CHAPTER 1

Action Space

1. Economic theory studies decision problems and strategic interactions, with the objective of understanding and/or predicting the behavior of agents involved in these situations.

Modeling a decision problem or a strategic interaction begins by specifying a set of *actions*, which agents may choose from, and a *payoff structure*, which is a formal description of how actions translate into payoffs for each agent. There exists an extraordinarily vast array of decision problems or strategic situations because, in principle, there is no *a priori* limit on the space of possible actions available, nor limits on the possible mappings from actions to payoffs.

Analysts have considerable freedom in choosing the action space and the payoff structure when constructing models, and a great achievement of game theory has been to identify, within that vast array of situations, simple situations that provide insight into important real-world problems, and for which behavior can be described or characterized. One such example is the Prisoner's Dilemma.

2. The prisoner's dilemma is a two-player game in which each player has only two actions (i.e., a two-by-two game), confess (C) or not confess (N), with the property that confessing is a better option for each individual whatever choice the other makes, and where *both not confessing* is a better outcome than *both confessing*.

The following matrix describes payoffs that are consistent with the properties described above:[1]

	C	N
C	$1,1$	$4,0$
N	$0,4$	$3,3$

Given these payoffs, each player individually finds that confessing is a better strategy, because $1 > 0$ and $4 > 3$. Confessing is said to be a dominant

[1] The matrix indicates that if player 1 chooses N and player 2 chooses C, player 1 gets 0 and player 2 gets 4.

strategy. The prediction is that both confess despite *both not confessing* being a better joint outcome, thereby reflecting the conflict between private objectives (confessing is individually better) and social objectives (both not confessing is jointly better).

3. A virtue of the formal description above is that it can be used across disparate applications: beyond the dilemma that prisoners might face, there are many interactions that fit naturally in the two-by-two game above, with “not confessing” characterizing a *cooperative* strategy, and “confessing” characterizing a *selfish* strategy. These broadly defined labels (cooperative and selfish) may capture different behaviors depending on context. But this is precisely why the model is useful, making it easily applicable across a large variety of situations.

4. Another virtue is that the prediction holds not just for a single specification of the payoff parameters, but for a large range of values, as long as the dominance relations hold. In particular, the agents need not have precise knowledge of the parameters of the model for the prediction to hold.

We emphasize the latter point, as this is central to the critique of the literature that we address in this book. In writing down a payoff structure, we take an outsider’s perspective, defining what each player gets as a function of the pair of actions played. In solving the game, we derive “optimal” choices for each player as though they knew the payoff structure. As analysts, we avoid complicating the model further; we avoid being precise about what each player actually knows about the payoff structure.[2]

The reason for doing so is parsimony. A model is meant to be an analyst’s tool, a parsimonious way to represent reality, helping to explain economic insights which seem relevant. For the sake of parsimony, we generally solve the model as if the agents had precise knowledge of its parameters, hoping (without formally verifying) that the insights drawn from the model do not hinge on this questionable assumption. The prisoner’s dilemma safely passes this test.

5. The restriction to two actions (“not confess,” the cooperative strategy, and “confess,” the selfish strategy) provides a parsimonious model of the conflict between private and social objectives. There are many contexts, however, in which one could imagine varying degrees of cooperative behavior, and where the restriction to two actions could be viewed as unrealistic. In an attempt to assess the strength of the forces away from “full” or efficient cooperation, one may want to enrich the model with multiple levels of cooperation.

But there is a tension: while the restriction to two actions enables the analyst to capture the basic strategic effect, further quantification of this effect

[2] Analysts sometimes take a different view, assuming that payoffs are precisely known, and known to be known, etc. We discuss this alternative view at the end of this chapter.

through a more "realistic" action space is subject to the criticism that the solution implicitly assumes that the agents of the model have substantial knowledge of the structure of the model. The next example illustrates that tension.

6. *A partnership game.* Consider a standard partnership problem in which each of two agents $i = 1,2$ picks an effort level e_i where the effort level can be any non-negative real number. Agent i's gain from the pair (e_1, e_2) is defined as $g_i^\gamma(e_1, e_2)$, with:

$$g_i^\gamma(e_1, e_2) = \gamma\sqrt{e_1 e_2} - e_i^2.$$

An *equilibrium* outcome in this model is a pair of actions (e_1^*, e_2^*) from which neither player wants to deviate unilaterally.[3] The equilibrium effort levels satisfy $e_i^* = \frac{1}{4}\gamma$, while the socially efficient levels would satisfy $e_i^{**} = \frac{1}{2}\gamma$.[4] The model allows one to quantify the effect of both agents following private objectives, each ignoring the positive externality on the other and the higher welfare that would result from a marginal increase in effort.

Despite being possibly more realistic in terms of the strategy space, the model implicitly makes implausible cognitive assumptions: the equilibrium outcome relies on agents behaving as though they knew the details of the model (e.g., the functional forms associated with gain and cost functions for each player) or as if they had learned which effort level was optimal among all possible levels.

7. *Coming to play an equilibrium.* Models are typically silent about how players come to play according to the equilibrium strategies the analyst identifies. One natural hypothesis is that equilibrium is the outcome of a learning process, the stable point from which individual experiments with alternative strategies are unprofitable. In this view, the cognitive assumption is not that players know the parameters of the model, but that they have learned which of their available strategies is best.

There are at least two dimensions that make learning difficult: the number of alternatives to be compared, and changes in the environment. The plausibility of the implicit cognitive assumption may therefore differ a great deal across models. In the prisoner's dilemma, the agent need only compare two actions, and best responses are unaffected by changes in the underlying payoff structure as long as the dominance relations continue to hold. In the partnership game, the cognitive assumption is stronger: the agent must compare many effort

[3] See the note at the end of this chapter for some history and motivation.

[4] Formally, an equilibrium is a pair (e_i^*, e_j^*) such that, for each i, the gain $g_i^\gamma(e_i, e_j^*)$ is maximum at e_i^*. The social optimum is a pair (e_i^{**}, e_j^{**}) such that the total gain $g_i^\gamma(e_i, e_j) + g_j^\gamma(e_i, e_j)$ is maximum at (e_i^{**}, e_j^{**}).

levels, and changes in the underlying model parameters result in changes in best responses.

This discussion relates to the well-known tradeoff between exploration and exploitation. In a changing environment, the sum of exploration and exploitation times is bounded, and finding the action best suited to the current environment is more difficult when the number of actions considered increases.

For the sake of parsimony, equilibrium analysis typically ignores these considerations. A prediction is obtained for each model specification, and as the specification varies (e.g., as γ varies), one predicts a different equilibrium outcome ($e^*(\gamma)$). In so doing, however, we run the risk of taking these predictions too seriously, forgetting that the relationship $e^*(\gamma)$ relies on agents quickly adjusting correctly and costlessly to variations in the payoff structure.

For example, in a changing environment (exhibiting somewhat persistent variations in γ), players might find it easier to track variations in γ by comparing only two effort levels, say $e_i \in \{1,4\}$, with the consequence that, in effect, the game actually played by agents is better described by a two-by-two game. Then, for example, when $\gamma = 6$, the game can be summarized by the matrix:

	$e_2 = 1$	$e_2 = 4$
$e_1 = 1$	$5,5$	$11,-4$
$e_1 = 4$	$-4,11$	$8,8$

and it has the structure of the prisoner's dilemma.

8. *Equilibrium as a shortcut.* Said differently, the payoff functions that we define are a modeling convenience, as is the static formulation that we adopt. These assumptions allow us to bypass the complex issues associated with learning and the possibility that the situations (hence the payoff functions) that agents face vary without their being precisely aware of these underlying variations. We solve for equilibrium as though agents knew γ precisely, and we obtain a prediction for each model specification (γ).

Thus, while the addition of strategies may make the model seem more realistic, this gives rise to predictions that are more finely tuned to the exact model specification assumed, ignoring the possibility that this addition actually diminishes agents' ability to tailor behavior to the underlying payoff structure.

In specifying payoff functions, we implicitly assume that there is a well-defined underlying payoff structure, and that this structure has some stability or permanence which enables some form of learning and some behavioral adjustment to that structure. With a richer underlying parameter space, however, the presumption that all parameters have permanence is less compelling.

9. The traditional answer to the concerns expressed above is that, if we think that agents do not know the structure of the model parameters, we should include in the model a description of what they are ignorant of and how that ignorance is accommodated. Modeling what agents ignore is a challenge that we will address. Let us simply observe for now that this traditional line of thought pushes the difficulty one step further. It produces models that lie in a richer parameter space (this richer parameter may be, for example, a distribution over possible realizations of γ), and strategies that also lie in a richer set, as we typically allow agents to condition behavior on whatever signal they get that might be correlated with γ.

10. Finally, we observe that the point we raise is not specific to games, but applies to decision problems as well. Setting aside questions of convergence to equilibrium, we make the simple observation that finding

$$a^*(s) = \arg\max_{a \in A} u(a, s)$$

is cognitively less demanding when A is a smaller set. Also, a richer action space comes with a richer parameter space (as all payoffs need to be specified), and the stability or permanence of these parameters over time becomes a strong assumption.

> **To summarize**: Models typically endow agents with the ability to ascertain which action is best without questioning how this is achieved. One possibility is that the agent knows the structure of the model itself; another is that, having faced related situations in the past, he has come to know which alternative is best. Whichever one finds more convincing, the agent is assumed to behave as if he knew how to compare the alternatives available, or as if model parameters had enough permanence to make learning plausible.
>
> One consequence is that there is a tension: one may add strategies to make the model descriptively more realistic, but this addition imposes a greater cognitive demand, or permanence of a larger number of model parameters, hence a possibly less realistic model.

Further Comments

Nash Equilibrium. *The equilibrium concept used throughout the book is called a Nash equilibrium, named after John Nash (1950). An earlier version of this idea was formulated by Antoine Augustin Cournot (1838) in his analysis of competition between two firms. With two players for example, it defines a pair of actions* (a_1^*, a_2^*) *from which neither player finds it attractive to change his behavior unilaterally.*

A common justification for equilibrium is that it is a plausible outcome of a learning process which would have converged. Once (a_1^*, a_2^*) *is played, experience cannot provide a player with incentives to change his behavior. It is a stable outcome.*

A second common justification is that players reach equilibrium play through introspection, rather than learning. This type of justification puts a heavy burden on the agents knowing in detail the structure of the game they are playing, and thinking about all the consequences of these details. For this reason, analysts following this path often start with the precautionary statement that the model is common knowledge among agents: each agent knows the model, knows that others know, etc.

Our view is that the learning interpretation is more plausible for many economic problems that employ Nash equilibria.

***Exploration, exploitation and the "considered" set**. The idea of the tradeoff between exploration and exploitation dates back to Thompson (1933), whose motivation came from clinical trials – when different treatments are available for a certain disease and one must decide which treatment to use on the next patient. In a seminal paper, Simon (1955) argued that a key aspect of decision making is the subset of actions that agents actually consider (out of those* a priori *available), and that this subset depends on the extent of exploration. In his view, exploration is driven by the speed with which aspirations decline or increase after a bad or good experience, with unmatched aspirations leading to exploration of new alternatives.*

***On the aim of modeling**. Osborne (2002, page 2) writes: "Game-theoretic modeling starts with an idea related to some aspect of the interaction of decision-makers. We express this idea precisely in a model, incorporating features of the situation that appear to be relevant. This step is an art. We wish to put enough ingredients into the model to obtain nontrivial insights, but not so many that we are led into irrelevant complications; we wish to lay bare the underlying structure of the situation as opposed to describe its every detail."*

References

Cournot, A. A. (1838). *Recherches sur les Principes Mathematiques de la Theorie des Richesses*, New York: L. Hachette.

Nash, J. (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1), 48–49.

Osborne, M. J. (2002). *An Introduction to Game Theory,* Oxford: Oxford University Press.

Simon, H. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118.

Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Bulletin of the American Mathematics Society*, 25, 285–294.