Mixed Strategy Equilibrium

Abstract: A mixed strategy is a probability distribution one uses to randomly choose among available actions in order to avoid being predictable. In a *mixed strategy equilibrium* each player in a game is using a mixed strategy, one that is best for him against the strategies the other players are using. In laboratory experiments the behavior of inexperienced subjects has generally been inconsistent with the theory in important respects; data obtained from contests in professional sports conforms much more closely with the theory.

In many strategic situations a player's success depends upon his actions being unpredictable. Competitive sports are replete with examples. One of the simplest occurs repeatedly in soccer (football): if a kicker knows which side of the goal the goaltender has chosen to defend, he will kick to the opposite side; and if the goaltender knows to which side the kicker will direct his kick, he will choose that side to defend. In the language of game theory, this is a simple 2×2 game which has no pure strategy equilibrium.

John von Neumann's (1928) theoretical formulation and analysis of such strategic situations is generally regarded as the birth of game theory. von Neumann introduced the concept of a *mixed strategy*: each player in our soccer example should choose his Left or Right action randomly, but according to some particular binomial process. Every *zero sum* two-person game in which each player's set of available strategies is finite must have a *value* (or *security level)* for each player and each player must have at least one *minimax* strategy – a strategy that assures him that no matter how his opponent plays, he will achieve at least his security level for the game, in expected value terms. In many such games the minimax strategies are pure strategies, requiring no mixing; in others, they are mixed strategies.

John Nash (1950) introduced the powerful notion of *equilibrium* in games (including non-zero-sum games and games with an arbitrary number of players): an equilibrium is a combination of strategies (one for each player) in which each player's strategy is a *best* strategy for him against the strategies all the other players are using. An equilibrium is thus a sustainable combination of strategies, in the sense that no player has an incentive to change unilaterally to a different strategy. A *mixed-strategy equilibrium* (MSE) is one in which each player is using a mixed strategy; if a game's only equilibria are mixed, we say it is an MSE game. In two-person zero-sum games there is an equivalence between minimax and equilibrium: it is an equilibrium for each player to use a minimax strategy, and an equilibrium can consist only of minimax strategies.

An example or two will be helpful. First consider the game Tic-tac-toe. There are three possible outcomes: Player A wins, Player B wins, or the game ends in a draw. Fully defining the players' possible strategies is somewhat complex, but anyone who has played the game more than a few times knows that each player has a strategy that guarantees him no worse than a draw. These are the players' respective minimax strategies and they constitute an equilibrium. Since they are *pure strategies* (requiring no mixing), Tic-tac-toe is not an MSE game.

A second example is the game called Matching Pennies. Each player places a penny either heads up or tails up; the players reveal their choices to one another simultaneously; if their choices match, Player A gives his penny to Player B, otherwise Player B gives his penny to Player A. This game has only two possible outcomes and it is obviously zero-sum. Neither of a player's pure strategies (Heads or Tails) ensures that he won't lose. But by choosing Heads or Tails randomly, each with a ½ probability (for example, by "flipping" the coin), he ensures that in expected value his payoff will be zero *no matter how his opponent plays*. This 50-50 mixture of Heads and Tails is thus a minimax strategy for each player, and it is an MSE of the game for each player to choose his minimax strategy.

Figure 1 provides a matrix representation of Matching Pennies. Player A, when choosing Heads or Tails, is effectively choosing one of the matrix's two rows; Player B chooses one of the columns; the cell at the resulting row-and-column intersection indicates Player A's *payoff*. Player B's payoff need not be shown, since it is the negative of Player A's (as always in a zero-sum game). Matching Pennies is an example of a 2×2 game: each player has two pure strategies, and the game's matrix is therefore 2×2.

- - Figures 1 and 2 should go about here - -

Figure 2 depicts our soccer example, another 2×2 MSE game. The kicker and the goalie simultaneously choose either Left or Right; the number in the resulting cell (at the row-and-column intersection) is the probability a goal will be scored, given the players' choices. The probabilities capture the fact that for each combination of choices by kicker and goalie the outcome is still random – a goal is less likely (but not impossible) when their choices match and is more likely (while not certain) when they don't. The specific probabilities will depend upon the abilities of the specific kicker and goalie: the probabilities in Figure 2 might represent, for example, a situation in which the kicker is more effective kicking to the left half of the goal than to the right half. For the specific game in Figure 2 it can be shown that the kicker's minimax strategy is a 50-50 mix between Left and Right and the goalie's minimax strategy is to defend Left 3/5 of the time and Right 2/5. The reader can easily see that the value of the game is therefore 3/5, *i.e.*, in the MSE the kicker will succeed in scoring a goal 60% of the time.

Non-zero-sum games and games with more than two players often have mixed strategy equilibria as well. Important examples are decisions whether to enter a competition (such as an industry, a tournament, or an auction), "wars of attrition" (decisions about whether and when to exit a competition), and models of price dispersion (which explain how the same good may sell at different prices), as well as many others.

How do people actually behave in strategic situations that have mixed strategy equilibria? Does the MSE provide an accurate description of people's behavior? Virtually from the moment Nash's 1950 paper was distributed in preprint, researchers began to devise experiments in which human subjects play games that have mixed strategy equilibria. The theory has not fared well in these experiments. The behavior observed in experiments typically departs from the MSE in two ways: participants do not generally play their strategies in the proportions dictated by the game's particular MSE probability distribution; and their choices typically exhibit negative serial correlation – a player's mixed strategy in an MSE requires that his choices be independent across multiple plays, but experimental subjects tend instead to switch from one action to another more often than chance would dictate. Experimental psychologists have reported similar "switching too often" in many experiments designed to determine people's ability to intentionally behave randomly. The evidence suggests that humans are not very good at behaving randomly.

The results from experiments were so consistently at variance with the theory that empirical analysis of the concept of MSE became all but moribund for nearly two decades, until interest was revived by Barry O'Neill's (1987) seminal paper. O'Neill pointed out that there were features of previous experiments that subtly invalidated them as tests of the theory of mixed strategy equilibrium, and he devised a clever but simple experiment that avoided these flaws. Although James Brown and Robert Rosenthal (1990) subsequently demonstrated that the behavior of O'Neill's subjects was still inconsistent with the theory, the correspondence between theory and observation was nevertheless closer in his experiment than in prior experiments.

Mark Walker and John Wooders (2001) were the first to use field data instead of experiments to evaluate the theory of mixed strategy equilibrium. They contended that while the rules and mechanics of a simple MSE game may be easy to learn quickly, as required in a laboratory experiment, substantial experience is nevertheless required in order to develop an understanding of the strategic subtleties of playing even simple MSE games. In short, an MSE game may be easy to play but not easy to play *well*. This fact alone may account for much of the theory's failure in laboratory experiments.

Instead of using experiments, Walker and Wooders applied the MSE theory to data from professional tennis matches. The "serve" in tennis can be described as a 2×2 MSE game exactly like the soccer example in Figure 2: the server chooses which direction to serve, the receiver chooses which direction to defend, and the resulting payoff is the probability the server wins the point. Walker and Wooders obtained data from matches between the best players in the world, players who have devoted their lives to the sport and should therefore be expert in the strategic subtleties of this MSE game. Play by these world-class tennis players was found to correspond quite closely to the MSE predictions. Subsequent research by others, with data from professional tennis and soccer matches, has shown a similar correspondence between theory and observed behavior.

Thus, the empirical evidence to date indicates that MSE is effective for explaining and predicting behavior in strategic situations at which the competitors are experts and that it is less effective when the competitors are novices, as experimental subjects typically are. This leaves several obvious open questions: In view of the enormous disparity in expertise between world-class athletes and novice experimental subjects, how can we determine, for specific players, whether the MSE yields an appropriate prediction or explanation of their play? And when MSE is not appropriate, what *is* a good theory of play? We clearly need a generalization of current theory, one that includes MSE, that tells us in addition when MSE is "correct," and that explains behavior when MSE is not correct. Moreover, the need for such a theory extends beyond MSE games to the theory of games more generally.

A more general theory will likely comprise either an alternative, more general notion of equilibrium or a theory of out-of-equilibrium behavior in which some players may, with enough experience, come to play as the equilibrium theory predicts. Recent years have seen research along both lines. Among the most promising developments are the notion of quantal response equilibrium introduced by Richard McKelvey and Thomas Palfrey (1995), the theory of level-n thinking introduced by Dale Stahl and Paul Wilson (1994), and the idea of reinforcement learning developed by Ido Erev and Alvin Roth (1998).

Mark Walker
John Wooders

References

Brown, J. and R. Rosenthal 1990. Testing the minimax hypothesis: a re-examination of O'Neill's game experiment. *Econometrica* 58, 1065-1081.

Erev, I. and A.E. Roth 1998. Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* 88, 848-881.

von Neumann, J. (1928). Zur Theorie der Gesellschaftsspiele, *Mathematische Annalen* 100, 295-320.

McKelvey, R. and T. Palfrey 1998. Quantal response equilibria for extensive-form games. *Experimental Economics* 1, 9-41.

Nash, J. F. 1950. Equilibrium points in N person games. *Proceedings of the National Academy of Sciences* 36, 48-49.

O'Neill, B. 1987. Nonmetric test of the minimax theory of two-person zerosum games. *Proceedings of the National Academy of Sciences* 84, 2106-2109.

Stahl, D. and P. Wilson (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization* 25, 309-327.

Walker, M. and J. Wooders 2001. Minimax play at Wimbledon. *American Economic Review* 91, 1521-1538.

Figure 1

Player B

Player A		H	T
	H	-1	1
	T	1	-1

Figure 2

Goaltender

Kicker		L	R
	L	.4	.9
	R	.8	.3