower integral. (Proposition 1 implies that this coincidence in fact holds for every one-player game with a bounded payoff function.) This fact illustrates the following result.

Theorem 2 In a game with bounded payoff functions, every best-response equilibrium σ is a justifiable equilibrium but not conversely.

Proof. The first assertion holds because, for every player i and strategy τ_i,

$$\overline{\int_S} u_i(s)\, d\sigma(s) \geq \int_{S_i} v_i(s_i)\, d\sigma_i(s_i) = \sup v_i \geq \overline{\int_{S_i}} v_i(s)\, d\tau_i(s_i) \geq \underline{\int_S} u_i(s)\, d(\tau_i, \sigma_{-i})(s),$$

where the equality follows from Proposition 1, the middle inequality is obvious and the other two follow from Lemma 1. The second assertion is proved by Example 8. ∎

The difference between justifiable equilibrium and best-response equilibrium goes beyond the former's use of the upper integral. It also reflects a radically different interpretation of mixed strategy. Justifiable equilibrium's perspective is an extension of the view that mixed strategies are strategies in the mixed extension of the game. This means that the mixed strategy each player plays is chosen from among, and is evaluated against, all mixed strategies. The different treatment of the chosen strategy and of the alternatives in (3) is only a concession to the potential non-integrability of the payoff function. Best-response equilibrium, by contrast, does not view players as *playing* mixed strategies. Indeed, these do not even have to be playable in any sense. A mixed strategy is an external, probabilistic and possibly incomplete description of a player's choice of action. The equilibrium condition is that it excludes actions that yield low expected payoff, where the expectation is with respect to the other players' mixed strategies (which reflects an assumption that the player's view of the others is also "external"; he has no special knowledge about their intentions). Best-response equilibrium thus describes rational choices of actions, or pure strategies, by the players. It is not interpreted as specifying choices of particular mixed, or randomized, strategies, and, correspondingly, no mixed extension of the original game is considered.

References

Dasgupta, P., Maskin, E., 1986. The existence of equilibrium in discontinuous economic games, I: Theory. *Review of Economic Studies* 53, 1–26.

Dastidar, K.G, 1995. On the existence of pure strategy Bertrand equilibrium. *Economic Theory* 5, 19–32.

Dastidar, K.G, 2011. Existence of Bertrand equilibrium revisited. *International Journal of Economic Theory* 7, 331–350.

Dunford, N., Schwartz, J.T., 1958. *Linear Operators. Part I.* Interscience, New York.

Eaton, B.C., Lipsey, R.G., 1975. The principle of minimum differentiation reconsidered: Some new developments in the theory of spatial competition. *Review of Economic Studies* 42, 27–49.

Flesch, J., Vermeulen, D., Zseleva, A., 2018. Existence of justifiable equilibrium. Mimeo.

Hoernig, S.H., 2007. Bertrand games and sharing rules. *Economic Theory* 31, 573–585.

Karlin, S., 1950. Operator treatment of minmax principle. In Kuhn, H. W., Tucker, A. W. (eds.), *Contributions to the Theory of Games* I, Annals of Mathematics Studies 24. Princeton University Press, Princeton, pp. 133–154.

Marinacci, M., 1997. Finitely additive and epsilon Nash equilibria. *International Journal of Game Theory* 26, 315–333.

Milchtaich, I., 2019. Polyequilibrium. *Games and Economic Behavior* 113, 339–355.

Osborne, M.J., Pitchik, C., 1986. The nature of equilibrium in a location model. *International Economic Review* 27, 223–237.

Shaked, A., 1982. Existence and computation of mixed strategy Nash equilibrium for 3-firms location problem. *Journal of Industrial Economics* 31, 93–96.

Sion, M., Wolfe, P., 1957. On a game without a value. In Dresher, M., Tucker, A.W., Wolfe, P. (eds.), *Contributions to the Theory of Games* III, Annals of Mathematics Studies 39. Princeton University Press, Princeton, pp. 299–306.

Vasquez, M.A., 2017. Optimistic equilibria in finitely additive mixed strategies. Mimeo.

Ville, J., 1938. Sur la théorie général des jeux où intervient l'habilité des joueurs. In *Traité du Calcul des Probabilités et de ses Applications* by É. Borel et al., vol. 2. Paris, pp. 105–113.

von Neumann, J., Morgenstern, O., 1953. *Theory of Games and Economic Behavior*, Third Edition. Princeton University Press, Princeton.

Yanovskaya, E.B., 1970. The solution of infinite zero-sum two-person games with finitely additive strategies. *Theory of Probability and its Applications* 15, 153−158.