

Solution Concepts

Jonathan Levin

April 2006

These notes discuss some of the central solution concepts for normal-form games: Nash and correlated equilibrium, iterated deletion of strictly dominated strategies, rationalizability, and self-confirming equilibrium.

1 Nash Equilibrium

Nash equilibrium captures the idea that players ought to do as well as they can given the strategies chosen by the other players.

Example 1 Prisoners' Dilemma

	<i>C</i>	<i>D</i>
<i>C</i>	1, 1	-1, 2
<i>D</i>	2, -1	0, 0

The unique Nash Equilibrium is (D, D) .

Example 2 Battle of the Sexes

	<i>B</i>	<i>F</i>
<i>B</i>	2, 1	0, 0
<i>F</i>	0, 0	1, 2

There are two pure Nash equilibria (B, B) and (F, F) and a mixed strategy equilibrium where Row plays $\frac{2}{3}B + \frac{1}{3}F$ and Column plays $\frac{1}{3}B + \frac{2}{3}F$.

Definition 1 A normal form game G consists of

1. A set of players $i = 1, 2, \dots, I$.
2. Strategy sets S_1, \dots, S_I ; let $S = S_1 \times \dots \times S_I$.

3. *Payoff functions: for each $i = 1, \dots, I$, $u_i : S \rightarrow \mathbb{R}$*

A (mixed) strategy for player i , $\sigma_i \in \Delta(S_i)$, is a probability distribution on S_i . A pure strategy places all probability weight on a single action.

Definition 2 *A strategy profile $(\sigma_1, \dots, \sigma_I)$ is a **Nash equilibrium** of G if for every i , and every $s_i \in S_i$,*

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(s_i, \sigma_{-i}).$$

And recall Nash's famous result:

Proposition 1 *Nash equilibria exist in finite games.*

A natural question, given the wide use of Nash equilibrium, is whether or why one should expect Nash behavior. One justification is that rational players ought somehow to reason their way to Nash strategies. That is, Nash equilibrium might arrive through introspection. A second justification is that Nash equilibria are self-enforcing. If players agree on a strategy profile before independently choosing their actions, then no player will have reason to deviate if the agreed profile is a Nash equilibrium. On the other hand, if the agreed profile is not a Nash equilibrium, some player can do better by breaking the agreement. A third, and final, justification is that Nash behavior might result from learning or evolution. In what follows, we take up these three ideas in turn.

2 Correlated Equilibrium

2.1 Equilibria as a Self-Enforcing Agreements

Let's start with the account of Nash equilibrium as a self-enforcing agreement. Consider Battle of the Sexes (BOS). Here, it's easy to imagine the players jointly deciding to attend the Ballet, then playing (B, B) since neither wants to unilaterally head off to Football. However, a little imagination suggests that Nash equilibrium might not allow the players sufficient freedom to communicate.

Example 2, cont. Suppose in BOS, the players flip a coin and go to the Ballet if the coin is Heads, the Football game if Tails. That is, they just randomize between two different Nash equilibria. This coin flip allows a payoff $(\frac{3}{2}, \frac{3}{2})$ that is *not* a Nash equilibrium payoff.

So at the very least, one might want to allow for randomizations between Nash equilibria under the self-enforcing agreement account of play. Moreover, the coin flip is only a primitive way to communicate prior to play. A more general form of communication is to find a mediator who can perform clever randomizations, as in the next example.

Example 3 This game has three Nash equilibria (U, L) , (D, R) and $(\frac{1}{2}U + \frac{1}{2}D, \frac{1}{2}L + \frac{1}{2}R)$ with payoffs $(5, 1)$, $(1, 5)$ and $(\frac{5}{2}, \frac{5}{2})$.

	L	R
U	$5, 1$	$0, 0$
D	$4, 4$	$1, 5$

Suppose the players find a mediator who chooses $x \in \{1, 2, 3\}$ with equal probability $\frac{1}{3}$. She then sends the following messages:

- If $x = 1 \Rightarrow$ tells Row to play U , Column to play L .
- If $x = 2 \Rightarrow$ tells Row to play D , Column to play L .
- If $x = 3 \Rightarrow$ tells Row to play D , Column to play R .

Claim. It is a Perfect Bayesian Equilibrium for the players to follow the mediator's advice.

Proof. We need to check the incentives of each player.

- If Row hears U , believes Column will play $L \Rightarrow$ play U .
- If Row hears D , believes Column will play L, R with $\frac{1}{2}, \frac{1}{2}$ probability \Rightarrow play D .
- If Column hears L , believes Row will play U, D with $\frac{1}{2}, \frac{1}{2}$ probability \Rightarrow play L .
- If Column hears R , believes Row will play $D \Rightarrow$ play R .

Thus the players will follow the mediator's suggestion. With the mediator in place, expected payoffs are $(\frac{10}{3}, \frac{10}{3})$, strictly higher than the players could get by randomizing between Nash equilibria.

2.2 Correlated Equilibrium

The notion of correlated equilibrium builds on the mediator story.

Definition 3 A *correlating mechanism* $(\Omega, \{H_i\}, p)$ consists of:

- A finite set of states Ω
- A probability distribution p on Ω .
- For each player i , a partition of Ω , denoted $\{H_i\}$. Let $h_i(\omega)$ be a function that assigns to each state $\omega \in \Omega$ the element of i 's partition to which it belongs.

Example 2, cont. In the BOS example with the coin flip, the states are $\Omega = \{\text{Heads}, \text{Tails}\}$, the probability measure is uniform on Ω , and Row and Column have the same partition, $\{\{\text{Heads}\}, \{\text{Tails}\}\}$.

Example 3, cont. In this example, the set of states is $\Omega = \{1, 2, 3\}$, the probability measure is again uniform on Ω , Row's partition is $\{\{1\}, \{2, 3\}\}$, and Column's partition is $\{\{1, 2\}, \{3\}\}$.

Definition 4 A *correlated strategy* for i is a function $f_i : \Omega \rightarrow S_i$ that is measurable with respect to i 's information partition. That is, if $h_i(\omega) = h_i(\omega')$ then $f_i(\omega) = f_i(\omega')$.

Definition 5 A strategy profile (f_1, \dots, f_I) is a *correlated equilibrium* relative to the mechanism $(\Omega, \{H_i\}, p)$ if for every i and every correlated strategy \tilde{f}_i :

$$\sum_{\omega \in \Omega} u_i(f_i(\omega), f_{-i}(\omega)) p(\omega) \geq \sum_{\omega \in \Omega} u_i(\tilde{f}_i(\omega), f_{-i}(\omega)) p(\omega) \quad (1)$$

This definition requires that f_i maximize i 's *ex ante* payoff. That is, it treats the strategy as a contingent plan to be implemented after learning the partition element. Note that this is equivalent to f_i maximizing i 's *interim* payoff for each H_i that occurs with positive probability — that is, for all i, ω , and every $s'_i \in S_i$,

$$\sum_{\omega' \in h_i(\omega)} u_i(f_i(\omega), f_{-i}(\omega')) p(\omega' | h_i(\omega)) \geq \sum_{\omega' \in h_i(\omega)} u_i(s'_i, f_{-i}(\omega')) p(\omega' | h_i(\omega))$$

Here, $p(\omega' | h_i(\omega))$ is the conditional probability on ω' given that the true state is in $h_i(\omega)$. By Bayes' Rule,

$$p(\omega' | h_i(\omega)) = \frac{\Pr(h_i(\omega) | \omega') p(\omega')}{\sum_{\omega'' \in h_i(\omega)} \Pr(h_i(\omega) | \omega'') p(\omega'')} = \frac{p(\omega')}{p(h_i(\omega))}$$

The definition of CE corresponds to the mediator story, but it's not very convenient. To search for all the correlated equilibria, one needs to consider millions of mechanisms. Fortunately, it turns out that we can focus on a special kind of correlating mechanism, called a *direct mechanism*. We will show that for any correlated equilibrium arising from some correlating mechanism, there is a correlated equilibrium arising from the direct mechanism that is precisely equivalent in terms of behavioral outcomes. Thus by focusing on one special class of mechanism, we can capture all possible correlated equilibria.

Definition 6 A *direct mechanism* has $\Omega = S$, $h_i(s) = \{s' \in S : s'_i = s_i\}$, and some probability distribution q over pure strategy profiles.

Proposition 2 Suppose f is a correlated equilibrium relative to $(\Omega, \{H_i\}, p)$. Define $q(s) \equiv \Pr(f(\omega) = s)$. Then the strategy profile \tilde{f} with $\tilde{f}_i(s) = s_i$ for all $i, s \in S$ is a correlated equilibrium relative to the direct mechanism $(S, \{\tilde{H}_i\}, q)$.

Proof. Suppose that s_i is recommended to i with positive probability, so $p(s_i, s_{-i}) > 0$ for some s_{-i} . We check that under the direct mechanism $(S, \{\tilde{H}_i\}, q)$, player i cannot benefit from choosing another strategy s'_i when s_i is suggested. If s_i is recommended, then i 's expected payoff from playing s'_i is:

$$\sum_{s_{-i} \in S_{-i}} u_i(s'_i, s_{-i})q(s_{-i}|s_i).$$

The result is trivial if there is only one information set H_i in the original mechanism for which $f_i(H_i) = s_i$. In this case, conditioning on s_i is the same as conditioning on H_i in the original. More generally, we substitute for q to obtain:

$$\frac{1}{\Pr(f_i(\omega) = s_i)} \cdot \sum_{\omega|f_i(\omega)=s_i} u_i(s'_i, f_{-i}(\omega))p(\omega).$$

Re-arranging to separate each H_i at which s_i is optimal:

$$\frac{1}{\Pr(f_i(\omega) = s_i)} \cdot \sum_{H_i|f_i(H_i)=s_i} \Pr(H_i) \left[\sum_{\omega \in H_i} u_i(s'_i, f_{-i}(\omega))p(\omega|H_i) \right]$$

Since $(\Omega, \{H_i\}, p, f)$ is a correlated equilibrium, each bracketed term for which $\Pr(H_i) > 0$ is maximized at $f_i(H_i) = s_i$. So s_i is optimal given recommendation s_i . *Q.E.D.*

Thus what really matters in correlated equilibrium is the probability distribution over strategy profiles. We refer to any probability distribution q over strategy profiles that arises as the result of a correlated equilibrium as a *correlated equilibrium distribution (c.e.d.)*.

Example 2, cont. In the BOS example, the c.e.d. is $\frac{1}{2}(B, B), \frac{1}{2}(F, F)$.

Example 3, cont. In this example, the c.e.d. is $\frac{1}{3}(U, L), \frac{1}{3}(D, L), \frac{1}{3}(D, R)$.

The next result characterizes correlated equilibrium distributions.

Proposition 3 *The distribution $q \in \Delta(S)$ is a correlated equilibrium distribution if and only if for all i , every s_i with $q(s_i) > 0$ and every $s'_i \in S_i$,*

$$\sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) q(s_{-i}|s_i) \geq \sum_{s_{-i} \in S_{-i}} u_i(s'_i, s_{-i}) q(s_{-i}|s_i). \quad (2)$$

Proof. (\Leftarrow) Suppose q satisfies (2). Then the “obedient” profile f with $f_i(s) = s_i$ is a correlated equilibrium given the direct mechanism $(S, \{H_i\}, q)$ since (2) says precisely that with this mechanism s_i is optimal for i given recommendation s_i . (\Rightarrow) Conversely, if q arises from a correlated equilibrium, the previous result says that the obedient profile must be a correlated equilibrium relative to the direct mechanism $(S, \{H_i\}, q)$. Thus for all i and all recommendations s_i occurring with positive probability, s_i must be optimal — i.e. (2) must hold. *Q.E.D.*

Consider a few properties of correlated equilibrium.

Property 1 Any Nash equilibrium is a correlated equilibrium

Proof. Need to ask if (2) holds for the probability distribution q over outcomes induced by the NE. For a pure equilibrium s^* , we have $q(s_{-i}^*|s_i^*) = 1$ and $q(s_{-i}|s_i^*) = 0$ for any $s_{-i} \neq s_{-i}^*$. Therefore (2) requires for all i, s_i :

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*).$$

This is precisely the definition of NE. For a mixed equilibrium, σ^* , we have that for any s_i^* in the support of σ_i^* , $q(s_{-i}|s_i^*) = \sigma_{-i}(s_{-i})$. This follows from the fact that in a mixed NE, the players mix independently. Therefore (2) requires that for all i, s_i^* in the support of σ_i^* , and s_i ,

$$\sum_{s_{-i} \in S_{-i}} u_i(s_i^*, s_{-i}) \sigma_{-i}(s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) \sigma_{-i}(s_{-i}),$$

again, the definition of a mixed NE. *Q.E.D.*

Property 2 Correlated equilibria exist in finite games.

Proof. Any NE is a CE, and NE exists. Hart and Schmeidler (1989) show the existence of CE directly, exploiting the fact that a CE is just a probability distribution q satisfying a system of linear inequalities. Their proof does not appeal to fixed point results! *Q.E.D.*

Property 3 The sets of correlated equilibrium distributions and payoffs are convex.

Proof. Left as an exercise.

2.3 Subjective Correlated Equilibrium

The definition of correlated equilibrium assumes the players share a common prior p over the set of states (or equivalently share the same probability distribution over equilibrium play). A significantly weaker notion of equilibrium obtains if this is relaxed. For this, let p_1, p_2, \dots, p_I be *distinct* probability measures on Ω .

Definition 7 The profile f is a **subjective correlated equilibrium** relative to the mechanism $(\Omega, \{H_i\}, p_1, \dots, p_I)$ if for every i, ω and every alternative strategy \tilde{f}_i ,

$$\sum_{\omega \in \Omega} u_i(f_i(\omega), f_{-i}(\omega))p_i(\omega) \geq \sum_{\omega \in \Omega} u_i(\tilde{f}_i(\omega), f_{-i}(\omega))p_i(\omega)$$

Example 3, cont. Returning to our example from above,

	L	R
U	$5, 1$	$0, 0$
D	$4, 4$	$1, 5$

Here, $(4, 4)$ can be obtained as a SCE payoff. Simply consider the direct mechanism with $p_1 = p_2 = \frac{1}{3}(U, L) + \frac{1}{3}(D, L) + \frac{1}{3}(D, R)$. This is a SCE, and since there is no requirement that the players have *objectively correct* beliefs about play, it may be that (D, L) is played with probability one!

2.4 Comments

1. The difference between mixed strategy Nash equilibria and correlated equilibria is that mixing is *independent* in NE. With more than two players, it may be important in CE that one player believes others are correlating their strategies. Consider the following example from Aumann (1987) with three players: Row, Column and Matrix.

0, 0, 3	0, 0, 0	2, 2, 2	0, 0, 0	0, 0, 0	0, 0, 0
1, 0, 0	0, 0, 0	0, 0, 0	2, 2, 2	0, 1, 0	0, 0, 3

No NE gives any player more than 1, but there is a CE that gives everyone 2. Matrix picks middle, and Row and Column pick (Up,Left) and (Down,Right) each with probability $\frac{1}{2}$. The key here is that Matrix must expect Row to pick Up precisely when Column picks Left.

2. Note that in CE, however, each agent uses a pure strategy — he just is uncertain about others’ strategies. So this seems a bit different than mixed NE if one views a mixed strategy as an explicit randomization in behavior by each agent i . However, another view of mixed NE is that it’s not i ’s actual choice that matters, but j ’s beliefs about i ’s choice. On this account, we view σ_i as what others expect of i , and i as simply doing some (pure strategy) best response to σ_{-i} . This view, which is consistent with CE, was developed by Harsanyi (1973), who introduced small privately observed payoff perturbations so that in pure strategy BNE, players would be uncertain about others behavior. His “purification theorem” showed that these pure strategy BNE are observably equivalent to mixed NE of the unperturbed game if the perturbations are small and independent.
3. Returning to our pre-play communication account, one might ask if a mediator is actually needed, or if the players could just communicate by flipping coins and talking. With two players, it should be clear from the example above that the mediator is crucial in allowing for messages that are not common knowledge. However, Barany (1992) shows that if $I \geq 4$, then any correlated equilibrium payoff (with rational numbers) can be achieved as the Nash equilibrium of an extended game where prior to play the players communicate through cheap talk. Girardi (2001) shows the same can be done as a sequential equilibrium provided $I \geq 5$. For the case of two players, Aumann and Hart (2003) characterize the set of attainable payoffs if players can communicate freely, but without a mediator, prior to playing the game.

3 Rationalizability and Iterated Dominance

Bernheim (1984) and Pearce (1984) investigated the question of whether one should expect rational players to introspect their way to Nash equilibrium play. They argued that even if rationality was common knowledge, this should not generally be expected. Their account takes a view of strategic behavior that is deeply rooted in single-agent decision theory.

To discuss these ideas, it's useful to explicitly define rationality.

Definition 8 *A player is rational if he chooses a strategy that maximizes his expected payoff given his belief about opponents' strategies.*

Note that assessing rationality requires defining beliefs, something that the formal definition of Nash equilibrium does not require. Therefore, as a matter of interpretation, if we're talking about economic agents playing a game, we might say that Nash equilibrium arises when each player is rational and know his opponents' action profile. But we could also talk about Nash equilibrium in an evolutionary model of fish populations without ever mentioning rationality.

3.1 (Correlated) Rationalizability

Rationalizability imposes two requirements on strategic behavior.

1. Players maximize with respect to their beliefs about what opponents will do (i.e. are rational).
2. Beliefs cannot conflict with other players being rational, and being aware of each other's rationality, and so on (but they need not be correct).

Example 4 In this game (from Bernheim, 1984), there is a unique Nash equilibrium (a_2, b_2) . Nevertheless a_1, a_3, b_1, b_3 can all be rationalized.

	b_1	b_2	b_3	b_4
a_1	0, 7	2, 5	7, 0	0, 1
a_2	5, 2	3, 3	5, 2	0, 1
a_3	7, 0	2, 5	0, 7	0, 1
a_4	0, 0	0, -2	0, 0	10, -1

- Row will play a_1 if Column plays b_3
- Column will play b_3 if Row plays a_3

- Row will play a_3 if Column plays b_1
- Column will play b_1 if Row plays a_1

This “chain of justification” rationalizes a_1, a_3, b_1, b_3 . Of course a_2 and b_2 rationalize each other. However, b_4 cannot be rationalized, and since no rational player would play b_4 , a_4 can’t be rationalized.

Definition 9 A subset $B_1 \times \dots \times B_I \subset S$ is a **best reply set** if for all i and all $s_i \in B_i$, there exists $\sigma_{-i} \in \Delta(B_{-i})$ to which s_i is a best reply.

- In the definition, note that σ_{-i} can reflect correlation — it need not be a mixed strategy profile for the opponents. This allows for more “rationalizing” than if opponents mix independently. More on this later.

Definition 10 The set of **correlated rationalizable strategies** is the component by component union of all best reply sets:

$$R = R_1 \times \dots \times R_I = \bigcup_{\alpha} B_1^{\alpha} \times \dots \times B_I^{\alpha}$$

where each $B^{\alpha} = B_1^{\alpha} \times \dots \times B_I^{\alpha}$ is a best reply set.

Proposition 4 R is the maximal best reply set.

Proof. Suppose $s_i \in R_i$. Then $s_i \in B_i^{\alpha}$ for some α . So s_i is a best reply to some $\sigma_{-i} \in \Delta(B_{-i}^{\alpha}) \subset \Delta(R_{-i})$. So R_i is a best reply set. Since it contains all others, it is maximal. *Q.E.D.*

3.2 Iterated Strict Dominance

In contrast to asking what players might do, iterated strict dominance asks what players *won't* do, and what they won't do conditional on other players not doing certain things, and so on. Recall that a strategy s_i is *strictly dominated* if there is some mixed strategy σ_i such that $u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$, and that iterated dominance applies this definition repeatedly.

- Let $S_i^0 = S_i$
- Let $S_i^k = \left\{ \begin{array}{l} s_i \in S_i^{k-1} : \text{There is no } \sigma_i \in \Delta(S_i^{k-1}) \text{ s.t.} \\ u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i}^{k-1} \end{array} \right\}$
- Let $S_i^{\infty} = \bigcap_{k=1}^{\infty} S_i^k$.

Iterated strict dominance never eliminates Nash equilibrium strategies, or any strategy played with positive probability in a correlated equilibrium (proof left as an exercise!). Indeed it is often quite weak. Most games, including many games with a unique Nash equilibrium, are not dominance solvable.

Example 4, cont. In this example, b_4 is strictly dominated. Eliminating b_4 means that a_4 is also strictly dominated. But no other strategy can be eliminated.

Proposition 5 *In finite games, iterated strict dominance and correlated rationalizability give the same solution set, i.e. $S_i^\infty = R_i$.*

This result is suggested by the following Lemma (proved in the last section of these notes).

Lemma 1 *A pure strategy in a finite game is a best response to some beliefs about opponent play if and only if it is not strictly dominated.*

Proof of Proposition.

$R \subset S^\infty$. If $s_i \in R_i$, then s_i is a best response to some belief over R_{-i} . Since $R_{-i} \subset S_{-i}$, Lemma 1 implies that s_i is not strictly dominated. Thus $R_i \subset S_i^2$ for all i . Iterating this argument implies that $R_i \subset S_i^k$ for all i, k , so $R_i \subset S_i^\infty$.

$S^\infty \subset R$. It suffices to show that S^∞ is a best-reply set. By definition, no strategy in S^∞ is strictly dominated in the game in which the set of actions is S^∞ . Thus, any $s_i \in S_i^\infty$ must be a best response to some beliefs over S_{-i}^∞ . *Q.E.D.*

3.3 Comments

1. Bernheim (1984) and Pearce (1984) originally defined rationalizability assuming that players would expect opponents to mix independently. So B is a best reply set if $\forall s_i \in B_i$, there is some $\sigma_{-i} \in \times_{j \neq i} \Delta(S_j)$ to which s_i is a best reply. For $I = 2$, this makes no difference, but when $I \geq 3$, their concept refines ISD (it rules out more strategies).
2. Brandenburger and Dekel (1987) relate correlated rationalizability to subjective correlated equilibrium. While SCE is more permissive than rationalizability, requiring players to have well defined conditional beliefs (and maximize accordingly) even for states $\omega \in \Omega$ to which they assign zero probability leads to a refinement of SCE that is the same as correlated rationalizability.

3. One way to categorize the different solution concepts is to note that if one starts with rationality, and common knowledge of rationality, the concepts differ precisely in how they further restrict the beliefs of the players about the distribution of play. Doug Bernheim suggests the following table:

	Different Priors	Common Prior
Correlation	Corr. Rationalizability/ ISD/Refined Subj CE	Correlated Equilibrium
Independence	Rationalizability	Nash Equilibrium

3.4 Appendix: Omitted Proof

For completeness, this section provides a proof of the Lemma equating dominated strategies with those that are never a best response. The proof requires a separation argument, so let's first recall the Duality Theorem for linear programming. To do this, start with the following problem:

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \sum_{j=1}^n c_j x_j & \quad (3) \\ \text{s.t. } \sum_{j=1}^n a_{ij} x_j \geq b_i & \quad \forall i = 1, \dots, m \end{aligned}$$

This problem has the same solution as

$$\max_{y \in \mathbb{R}_+^m} \left(\min_{x \in \mathbb{R}^n} \sum_{j=1}^n c_j x_j + \sum_{i=1}^m y_i \left(b_i - \sum_{j=1}^n a_{ij} x_j \right) \right). \quad (4)$$

Rearranging terms, we obtain

$$\max_{y \in \mathbb{R}_+^m} \left(\min_{x \in \mathbb{R}^n} \sum_{j=1}^n \left(c_j - \sum_{i=1}^m y_i a_{ij} \right) x_j + \sum_{i=1}^m y_i b_i \right). \quad (5)$$

Swapping the order of optimization give us

$$\min_{x \in \mathbb{R}^n} \left(\max_{y \in \mathbb{R}_+^m} \sum_{j=1}^n \left(c_j - \sum_{i=1}^m y_i a_{ij} \right) x_j + \sum_{i=1}^m y_i b_i \right), \quad (6)$$

which can be related to the following “dual” problem:

$$\begin{aligned} & \max_{y \in \mathbb{R}_+^m} \sum_{i=1}^m y_i b_i \\ \text{s.t.} \quad & \sum_{j=1}^n \left(c_j - \sum_{i=1}^m y_i a_{ij} \right) = 0 \quad \forall j = 1, \dots, n. \end{aligned} \tag{7}$$

Theorem 1 *Suppose problems (3) and (7) are feasible (i.e. have non-empty constraint sets). Then their solutions are the same.*

We use the duality theorem to prove the desired Lemma.

Lemma 2 *A pure strategy in a finite game is a best response to some beliefs about opponent play if and only if it is not strictly dominated.*

Proof (Myerson, 1991). Let $s_i \in S_i$ be given. Our proof will be based on a comparison of two linear programming problems.

Problem I:

$$\begin{aligned} & \min_{\sigma_{-i}, \delta} \delta \\ \text{s.t.} \quad & \sigma_{-i}(s_{-i}) \geq 0 \quad \forall s_{-i} \in S_{-i} \\ & \sum_{s_{-i}} \sigma_{-i}(s_{-i}) \geq 1 \quad \text{and} \quad - \sum_{s_{-i}} \sigma_{-i}(s_{-i}) \geq -1 \\ & \delta + \sum_{s_{-i}} \sigma_{-i}(s_{-i}) [u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})] \geq 0 \quad \forall s'_i \in S_i \end{aligned}$$

Observe that s_i is a best response to some beliefs over opponent play if any only if the solution to this problem is less than or equal to zero.

Problem II:

$$\begin{aligned} & \max_{\eta, \varepsilon_1, \varepsilon_2, \sigma_i} \varepsilon_1 - \varepsilon_2 \\ \text{s.t.} \quad & \varepsilon_1, \varepsilon_2 \in \mathbb{R}_+, \sigma_i \in \mathbb{R}_+^{|S_i|}, \eta \in \mathbb{R}_+^{|S_{-i}|} \\ & \sum_{s'_i} \sigma_i(s'_i) \geq 1 \\ & \eta(s_{-i}) + \varepsilon_1 - \varepsilon_2 + \sum_{s'_i} \sigma_i(s'_i) [u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})] = 0 \quad \forall s_{-i} \in S_{-i} \end{aligned}$$

Observe that s_i is strictly dominated if and only if the solution to this problem is strictly greater than zero — i.e. s_i is *not* strictly dominated if and only if the solution to this problem is less than or equal to zero.

Finally, the Duality Theorem for linear programming says that so long as these two problems are feasible (have non-empty constraint sets), their solutions must be the same, establishing the result. *Q.E.D.*

4 Self-Confirming Equilibria

The third possible foundation for equilibrium is learning. We'll look at explicit learning processes later; for now, we ask what might happen as the end result of a learning processes. For instance, if a learning process settles down into steady-state play, will this be a Nash Equilibrium? Fudenberg and Levine (1993) suggest that a natural end-result of learning is what they call *self-confirming equilibria*. In a self-confirming equilibrium:

1. Players maximize with respect to their beliefs about what opponents will do (i.e. are rational).
2. Beliefs cannot conflict with the empirical evidence (i.e. must match the empirical distribution of play).

The difference with Nash equilibrium and rationalizability lies in the restriction on beliefs. In a Nash equilibrium, players hold correct beliefs about opponents' strategies and hence about their behavior. By contrast, with rationalizability, beliefs need not be correct, they just can't conflict with rationality. With SCE, beliefs need to be consistent with available data.


4.1 Examples of Self-Confirming Equilibria

In a simultaneous game, assuming actions are observed after every period, every Nash equilibrium is self-confirming. Moreover, any self-confirming equilibrium is Nash.

Example 1 Consider matching pennies. It is an SCE for both players to mix 50/50 and to both believe the other is mixing 50/50. On the other hand, if player i believes anything else, he must play a pure strategy. But then player j must believe i will play this strategy or else he would eventually be proved wrong. So the only SCE is the same as the NE.

In extensive form games, the situation is different, as the next example shows.

Example 2 In the entry game below, the only Nash equilibria are (Enter, Don't Fight) and (Don't Enter, Fight). These equilibria, with correct beliefs, are also self-confirming equilibria.




An Entry Game

There is also another SCE, however, where Player 1 plays Don't Enter and believes Fight, while Player 2 player plays Don't Fight and believes Don't Enter. In this SCE, player 1 has the wrong beliefs, but since he never enters, they're never contradicted by the data!

In the entry game, the non-Nash SCE is indistinguishable from a Nash equilibrium in terms of observed behavior. But even that need not be the case, as the next example shows.

Example 3 (Fudenberg-Kreps, 1993) Consider the three player game below.



SCE are different than NE

In this game, there is a self-confirming equilibrium where (A_1, A_2) is played. In this equilibrium, player 1 expects player 3 to play R , while player 2 expects 3 to play L . Given these beliefs, the optimal strategy for player 1 is to play A_1 , while the optimal strategy for player 2 is to play A_2 . Player 3's beliefs and strategy can be arbitrary so long as the strategy is optimal given beliefs.

The key point here is that there is *no* Nash equilibrium where (A_1, A_2) is played. The reason is that in any Nash equilibrium, players 1 and 2 *must* have the *same* (correct) beliefs about player 3's strategy. But if they have the same beliefs, then at least one of them must want to play D .

The distinction between Nash and self-confirming equilibria in extensive form games arises because players do not get to observe all the relevant information about their opponents' behavior. The same issue can arise in simultaneous-move games as well if the information feedback that players' receive is limited.

Example 4 Consider the following two-player game. Suppose that after the game is played, the column player observes the row player's action, but the row player observes only whether or not the column player chose R , and gets no information about her payoff.

	L	M	R
U	2, 0	0, 2	0, 0
D	0, 0	2, 0	3, 3

This game has a unique Nash equilibrium, (D, R) ; indeed the game is dominance-solvable. The profile (D, R) is self-confirming too; however that is not the only self-confirming profile. The profile (U, M) is also self-confirming. In this SCE, column has correct beliefs, but row believes that column will play L . This mistaken belief isn't refuted by the evidence because all row observes is that column does not play R . There are also SCE where row mixes, and where both players mix.

Note that in the (U, M) SCE, row's beliefs do not respect column's rationality. This suggests that one might refine SCE by further restricting beliefs to respect rationality or common knowledge of rationality — Dekel, Fudenberg and Levine (1998) and Esponda (2006) explore this possibility.

The above example is somewhat contrived, but the idea that players might get only partial feedback, and this might affect the outcome of learning, is natural. For instance, in sealed-bid auctions it is relatively common to announce only the winner and possibility not even the winning price, so the information available to form beliefs is rather limited.

4.2 Formal Definition of SCE

To define self-confirming equilibrium in extensive form games, let s_i denote a strategy for player i , and σ_i a mixture over such strategies. Let H_i denote the set of information sets at which i moves, and $H(s_i, \sigma_{-i})$ denote the set of information sets that can be reached if player i plays s_i and opponents play σ_{-i} . Let $\pi_i(h_i|\sigma_i)$ denote the mixture over actions that results at information set h_i , if player i is using the strategy σ_i (i.e. π_i is the behavior strategy induced by the mixed strategy σ_i). Let μ_i denote a belief over $\Pi_{-i} = \times_{j \neq i} \Pi_j$ the product set of other players' behavior strategies.

Definition 11 *A profile σ is a Nash equilibrium if for each $s_i \in \text{support}(\sigma_i)$, there exists a belief μ_i such that (i) s_i maximizes i 's expected payoff given beliefs μ_i , and (ii) player i 's beliefs are correct, for all $h_j \in H_{-i}$*

$$\mu_i [\{\pi_{-i} \mid \pi_j(h_j) = \pi_j(h_j|\sigma_j)\}] = 1.$$

The way to read this is that for each information set at which some player $j \neq i$ moves, player i 's belief puts a point mass on the probability distribution that exactly coincides with the distribution induced by j 's strategy. Thus i has correct beliefs at all opponent information sets.

Definition 12 *A profile σ is a Self-Confirming equilibrium if for each $s_i \in \text{support}(\sigma_i)$, there exists a belief μ_i such that (i) s_i maximizes i 's expected payoffs given beliefs μ_i , and (ii) player i 's beliefs are empirically correct, for all histories $h_j \in H(s_i, \sigma_{-i})$ and all $j \neq i$*

$$\mu_i [\{\pi_{-i} \mid \pi_j(h_j) = \pi_j(h_j|\sigma_j)\}] = 1.$$

The difference is that in self-confirming equilibrium, beliefs must be correct only for reachable histories. This definition assumes that the players observe their opponents' actions perfectly, in contrast to the last example above; it's not hard to generalize the definition. Note that this definition formally encompasses games of incomplete information (where Nature moves first); Dekel et. al (2004) study SCE in these games.

References

- [1] Aumann, R. (1974) "Subjectivity and Correlation in Randomized Strategies," *J. Math. Econ.*
- [2] Aumann, R. (1987) "Correlated Equilibrium as an Extension of Bayesian Rationality," *Econometrica*, 55, 1–18.
- [3] Aumann, R. and S. Hart (2003) "Long Cheap Talk," *Econometrica*.
- [4] Barany, I. (1992) "Fair Distribution Protocols or How Players Replace Fortunes," *Math. Op. Res.*, 327–340.
- [5] Bernheim, D. (1984) "Rationalizable Strategic Behavior," *Econometrica*, 52, 1007–1028.
- [6] Brandenburger, A. and E. Dekel (1987) "Rationalizability and Correlated Equilibrium," *Econometrica*, 55, 1391–1402
- [7] Dekel, E., D. Fudenberg and D. Levine (2004) "Learning to Play Bayesian Games," *Games and Econ. Behav.*, 46, 282–303.
- [8] Fudenberg, D. and David Kreps (1993) "Learning Mixed Equilibria," *Games and Econ. Behav.*, 5, 320–367.
- [9] Fudenberg, D. and D. Levine (1993) "Self-Confirming Equilibrium," *Econometrica*, 61, 523–546.
- [10] Gerardi, D. (2001) "Unmediated Communication in Games with Complete and Incomplete Communication," *J. Econ. Theory*.
- [11] Harsanyi, J. (1973) "On Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points," *Int. J. Game Theory*, 2, 1–23.
- [12] Hart, S. and D. Schmeidler (1989) "Existence of Correlated Equilibria," *Math. Op. Res.* 14, 18–25.
- [13] Holt, C. and A. Roth (2004) "The Nash Equilibrium: A Perspective," *PNAS*, 101(12), 3999–4002.
- [14] Pearce, D. (1984) "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029–1050.