If we allow players to randomize over their strategies, then a vector of mixed strategies, $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_N)$ is a mixed-strategy Nash equilibrium if for every player i, given a profile of mixed strategies of the other players $\boldsymbol{\sigma}_{-i}$, player i's mixed strategy σ_i is optimal. A game can have no Nash equilibrium, a unique Nash equilibrium, or many equilibria (pure-strategy or mixed).

Here are two basic results on the existence of pure and mixed-strategy Nash equilibria.

PROPOSITION F.16 *Every game with finite strategy sets for all the players has a mixed-strategy Nash equilibrium.*

PROPOSITION F.17 *If the strategy sets for the players S_i are nonempty, convex and compact subsets of $\Re^m$, $u_i(\mathbf{s})$ is continuous in* s *and quasiconcave in s_i for all i, then the game has a pure-strategy Nash equilibrium.*

However, in either case, there is no guarantee that the equilibrium is unique.

Bayesian Nash Equilibrium

Games with *incomplete information* model situations where the players do not know with certainty what the other players' strategy sets, parameters, and utility functions are. Each forms a probabilistic view of the other players' private information (akin to a Bayesian prior; this probabilistic view may be updated in a repeated game as the game reveals more information to the players).

The model is as follows: Player i's payoff function is now given by $u_i(\mathbf{s}, \theta_i)$, where $\theta_i \in \Theta_i$ is a random variable whose realization is observed only by player i. Let $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$ and $\boldsymbol{\Theta} = \Theta_1 \times \Theta_2 \times \cdots \times \Theta_N$. However, the joint probability distribution of $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, $F(\boldsymbol{\theta})$ is common knowledge among the players. The Bayesian game is then $\Gamma = [N, \{S_i\}, \{u_i(\cdot)\}, \boldsymbol{\Theta}, F(\cdot)]$.

A pure-strategy for player i is in this case a decision rule $s_i(\theta_i)$. His strategy is a function of the realization of his θ_i. Given a vector of pure strategies for all the players $(s_1(), \ldots, s_N(\cdot)$, player i's payoff is given by the expectation over the θ's:

$$\tilde{u}_i(s_i(\cdot), \mathbf{s}_{-i}(\cdot)) = E_{\boldsymbol{\theta}}[u_i(s_i(\cdot), \mathbf{s}_{-i}(\cdot))].$$

The extension to the Nash equilibrium concept is then as follows. A *(pure-strategy) Bayesian Nash equilibrium* is a vector of decision rules $\mathbf{s}(\cdot) = (s_1(\cdot), \ldots, s_N(\cdot)$ such that

$$\tilde{u}_i(s_i(\cdot), \mathbf{s}_{-i}(\cdot)) \geq \tilde{u}_i(s_i^{'}(\cdot), \mathbf{s}_{-i}(\cdot)).$$

Repeated Games

Finite repeated games are one-shot games that are repeated over a number of periods. At the beginning of each period, firms are aware of the others' past moves and make their decisions simultaneously and noncooperatively for that period.

For example, a repeated Bertrand game would have firms setting prices simultaneously at the beginning of each period, and a repeated Cournot game would have firms deciding how much to produce at the beginning of each period. For instance, when all the firms in the market post their prices on a centralized industrywide reservation system every day, the repeated game's period is one day, and at the beginning of each day firms set prices simultaneously without knowing how the other firms will choose their prices that day.