Economics 109, Game Theory, Spring 2002, Vincent Crawford

A *game* is a multi-person decision situation, defined by its players, its "rules" (the order of players' decisions, their feasible decisions at each point, and the information they have when making them); how players' decisions determine the outcome; and players' preferences over outcomes. I call these things the game's *structure*.

Analyses of games must confront all the issues that arise with individual decisions, plus one that is unique to games: Because the outcome is influenced by other players' decisions as well as your own, to do well in a game you need to predict others' decisions, taking their incentives into account. This may require a mental model of other players (including a model of their models of you).

Examples to illustrate issues a theory of games should address:

	L	**R**
T	**2** **2**	**1** **2**
B	**2** **1**	**1** **1**

Crusoe vs Crusoe

	Confess	**Don't**
Con-fess	**-5** **-5**	**-10** **-1**
Don't	**-1** **-10**	**-2** **-2**

Prisoner's Dilemma

Crusoe vs. Crusoe is just two decision problems, not really a game; each player has a best decision independent of other's (*dominant*).

In Prisoner's Dilemma players' decisions affect each other's payoffs but each still has a dominant decision; individual optimal decisions with payoff interactions yield a Pareto-inefficient outcome. Don't need special theory for games like this (but Prisoner's Dilemma gets more interesting when we study ways to improve its outcome).

	Push	Wait
Push	1 5	5 3
Wait	-1 9	0 0

Pigs in a Box

	Heads	Tails
Heads	-1 1	1 -1
Tails	1 -1	-1 1

Matching Pennies

In Pigs in a Box, Row (R) is a big pig and Column (C) is a little pig. The box is a Skinner box, named for psychologist B.F. Skinner. Pushing a lever at one end yields 10 units of grain at the other. Pushing "costs" either pig 2 units of grain. If R pushes while C waits, C can eat 5 units before R lumbers down and shoves C aside. If C pushes while R waits, C cannot shove R aside and R gets all but one unit of grain. If both push, and then arrive at the grain together, C gets 3 units and R gets 7. If both wait, both get 0.

In experiments with real pigs, if they settle down, it tends to be at (R Push, C Wait). C does better, although R can do anything C can do! This couldn't happen in an individual decision problem. It happens here because Wait dominates Push for C, but not for R: the way they interact in this game, only R has an incentive to Push.

In games, (the right kind of) weakness can be an advantage! R might do better if he could commit himself to giving C more grain if C Pushed. Understanding this should help to understand many things in economics. E.g. corporations as legal "persons" have the right to be sued. This is a "right" because it helps enforce contracts.

(If the pigs had studied game theory, they wouldn't have to "settle down": they could just figure out that they should play (R Push, C Wait). That they got there anyway suggests that learning and rationality arguments yield the same conclusions in the long run.)

Matching Pennies has no good *pure* decisions (often called *strategies*), but a unique good *mixed* (randomized) strategy. How would you play? What if the 1 (–1) for (Heads, Heads) were 2 (–2)?

	L	M	R
T	0 7	5 0	3 0
M	0 5	2 2	0 5
B	7 0	5 0	3 7

Dominance-solvable Game

This 3x3 game has a more complex pattern of *iterated* dominance.

	L	C	R
T	10 4	0 3	3 1
B	0 0	10 2	3 10

Domination Via Mixed Strategies

This 2x2 game has dominance only if we consider mixed strategies.

	L	M	R
T	0 7	5 0	7 0
M	0 5	2 2	0 5
B	7 0	5 0	0 7

Unique Equilibrium without Dominance

This 3x3 game has a unique *equilibrium* combination of strategies such that each player's is best for him, given the other's; but no dominance. It shows that we will need a way to analyze players' decisions that takes their interdependence fully into account.

	Go	Wait
Go	0, 0	1, 1
Wait	1, 1	0, 0

Alphonse and Gaston

	Fights	Ballet
Fights	2, 1	0, 0
Ballet	0, 0	1, 2

Battle of the Sexes

Alphonse and Gaston's problem is that there are *two* ways to solve their coordination problem…and therefore maybe no good way! Each of the two ways requires them to behave differently when there are no clues to distinguish their roles.

(In the early 1900s Frederick B. Opper created the *Alphonse and Gaston* comic strip, with two excessively polite fellows saying "after you, my dear Gaston" or "…Alphonse" and never getting through the doorway. They are mostly forgotten, but we have Alphonse-Gaston games in dual-control lighting circuits in our homes.)

Alphonse and Gaston

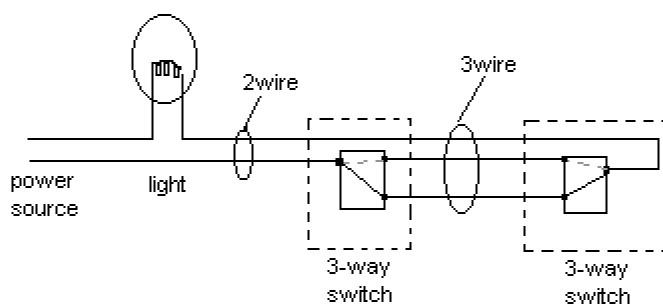

Alphonse and Gaston in your home

Coordination games like Alphonse and Gaston show that players may have problems even if preferences are the same. If economics is "about" coordination, we should study such problems—not just the coordination that happens in competitive markets. (Mixed strategies can help Alphonse and Gaston learn to coordinate if they play the same game over and over; but the mixed strategies serve a completely different purpose than in Matching Pennies.)

Battle of the Sexes complicates Alphonse and Gaston's problem with different preferences about *how* to coordinate (the Hawk-Dove game from evolutionary game theory is Battle of the Sexes with different labels that highlight the problem of breaking symmetry). How would you play Battle of the Sexes once? Repeatedly? Would you play differently if the 2 for Row at (Fights, Fights) were a 3?

Other Player

	Stag	**Rabbit**
Stag	2, 2	0, 1
Rabbit	1, 0	1, 1

Two-Person Stag Hunt

All Other Players

	All-Stag	**Not All-Stag**
Stag	2	0
Rabbit	1	1

n-Person Stag Hunt

In Stag Hunt (Rousseau's story, assembly line, meeting), with two or *n* players, there are two symmetric, Pareto-ranked, pure-strategy equilibria, "all-Stag" and "all-Rabbit". There's also an uninteresting mixed-strategy equilibrium. All-Stag is better for all than all-Rabbit; but Stag is riskier in that unless all others play Stag, a player does better with Rabbit. The game is like a choice between participating in a highly productive but fragile society and autarky, which is less rewarding but safer because less dependent on coordination.

Terminology and key concepts

There are two leading frameworks for analyzing games:

- *Cooperative* game theory assumes rationality, unlimited communication and ability to make agreements. It assumes Pareto-efficiency and sometimes symmetry across players to characterize possible outcomes of rational bargaining, filtering out many details.

- *Noncooperative* game theory also assumes rationality, but replaces the assumptions of unlimited communication and ability to make agreements with a detailed model of the situation. It uses rationality, augmented by the notion of equilibrium seen above, to explain outcomes (sometimes including cooperation). Need a clear distinction between behavioral assumptions and structure, like separating preferences and feasibility in consumer theory.

Like most applications, I focus on noncooperative game theory. Recall that a *game* is a multi-person decision situation, defined by the game's *structure*: its players, its "rules" (the order of players' decisions, their feasible decisions at each point, and the information they have when making them); how players' decisions determine the outcome; and players' preferences over outcomes. (Handle uncertainty about the outcome by assigning *payoffs* (*vN-M utilities*) to outcomes and assuming players maximize expected payoffs.

Assume numbers of players, decisions, periods are finite; can relax. Assume the game is a complete model of the situation; if not, make it one, e.g. by including decision to participate.

"Game" is a misnomer: frivolous or not, most interactions are games. "Noncooperative" is also a misnomer: Noncooperative game theory spans the entire range of multi-person decision situations from pure conflict (as in *zero-sum* parlor games) to pure coordination. Most applications have both conflict and coordination, and many involve figuring out how to support cooperation.

A *static* or *simultaneous-move game* has one stage, at which players make simultaneous decisions, like those discussed above.

A *dynamic game* has some sequential decisions.
E.g. Ultimatum Contracting with two feasible contracts, X and Y:

R proposes X or Y to C, who must either accept (a) or reject (r).
If C accepts, the proposed contract is enforced.
If C rejects, the outcome is a third alternative, Z.
R prefers Y to X to Z, and C prefers X to Y to Z.
R's payoffs: u(Y)=2, u(X)=1, u(Z)=0; C's: v(X)=2, v(Y)=1, v(Z)=0.

The game actually depends on whether C can observe R's proposal before deciding whether to accept. With observable proposal it's dynamic; with unobservable proposal it's static.

Can represent either game by its *extensive form* or *game tree*, which shows its sequence of decisions, outcomes, and payoffs. Order of *decision nodes* must respect timing of moves. Each node belongs to an *information set* (represented by circles), the nodes the player whose decision it is cannot distinguish (and at which he must therefore make the same decision). All such nodes must belong to same player and have same feasible decisions.

(A game of *perfect information* is one in which a player making a decision can always observe all previous decisions, so every information set contains one decision node, as in Ultimatum Contracting with Observable Proposal.)

For dynamic games it is important to distinguish *strategies* from *decisions* or *actions*. A *strategy* is a complete contingent plan that specifies a decision for each of a player's decision nodes and information sets (like a chess textbook, *not* a move in chess). In a static game a strategy reduces to a *decision* or *action*. (These definitions apply equally well to *mixed* or *pure* strategies.)

Specifying a strategy profile—one for each player—must determine an outcome (or probability distribution over outcomes).

A player's strategy (or decision) must be feasible independent of others' strategies; e.g. "wrestle" is not a well-defined strategy.

Players *must* be thought of as choosing strategies *simultaneously* (without observing others' strategies) at the start of play. Rational, perfect foresight implies that simultaneous choice of strategies yields the same outcome as decision-making in "real time" (this is a testable prediction, which can fail, and does for some real people).

We need complete contingent plans (even for nodes ruled out by prior decisions) to evaluate consequences of alternative strategies, to formalize the idea that the predicted strategy choice is optimal. (0-probability events are endogenously determined by decisions, and so cannot be ignored here as they are in individual decisions.)

With the concept of strategy, we can also represent a game, static or dynamic, by the relationship between its strategy profiles and payoffs: *normal form*, *payoff function*, or (if 2 people) *payoff matrix*.

	a	**r**
X	2 1	0 0
Y	1 2	0 0

Ultimatum Contracting with Unobservable Proposal

	a, a	**a, r**	**r, a**	**r, r**
X	2 1	2 1	0 0	0 0
Y	1 2	0 0	1 2	0 0

Ultimatum Contracting with Observable Proposal

In Ultimatum Contracting, whether or not C can observe R's proposal, R has two pure strategies, "(propose) X" and "Y."
If C cannot observe R's proposal, C has two pure strategies, "a(ccept)" and "r(eject)". If C can, C has four pure strategies, "a (if X proposed), a (if Y proposed)", "a, r", "r, a", and "r, r."

C's additional information in Ultimatum Contracting with Observable Proposal "shows up" only in the form of extra strategies for C. But when the players are rational, this can affect the outcome.

Suppose the payoffs are as above (R's: u(Y)=2, u(X)=1, u(Z)=0; C's: v(X)=2, v(Y)=1, v(Z)=0). Then C prefers either X or Y to Z, so C will accept either X or Y whether or not C can observe R's proposal. R will then propose Y, his favorite contract, and C will accept.

Now suppose C's payoffs are changed to: v(X)=2, v(Y)= 0, v(Z)=1, so that C now prefers X to Z, but not Y to Z (R's are unchanged). If C can observe R's proposal, C will accept X but not Y. R will then propose X, which he prefers to Z, and C will accept. But if C cannot observe R's proposal, C must accept or reject what R proposes without regard to what it is. If C accepted, R would propose Y, which is worse for C than Z, so C will reject whatever R proposes. We will see below how to do this kind of analysis more generally.

Behavioral assumptions

Something is *mutual knowledge* if all players know it, and *common knowledge* if all know it, all know that all know it, and so on. Focus on problem of predicting others' responses by assuming common knowledge of structure (allows uncertainty with commonly known distributions), as in games of *complete information*.

Assume that players are *rational*, i.e. that they maximize expected *payoffs* given *beliefs* (subjective probability distributions) about other players' decisions or strategies that are not logically inconsistent with anything they know, and that follow Bayes' Rule.

Define *strictly and weakly dominant and dominated strategies* (e.g. Prisoner's Dilemma). Dominance for pure implies dominance for mixed strategies, but can have dominance by mixed without dominance by pure strategies (e.g. above 2x3 game).

Define *iterated elimination* or *deletion* of strictly dominated strategies (*iterated dominance*; e.g. Pigs in a Box). Order-independent, but this doesn't work for iterated *weak* dominance.

Define *dominance*-solvability (e.g. first 3x3 game above).

Rationalizability

Define *rationalizable* strategies: survive iterated elimination of never *(weak) best responses* (same as iterated *strict* dominance in two-person games, but not quite in general; order-independent)

For the original payoffs in Ultimatum Contracting, a rational C will accept whether or not C R's proposal is observable, whatever R proposes. A rational R will therefore propose Y. Y and a (or (a,a) if the proposal is observable) are the only rationalizable strategies. The set of rationalizable strategies can't be larger than the set that survive iterated strict (but not weak) dominance, because strictly dominated strategies can never be weak best responses.

In two-person games the two sets are the same because never-weak-best-response strategies are exactly those that are strictly dominated. E.g. *any* combination of strategies in Matching Pennies, Battle of the Sexes, or the second 3x3 game is rationalizable.

Rationalizability has exactly the implications of common knowledge of the structure and players' rationality (latter is key assumption).

Theorem: Common knowledge of the structure and players' rationality implies that players will choose rationalizable strategies, and any profile of rationalizable strategies is consistent with common knowledge of the structure and rationality.

Proof: Illustrate in first direction for 3x3 games, in second by building "towers" of beliefs to support outcomes in 3x3 games.

The number of levels of iterated knowledge of rationality needed is just the number of rounds of iterated dominance; need full *common* knowledge only for indefinitely "large" games.

E.g. Guessing Game: *n* players simultaneously *guess* from 0 to 100. Whoever's guess is closest to 70% of the average wins $20.

Questions: (i) What is your Nash equilibrium guess?
(ii) What would you guess in this group (excluding me)?

Answers: (i) "All-0" is the unique rationalizable outcome (infinite iterated dominance and common knowledge of rationality.)
(ii) Nobody (not even I) would guess 0, so it would be stupid for you to guess 0. In this group, 25 is more like it. (Answer varies some, but not much, with group's size and social or other characteristics.)

Nash Equilibrium

Most economic games have many rationalizable outcomes (e.g. coordination games, second 3x3 game). So players must base decisions on beliefs not given by common knowledge of rationality.

Game theory assumes players' strategies are in *Nash equilibrium*, so each player's maximizes expected payoff, given the others'. (Can generalize to *equilibrium in beliefs*, useful describing equilibria in populations randomly matched to play game; beliefs determine players' best-response correspondences, all the theory predicts.)

Nash equilibrium is a kind of "rational expectations" equilibrium. If all players expect the same strategies and choose best responses, their beliefs are confirmed iff they are in equilibrium (differs in that individual decisions are predicted, and players' beliefs interact).

Any equilibrium strategy is rationalizable (why?). In dominance-solvable games, equilibrium is equivalent to rationalizability (why?).

In non-dominance-solvable games, equilibrium is consistent with rationalizability, but it also requires that players' strategies are best responses to *correct* beliefs about others' strategies (which must then be the same for all)—not just some beliefs consistent with common knowledge of rationality (3x3 game). Equilibrium is surely at least a *necessary* condition for a rational prediction about behavior, but why should players have correct beliefs?

- Traditional rationale: Players form correct (self-fulfilling) beliefs about each other's strategies when they first play a game, and so (if rational) play equilibrium immediately.
- Adaptive rationale: Players (like the pigs in the above example) learn to predict others' strategies in repeated play of analogous games, and so (if rational) converge to equilibrium over time.

Traditional game theory focuses on rationality-based reasoning, while adaptive learning models make assumptions directly about how players adjust strategies over time. Both agree that possible limiting outcomes are equilibria (in the game that is repeated, but maybe not in the game that describes the entire process).

Applications (outline only)

Chapter 4: Continuous pure strategies, best-response curves,
Stackelberg, Cournot, and Bertrand duopoly
Chapter 5: Mixed-strategy equilibria in 0- and non-0-sum games
Chapter 8: Prisoner's Dilemma, Hofstadter/Axelrod and tit-for-tat
Chapter 9: Strategic moves, Schelling
Chapter 10: Evolutionary games, Hawk-Dove, sex ratio, Stag Hunt
Chapter 11: Collective-action games, Schelling (optional)
Chapter 12: Uncertainty and information, signaling, screening, lying
Chapter 13: Brinkmanship (optional)
Chapter 14: Strategy and voting (optional)
Chapter 15: Bidding and auctions (optional)
Chapter 16: Bargaining, ultimatum and alternating-offers, Nash's
demand game