Lecture Slides: Game Theory
Vincent P. Crawford

University of Oxford
Michaelmas Term 2013

Outline

1. Taxonomy of games: cooperative and noncooperative

2. Describing noncooperative games and strategic behavior: rationality, dominance, iterated dominance, and Nash equilibrium

3. Game experiments: guessing and coordination games

4. Repeated games: supporting cooperation via credible threats

5. Examples of cooperative games: marriage, college admissions, and matching markets (slides excerpted from Jonathan Levin's)

1. Taxonomy of games: cooperative and noncooperative

A game is a multi-person decision situation, in which a person's outcome is influenced by other people's decisions as well as his own.

Thus almost all economic interactions are games.

There are two leading frameworks for analyzing games: *cooperative* and *noncooperative*.

- *Cooperative* game theory assumes rational strategic behavior, unlimited communication, and unlimited ability to make agreements.

- It sidesteps the details of the structure by assuming that players reach a Pareto-efficient agreement, which is sometimes further restricted, e.g. by requiring symmetry of utility outcomes for symmetrically situated players.

- Its goal is to characterize the limits of the set of possible agreements that might emerge from rational (possibly implicit) bargaining.

● *Noncooperative* game theory also assumes strategic rationality.

But by contrast:

● Noncooperative game theory replaces cooperative game theory's assumptions of unlimited communication and ability to make agreements with a detailed model of the situation and of how rational players will behave in it.

● Its goal is to use rationality, augmented by the "rational expectations" notion of Nash equilibrium, to predict or explain outcomes in a situation.

This course focuses on noncooperative game theory, which dominates applications.

But cooperative game theory is better suited to some applications, e.g. where the structure of the game is unclear or unobservable, and it is desired to make predictions that are robust to it.

I will give a whirlwind tour of an example of cooperative game theory at the end.

Like the term “game” itself, the term “noncooperative” is a misnomer:

- Noncooperative game theory spans the entire range of *multi-person* or *interactive decision situations*.

- Although *zero-sum games*, whose players have perfectly conflicting goals, played a leading role in the development of the theory, most applications combine elements of conflict with elements of coordination, and some involve predicting which settings are better for fostering cooperation.

- This is done by making behavioral assumptions at the individual level (“methodological individualism”), thereby requiring cooperation to emerge (if it does) as the outcome of explicitly modeled, independent decisions by individuals in response to explicitly modeled institutions.

- By contrast, cooperative game theory makes the group-level assumption that the outcome will be Pareto-efficient, and (with some important exceptions) downplays the incentive and coordination issues that are the focus of noncooperative analyses of cooperation.

In these lectures, we will first describe the structure of a noncooperative game.

We will then introduce assumptions about strategic behavior, gradually refining the notion of what it means to make rational strategic decisions.

In the process we will show how game theory can elucidate economic questions.

As you learn to describe the structure, bear in mind that the goal is to provide enough information about the game to formalize the idea of a rational decision.

(This may foster patience about not yet knowing what it means to be rational.)

From the noncooperative point of view, a *game* is a multi-person decision situation defined by its *structure*, which includes:

- The *players*, independent decision makers (e.g. bridge has first four players, then three, not two, even though partners have the same goals)

- The *rules*, which specify the order of players' decisions, their feasible decisions at each decision point, and their information at each decision point.

- How players' decisions jointly determine the physical *outcome*

- Players' preferences over outcomes (or probability distributions of outcomes)

Players' preferences over outcomes are modeled just as in decision theory.

Preferences can be extended to handle shared uncertainty about how decisions determine the outcome just as in decision theory: by assigning von Neumann-Morgenstern utilities, or *payoffs*, to outcomes and assuming that players maximize their expected payoffs.

Something is *mutual knowledge* if all players know it.

Something is *common knowledge* if all players know it, all know that all know it, and so on.

Assume here, for simplicity, that the structure of the game is *common knowledge*, except possibly for shared uncertainty about how decisions determine the outcome, with the probability distributions common knowledge.

(Later parts of the course will relax the assumption that the structure is common knowledge, developing ways to model asymmetric information.)

I also assume, for simplicity, that players make single, simultaneous decisions.

Simultaneous decisions need not be synchronous, but they must be *strategically simultaneous* in that players cannot observe each other's decisions in time to react.

(Later parts of the course will relax this too.)

Aside on games with sequential decisions

Noncooperative game theory's methods for analyzing rational decisions in games in which players make simultaneous decisions can (and will, later) be extended to games in which some decisions are sequential, and reactions are possible.

Generalize the notion of a decision to a *decision rule* or *strategy*, a complete contingent plan for playing the game that allows a player's decisions to respond to others' decisions when he can observe them before making his own decisions.

Players must be thought of as choosing their strategies strategically simultaneously (without observing others' strategies) at the start of play.

Rational, perfect foresight (strong assumption!) implies that simultaneous choice of such strategies yields the same outcome as decision-making in real time.

Complete contingent plans are needed (even for decision points ruled out by prior decisions) to evaluate consequences of alternative strategies and formalize the idea that the predicted strategy choice is optimal. (Zero-probability events are endogenously determined by players' decisions, so cannot be ignored in games.)

End of aside

Return to games in which players make single, simultaneous decisions.

● In these and other games, a player's decisions must be feasible *independent* of others' decisions; e.g. "wrestle with player 2" is not a well-defined decision, although "try to wrestle with player 2" could be, if what happens if 2 doesn't also try is clearly specified.

● Specifying all of each player's decisions must completely determine an outcome (or at least a probability distribution over outcomes).

If a game model does not pass these tests, it must be modified until it does.

E.g. If you object to my game analysis on the grounds that players don't really have to play "my" game, my (only!) remedy is to add to my game's rules a player's decision whether to participate, and then to insist that that decision be explained by the same behavioral assumptions as players' other decisions.

2. Describing noncooperative games and strategic behavior: rationality, dominance, iterated dominance, and Nash equilibrium

These examples will show how to describe games with simultaneous decisions in “normal” or “payoff-function” form, how to describe strategic behavior, and introduce the issues a theory of strategic behavior should address.

(Later in the course you will learn how to use the normal form and the alternative, *extensive form*, to model games with some sequential decisions, using the notion of strategy mentioned above.)

Crusoe versus Crusoe is not really a game, just two independent decision problems; and we don’t need any theory to predict that rational players will choose (T, L).

	L	R
T	2, 2	2, 1
B	1, 2	1, 1

Crusoe vs Crusoe

Prisoner's Dilemma *is* a game, because players' decisions affect each other's payoffs; but we still don't need a new theory to predict that rational players will choose (Confess, Confess). ("Confess" = "Defect"; "Don't" = "Cooperate".)

(The "'s" in Prisoner's Dilemma (not "s'") signals methodological individualism.)

	Don't	**Confess**
Don't	3, 3	0, 5
Confess	5, 0	1, 1

Prisoner's Dilemma

- A *strictly dominant* decision is a decision that yields a player strictly higher payoff, no matter which decision(s) the other player(s) choose.

  E.g. T for Row or L for Column in Crusoe vs Crusoe, or Confess for either player in Prisoner's Dilemma.

	Don't	**Confess**
Don't	**3** **3**	**5** **0**
Confess	**0** **5**	**1** **1**

Prisoner's Dilemma

A rational player must choose a strictly dominant decision if he has one.

- A *strictly dominated* decision is a decision that yields a player strictly lower payoff than another feasible decision, no matter which decisions the others choose.

  E.g. B for Row or R for Column in Crusoe vs Crusoe, or Don't in Prisoner's Dilemma.

	Don't	**Confess**
Don't	3 3	5 0
Confess	0 5	1 1

Prisoner's Dilemma

A rational player will never play a strictly dominated decision, because there are no beliefs about other players' decisions that make it a best response.

Although it doesn't happen in Crusoe vs Crusoe or Prisoner's Dilemma, there can be dominated decisions without a dominant decision, which makes the notion of dominated decision more useful than the notion of dominant decision.

Because of the way the prisoners' payoffs interact, individually rational decisions yield a collectively suboptimal (Pareto-inefficient, in the prisoners' view) outcome.

Note that what's Pareto-efficient for the prisoners need not be good for society.

	Don't	**Confess**
Don't	3 3	5 0
Confess	0 5	1 1

Prisoner's Dilemma

Prisoner's Dilemma highlights a flaw in libertarianism: an enforceable law against confessing would make both prisoners better off, while limiting their freedom.

(The grain of truth in libertarianism is that it would be beneficial to be the *only one* allowed to break the law.

But that can't yield a universal rule by which to organize society, "universal" as in: "Act only according to that maxim whereby you can, at the same time, will that it should become a universal law."—Kant)

Prisoner’s Dilemma’s tension between individual rationality and Pareto-efficiency makes it the simplest possible model of incentive problems, which makes it a popular platform for the analyses of institutions that overcome such problems.

The positive flip side of my caveat about modeling situations as games, e.g.:

“If you object to my game analysis on the grounds that players don’t really have to play ‘my’ game, my (only!) remedy is to add to my game’s rules a player’s decision whether to participate, and then to insist that that decision be explained by the same behavioural assumptions as players’ other decisions.”

This insistence is an important constraint on analysis: otherwise there is nothing to pin down the assumptions implicit in speculative “solutions” to problems.

Later in these lectures we will see examples of how repeated interaction can sometimes support cooperation despite incentive problems in the short run.

Yet a Prisoner's Dilemma model of incentive problems is too simple: it ignores the difficulty of coordination and conflicts between different ways to cooperate.

In Pigs in a Box, Row (R) is a big (“dominant”) pig and Column (C) a little (“subordinate”) pig. The box is a “Skinner box”, named for B.F. Skinner.

- Pushing a lever at one end of the box 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 comes 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.

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

Pigs in a Box

Here rational strategic behavior is more subtle, in that for the first time, it requires at least one player to predict the other's response to the game.

Its consequences are also a bit surprising:

● R can do anything C can do, which in an individual decision problem would ensure that R does better.

● But in the lab pigs tend to settle down at (R Push, C Wait): C does better!

● In games, evidently, (the right kind of) weakness might be an advantage.

To see why, and to see which kinds of weakness might be advantageous, we need more theory.

Recall that the structure is assumed for simplicity, to be *common knowledge*.

- A player is *rational* (in the decision-theoretic sense) if he maximizes his payoff given *beliefs* (subjective probability distributions) about other players' decisions that are not inconsistent with anything he knows.

- A player who faces uncertainty about the consequences of his decisions is *rational* if he maximizes his expected payoff (*vN-M utility*).

A first guess at how to formalize the idea of rational decisions in games is that assuming that players are rational will suffice for a useful theory of behavior.

That guess is correct for games like Crusoe v. Crusoe and Prisoner's Dilemma.

But that guess fails badly in slightly more complex games, such as Pigs in a Box.

A second guess is that assuming that players are rational *and* that that fact is common or at least mutual knowledge is enough to yield a useful theory.

That guess works in some games, in which common knowledge of rationality yields a unique prediction.

But even that guess fails badly in many economically interesting games.

A third guess is that assuming that players are rational and that players' decisions are best responses to *correct* beliefs about others' decisions (which must then be the same for all players) is enough to yield a useful theory.

That guess leads to the idea of Nash equilibrium, which is the leading behavioral assumption in noncooperative game theory; but even it has some drawbacks.

To see how common or mutual knowledge of rationality works, imagine that the pigs are as good at reasoning about others' responses to incentives as (some) humans seem to be.

They can then use rationality and knowledge of others' rationality—in this case mutual knowledge is enough—to figure out they should play (R Push, C Wait).

If they have mutual knowledge of rationality, the reasoning goes as follows:

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

Pigs in a Box

- No rational C will choose Push.

- Therefore no rational R who knows that C is rational will play Wait.

- (R Push, C Wait) is the only possible outcome.

In economic terms, because Wait strictly dominates Push for C, only R has an incentive to Push.

This incentive effect is what turns R's greater strength into a weakness.

R *might* do better if he can change the game in a way that gives C an incentive to Push, at least some of the time; e.g. by committing himself to giving C more grain if C Pushed. (There's still a coordination problem; more below.)

Understanding which kinds of games commitments help in, and what kinds of commitments help, should help us to understand the usefulness of contracts and other ways to change how relationships are governed.

(As legal "persons", corporations have the "right" to be sued. This is a "right", not simply a liability, because it may allow mutually beneficial contracts that would not be in the other party's interests if it could not sue for breach.)

Aside on learning

Pigs are probably not really as good as humans at reasoning about others' likely decisions. So why do they still tend to settle down at (R Push, C Wait)?

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

Pigs in a Box

- In repeated play, because Push is strictly dominated for player C, it must do worse on average for C than Wait.

- Thus even a C that reacts unthinkingly to rewards will "learn" to choose Wait with higher and higher probability over time.

- Once the probability that C chooses Wait is high enough (> 4/7), player R will learn to choose Push with higher and higher probability over time.

- They will eventually settle down at (R Push, C Wait): Learning yields the same outcome in the limit as rationality-based reasoning does: a general result.

End of aside

We can characterize the implications of common knowledge of rationality via

- *Iterated deletion* of strictly dominated decisions (often called “iterated strict dominance”): eliminating strictly dominated decisions for one or both players, then eliminating decisions that become strictly dominated once players’ strictly dominated decisions are eliminated, and so on ad infinitum.

If iterated strict dominance reduces the game to a single decision for each player—as in Pigs in a Box eliminating Push for C and then Wait for R reduces the game to (Push, Wait)—the game is said to be *dominance-solvable*.

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

Pigs in a Box

The outcome of iterated strict dominance is independent of the order in which it is done, so in a dominance-solvable game it always leads to the same decisions.

Aside on weak dominance

- A *weakly dominant* decision is a decision that yields a player weakly higher payoff (or expected payoff), no matter which decisions the other players choose.

- A *weakly dominated* decision is a decision that yields a player weakly lower payoff (or expected payoff), no matter which decisions the other players choose.

The outcome of iterated weak dominance is *not* independent of the order in which it is done. E.g. in Give Me a Break, it leads to (T, L), (T, R), or (B, R).

	L	**R**
T	**1**, **1**	**0**, **1**
B	**0**, **0**	**0**, **0**

Give Me a Break

Moreover, a rational player with sharply focused beliefs need not choose a weakly dominant decision, and might choose a weakly dominated decision.
End of aside

Iterated strict dominance is linked to common knowledge of rationality via the notion of *rationalizable* decisions.

- A *rationalizable* decision is one that survives iterated elimination of *never (weak) best responses*, those decisions that are not even tied for being a best response to any beliefs.

The set of rationalizable decisions can't be larger than the set that survive iterated strict dominance, because strictly dominated decisions can never be weak best responses (that is why the notion builds on *strict* dominance).

In a two-person game, a rationalizable decision is exactly one that survives iterated strict dominance; well defined because the latter is order-independent.

In games with more than two players, the two notions are not quite the same because players can have correlated beliefs about others' strategies.

It will be shown later in the course that if the structure and players' rationality are common knowledge, then players must choose rationalizable decisions, and that any profile of rationalizable decisions can be supported by beliefs that are consistent with common knowledge of the structure and of players' rationality.

E.g., M and C are the only rationalizable decisions in (how support them?):

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

Dominance-solvable game

But any decisions are rationalizable in (how support them?):

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

Unique equilibrium but no dominance

Nash Equilibrium

Most economically interesting games have multiple rationalizable outcomes, so players' decisions are not *dictated* by common knowledge of rationality, and the guess that it will yield a useful theory of strategic behavior fails badly.

- To make sharper predictions, noncooperative game theory assumes that players' decisions are in *Nash equilibrium*, that is, that each player's decision maximizes his payoff or expected payoff, given the others' decisions.

Any equilibrium decision is rationalizable (why?).

It can be shown that an equilibrium always exists in non-pathological games.

Therefore in a dominance-solvable game, players' unique rationalizable decisions are in equilibrium (why?).

In non-dominance-solvable games, however, equilibrium also effectively requires that players' decisions are best responses to *correct* beliefs about others' decisions, which must then be the same for all players, e.g.:

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

Unique equilibrium but no dominance

Nash equilibrium is a kind of "rational expectations" equilibrium, in that if players are rational, and all expect the same decisions and best respond to those beliefs, then their beliefs are self-confirming if and only if they are in Nash equilibrium.

This goes far beyond rationality, or even common knowledge of rationality.

Why might players have correct beliefs about each other’s decisions?

There are two possible justifications, which generalize those mentioned in connection with Pig in a Box.

- Thinking: If players are rational and have perfect models of each other’s decisions, strategic thinking will lead them to have the same beliefs immediately, and so play an equilibrium even in their initial responses to a game.

- Learning: Even without perfect models, if players are rational and repeatedly play analogous games, experience will eventually allow them to predict each others’ decisions well enough to play an equilibrium in the game that is repeated.

Aside on mixed strategies
In game theory it is useful to extend the idea of decision, or strategy, from the unrandomized (*pure*) notion to allow randomized (*mixed*) decisions or strategies.

E.g. Matching Pennies has no appealing pure strategies, but there is an appealing way to play using mixed strategies: randomizing 50-50. (Why?)

	Heads	**Tails**
Heads	**1**, **-1**	**-1**, **1**
Tails	**-1**, **1**	**1**, **-1**

Matching Pennies

Our definitions apply to mixed as well as pure strategies, if the uncertainty mixed strategies cause is handled as for other kinds of uncertainty, by assigning payoffs to outcomes so that rational players maximize their expected payoffs.

Mixed strategies ensure that "well-behaved" games always have rational-expectations strategy combinations: i.e. that *Nash equilibria* always exist.
End of aside

Nonuniqueness of Equilibrium and Coordination

	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

In the early 1900s Frederick B. Opper created the comic strip *Alphonse and Gaston*, with two excessively polite fellows saying "after you, my dear Gaston" or "after you, my dear Alphonse" and thus never getting through a doorway. Alphonse and Gaston live on in the dual-control lighting circuits in our homes.

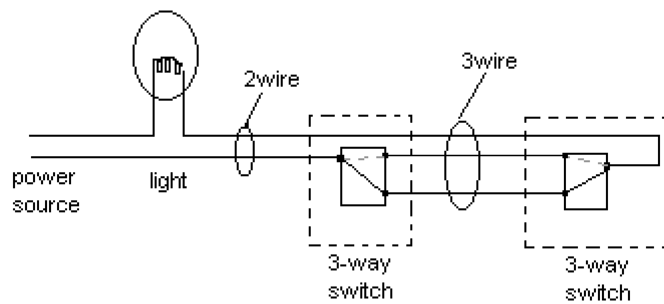

Alphonse and Gaston's problem is that there are *two* good ways to solve their coordination problem…and therefore maybe no good way.

Each way requires them to decide differently, but the setting provides no clue to break the symmetry of their roles.

Battle of the Sexes—the simplest possible bargaining problem—adds to the difficulty of coordination by giving players different preferences about *how* to coordinate, but still no clue about how to break the symmetry.

These games are popular platforms for the analyses of institutions that overcome such problems, e.g. via conventions that use labels to break the symmetry of players' roles, such as "defer to short people" or "defer to women".

In Stag Hunt (Rousseau's story), with two or *n* players, there are two symmetric, Pareto-ranked, pure-strategy equilibria, “all-Stag” and “all-Rabbit”.

2-person Stag Hunt

	Other Player: Stag	**Other Player: Rabbit**
Stag	2, 2	0, 1
Rabbit	1, 0	1, 1

n-Person Stag Hunt

	All Other Players: Stag	**All Other Players: Rabbit**
Stag	**2**	**0**
Rabbit	**1**	**1**

All-Stag is better for all than all-Rabbit: Kant would have no trouble here.

But Stag is riskier in that unless all others play Stag, a player does better with Rabbit.

Stag Hunt is like a choice between autarky and participating in a highly productive but brittle society, which is more rewarding but riskier because productivity depends on perfect coordination.

Stag Hunt is a special case of Larry Summers's Bank Runs example:

"A crude but simple game, related to Douglas Diamond and Philip Dybvig's [1983 *JPE*] celebrated analysis of bank runs, illustrates some of the issues involved here. Imagine that everyone who has invested $10 with me can expect to earn $1, assuming that I stay solvent. Suppose that if I go bankrupt, investors who remain lose their whole $10 investment, but that an investor who withdraws today neither gains nor loses. What would you do? Each individual judgment would presumably depend on one's assessment of my prospects, but this in turn depends on the collective judgment of all of the investors.

Suppose, first, that my foreign reserves, ability to mobilize resources, and economic strength are so limited that if any investor withdraws I will go bankrupt. It would be a Nash equilibrium (indeed, a Pareto-dominant one) for everyone to remain, but (I expect) not an attainable one. Someone would reason that someone else would decide to be cautious and withdraw, or at least that someone would reason that someone would reason that someone would withdraw, and so forth. This…would likely lead to large-scale withdrawals, and I would go bankrupt. It would not be a close-run thing. …Keynes's beauty contest captures a similar idea.

Now suppose that my fundamental situation were such that everyone would be paid off as long as no more than one-third of the investors chose to withdraw. What would you do then? Again, there are multiple equilibria: everyone should stay if everyone else does, and everyone should pull out if everyone else does, but the more favorable equilibria seems much more robust."

—Lawrence Summers, "International Financial Crises: Causes, Prevention, and Cures," (2000 *AER*).

The game Summers describes can be represented by a payoff table as follows:

		Summary statistic	
		In	**Out**
Representative player	**In**	**1**	**-10**
	Out	**0**	**0**

Bank Runs

The summary statistic is a measure of whether or not the required number of investors stays In.

		Summary statistic	
		In	Out
Representative player	In	1	-10
	Out	0	0

Bank Runs

In Summers's first example, all investors must stay In to prevent the bank from collapsing, so the summary statistic takes the value In if and only if all but the representative player stay In.

In Summers's second example, two-thirds of the investors need to stay In, so the summary statistic takes the value In if and only if that is the case, adding in the representative player.

In each example there are two pure-strategy equilibria: "all-In" and "all-Out".

(There is also a behaviorally implausible mixed-strategy equilibrium.)

In this simplified static model, what happens depends on players' initial responses to the game as shaped by their strategic thinking: specifically, which equilibrium's basin of attraction, “all-In” or “all-Out”, the initial responses fall into.

The leading models of initial responses for games like this are Harsanyi and Selten’s (1988) notions of *payoff-dominance* and *risk-dominance*.

Payoff-dominance favors equilibria that are Pareto-superior to other equilibria.

Hence here it selects the all-In equilibrium, for any value of the population size n and deviation cost (here, the -10).

But that seems behaviorally unlikely, even for small n and “small -10”.

Risk-dominance favors the equilibrium with (roughly) the largest “basin of attraction” in beliefs space.

		Summary statistic In	Summary statistic Out
Representative player	**In**	**1**	**-10**
	Out	**0**	**0**

Bank Runs

In games like this one, that turns out to be the same as selecting the equilibrium that results if each player best responds to a uniform prior over others’ decisions.

Assuming independence of others’ decisions, with these payoffs risk-dominance favors the all-Out equilibrium for any n, even if only two-thirds need to stay In.

That again seems behaviorally unlikely for small n.

3. Game experiments: guessing and coordination games

“...professional investment may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole; so that each competitor has to pick, not those faces which he himself finds prettiest, but those which he thinks likeliest to catch the fancy of the other competitors, all of whom are looking at the problem from the same point of view. It is not a case of choosing those which, to the best of one’s judgment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees.”

—John Maynard Keynes, *The General Theory of Employment, Interest, and Money*, 1936

n-Person guessing games

Imagine that you and every other member of the class are making simultaneous guesses, between limits 0 and 100.

The member whose guess is closest to the target of 1/2 the class average guess wins a £20 prize (imagine!), with ties broken randomly.

What is your guess?

Now imagine instead that the member whose guess is closest to the target of 2/3 the class average guess wins a £20 prize, with ties broken randomly.

What is your guess?

Let's start by identifying your rationalizable and Nash equilbrium guesses.
(I ignore your own small influences on the class average for simplicity.)

The answers are the same, because the game is strictly dominance-solvable.

For example, if the target p = 1/2:

- It's strictly dominated to guess more than 50 (because 1/2 × 100 ≤ 50).

- Unless you think that other people will make strictly dominated guesses, it's also strictly dominated to guess more than 25 (because 1/2 × 50 ≤ 25).

- And so on, down to 12.5, 6.25, 3.125, and eventually to 0.

Thus "all–0" is the unique Nash equilibrium.

The argument for this equilibrium depends "only" on common knowledge of rationality, not on the assumption that players have the same correct beliefs.

Thus the game provides a direct test of this kind of rationality-based reasoning.

In Nagel's (1995 *AER*) *n*-person guessing game design:

- 15-18 subjects simultaneously guessed between [0,100].
- The subject whose guess was closest to a target *p* (= 1/2 or 2/3), times the group average guess won a prize, roughly £20.
- The structure was publicly announced.

Nagel's subjects played these games repeatedly, but we can view their initial guesses as responses to games played as if in isolation if they treated their influences on the future as negligible, which is plausible in groups of 15 to 18.

In those initial responses, subjects never played their equilibrium strategies.

Instead there were spikes, which compared across treatments, appear to reflect a distribution of discrete "level-*k*" rules of thumb, whereby subjects start with a naïve prior that the average guess will be random on [0,100], and then iterate best responses one to three times (Crawford, Costa-Gomes, and Iriberri, "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications," *Journal of Economic Literature* 51 (March 2013), 5-62; http://weber.ucsd.edu/~vcrawfor/CCGIJEL13.pdf)

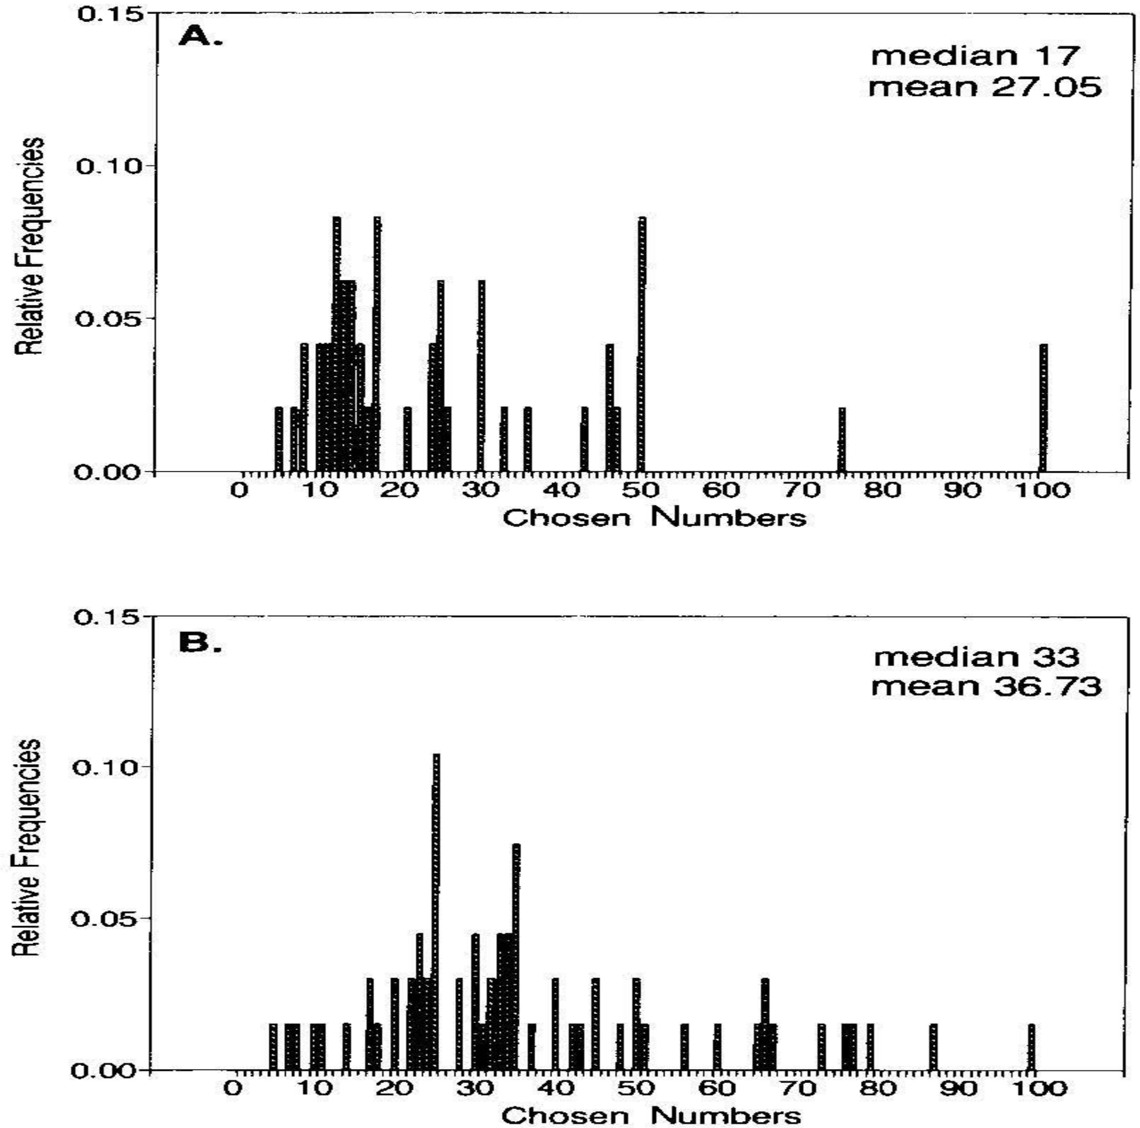

Part of Nagel's (1995 *AER*) Figure 1: top of figure p = 1/2, bottom p = 2/3.

Costa-Gomes and Crawford (2006 *AER*) studied two-person guessing games, with subjects randomly and anonymously paired to play a series of 16.

Subjects played the series only once, and the design suppresses learning and repeated-game effects in order to elicit their initial responses to each game played as if in isolation, game by game.

The profile of a subject's guesses in the 16 games forms a "fingerprint" that helps to identify his strategic thinking more precisely than is possible by observing his responses to a series of games with small strategy spaces (as in Stahl and Wilson 1995 *GEB*) or a single game with large strategy space (Nagel 1995 *AER*).

In Costa-Gomes and Crawford’s guessing games, each player has his own lower and upper limit, both strictly positive.

Each player also has his own target.

Players make simultaneous guesses, and each player’s payoff increases with the closeness of his guess to his target times the other’s guess.

The targets and limits vary independently across players and games, with targets both less than one, both greater than one, or “mixed”.

(In Nagel’s experiments, the targets and limits were always the same for both players, and they varied at most across treatments with different subject groups.)

For example, consider this game, and imagine that each player earns 1000p minus the distance between her/his guess and the product of her/his target times the other's guess.

	Lower Limit	Target	Upper Limit
Player 1	200	0.7	600
Player 2	400	1.5	700

If you are Player 1, playing with one randomly selected member of the class as Player 2, what is your guess?

If you are Player 2, playing with one randomly selected member of the class as Player 1, what is your guess?

Costa-Gomes and Crawford's games are all finitely dominance-solvable (Nagel's were infinitely dominance-solvable), with unique Nash equilibria determined by players' lower (upper) limits when the product of targets is less (greater) than one.

Consider a different game, in which players' targets are 0.7 and 1.5, the first player's limits are [300, 500], and the second player's are [100, 900].

No guess is dominated for the first player, but any guess outside [450, 750] is dominated for the second player.

Given this, any guess outside [315, 500] becomes dominated for the first player.

Given this, any guess outside [472.5, 750] becomes dominated for the second player.

And so on until we reach (500, 750) after 22 rounds.

(500, 750) is the (unique) Nash equilibrium, as is easily checked directly.

It can be shown that because the product of players' targets is 1.05 > 1, the Nash equilibrium (500, 750) is determined by players' upper limits.

But not directly: In equilibrium the first player guesses his upper limit of 500, but the second guesses 750 (= 500 × his target 1.5), below his upper limit of 900.

When the product of targets is < 1, the equilibrium is determined by players' lower limits in a similar way.

In Costa-Gomes and Crawford's experiment, only 1/8th of the subjects played close to their equilibrium guesses.

Most subjects instead closely followed level-*k* rules of thumb, with levels concentrated on 1, 2, or in a few cases 3.

Miguel Costa-Gomes and Vincent Crawford, "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *American Economic Review* 96 (2006), 1737-1768; http://dss.ucsd.edu/~vcrawfor/CGCAER06.pdf.

The “continental divide” game

In Van Huyck, Cook, and Battalio’s (1997 *JEBO*) experiment, 7 subjects chose simultaneously among efforts from 1 to 14, with each subject’s payoff determined by his own effort and a summary statistic, the median, of all players’ efforts.

After subjects chose their efforts, the group median was publicly announced, subjects chose new efforts, and the process continued.

The relation between a subject’s effort, the median effort, and his payoff was publicly announced via a table as on the next slide.

The payoffs of a player’s best responses to each possible median are highlighted in bold in the table as displayed here (but not as displayed to subjects).

The payoffs of the (symmetric, pure-strategy) Nash equilibria, “all–3” and “all–12”, are highlighted in large bold.

Continental divide game payoffs

Your Choice	Median Choice 1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	45	49	52	55	56	55	46	-59	-88	-105	-117	-127	-135	-142
2	**48**	53	58	62	65	66	61	-27	-52	-67	-77	-86	-92	-98
3	**48**	**54**	**60**	**66**	70	74	72	1	-20	-32	-41	-48	-53	-58
4	43	51	58	65	**71**	**77**	80	26	8	-2	-9	-14	-19	-22
5	35	44	52	60	69	**77**	**83**	46	32	25	19	15	12	10
6	23	33	42	52	62	72	82	62	53	47	43	41	39	38
7	7	18	28	40	51	64	78	75	69	66	64	63	62	62
8	-13	-1	11	23	37	51	69	83	81	80	80	80	81	82
9	-37	-24	-11	3	18	35	57	88	89	91	92	94	96	98
10	-65	-51	-37	-21	-4	15	40	**89**	**94**	98	101	104	107	110
11	-97	-82	-66	-49	-31	-9	20	85	**94**	**100**	105	110	114	119
12	-133	-117	-100	-82	-61	-37	-5	78	91	99	**106**	**112**	**118**	**123**
13	-173	-156	-137	-118	-96	-69	-33	67	83	94	103	110	117	**123**
14	-217	-198	-179	-158	-134	-105	-65	52	72	85	95	104	112	120

There were ten sessions, each with its own separate subject group.

Half the groups happened to have an initial median of eight or above, and half happened to have an initial median of seven or below.

(The experimenters probably chose the design to try to make this happen, but this kind of variation across sessions is not uncommon.)

The results are graphed on the next slide:

- The median-eight-or-above groups converged almost perfectly to the all–12 equilibrium.

- By contrast, the median-seven-or-below groups converged almost perfectly to the all–3 equilibrium.

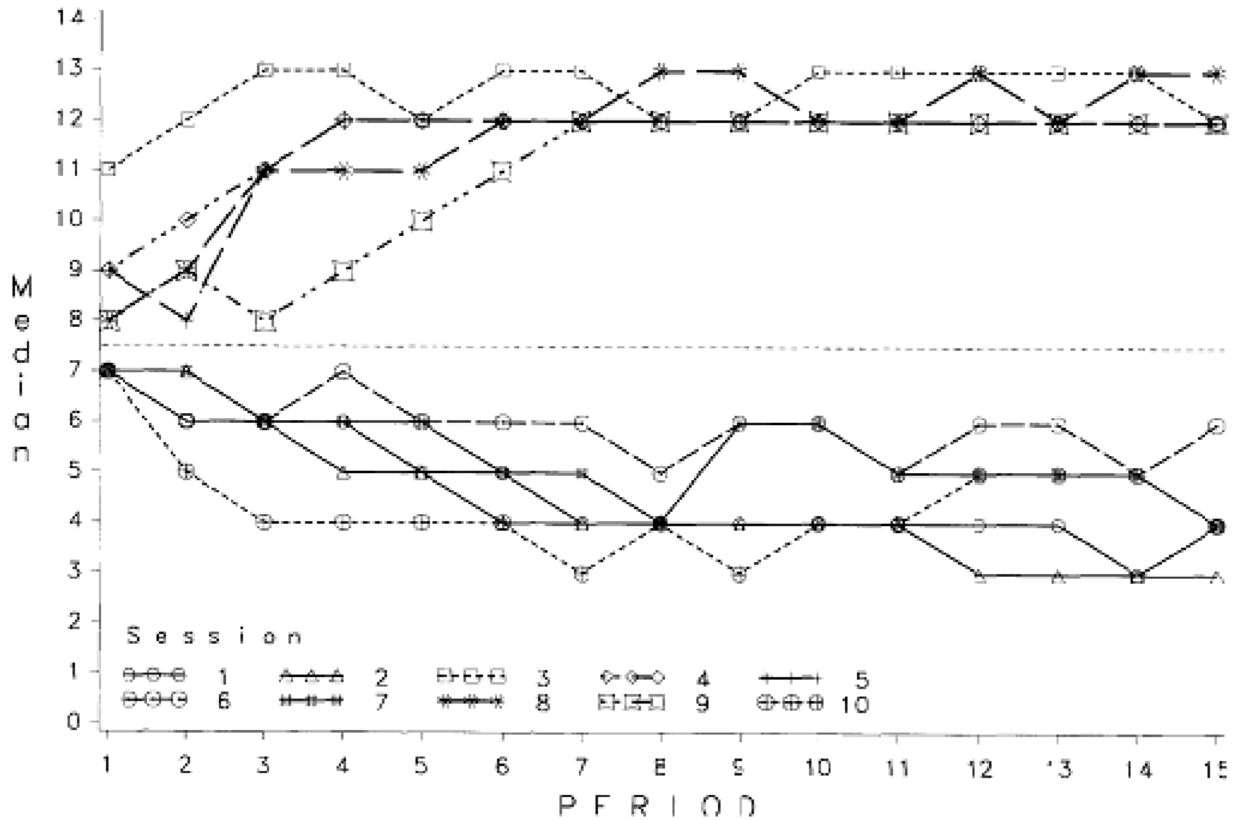

Fig. 3. Median choice in sessions 1 to 10 by period.

Van Huyck, Cook, and Battalio's Figure 3

The results strongly suggest that, as soon as people can observe other people's decisions in analogous games, strategic thinking is eclipsed by adaptive learning.

In adaptive learning, players adjust their decisions in the stage game in a direction that would increase payoffs, other things equal, given the current state.

But, even though players converge to a Nash equilibrium in the stage game, the outcome is highly history-dependent, determined by which equilibrium's basin of attraction subjects' initial responses to the stage game falls into (in the case, their median initial response is all that matters).

Thus even if we care only about the limiting outcome, we need to understand both strategic thinking and learning to predict it.

John Van Huyck, Joseph Cook, and Raymond Battalio (1997): "Adaptive Behavior and Coordination Failure," *Journal of Economic Behavior and Organization* 32, 483-503 (http://www.sciencedirect.com/science/journal/01672681)

4. Repeated games: supporting cooperation via credible threats

A *repeated game* is a dynamic game in which same *stage game* is played over and over again each period by the same players.

The repeated Prisoner's Dilemma is the canonical (but overworked) model of using repeated interaction to overcome short-run incentive problems.

It provides an opportunity to see if repeated play has a chance to overcome short-run incentives to cheat via credible threats to end a relationship.

In economics this is called an “implicit contract”: there is no true altruism, just “reciprocal altruism” supported via purely self-interested behavior.

	Don’t	**Confess**
Don’t	3, 3	0, 5
Confess	5, 0	1, 1

Prisoner’s Dilemma

From now on, Confess = Defect, Don't = Cooperate.

First consider the game repeated a commonly known, finite number of times, T.

However large T is, the strategies (complete contingent plans) (Defect no matter what, Defect no matter what) still form the only Nash equilibrium.

- When T comes, Defect is a strictly dominant decision for both players.

- Even if players contemplate repeated-game strategies that make their decisions functions of past history, if players are rational, nothing can avoid (Defect, Defect) in period T.

- Given that, and that players take the future consequences of current decisions rationally into account, when T-1 comes, Defect is a strictly dominant decision for each player.

- And so on, all the way back to the first period.

Now consider the game repeated an *infinite* number of times.

(View the infinite horizon as only *potentially* infinite, with conditional probabilities of continuation bounded above zero, and perhaps discounting too.)

Assume that players' preferences aggregate their payoffs over plays.

But for this to yield a well-defined preference ordering, they must discount future payoffs; say with a constant discount factor $\delta < 1$, say equal for both.

Thus, for example, a payoff of 1 in every period is valued at $1 + \delta 1 + \delta^2 1 + \ldots$, which is finite when $\delta < 1$. (It equals $1/(1- \delta$.)

We also need to formalize the idea of “credible threats”, for as you will see later in the course, there can be Nash equilibria in sequential games in which one player believes the other will do something crazy in some contingencies, but the cost of such craziness to the first player makes him avoid the contingency, so that the second player’s planned craziness doesn’t cost him anything.

E.g. if I believe you will blow us both up if I don’t give you £20, even though it would hurt you as well as me, and so be (intuitively) irrational for you, it will be (decision-theoretically) rational for me to give you the £20, hence rational for you to blow us both up if I don’t, because you never have to do it; thus such a contingent plan is consistent with Nash equilibrium.

The standard remedy is a *refinement* of Nash equilibrium, *subgame-perfection*.

● A *subgame* is (roughly) a subset of a game that starts with a single decision point, contains all and only that point’s successors in the decision tree, and which all players have enough information to identify.

● A *subgame-perfect equilibrium* is a strategy profile that induces an equilibrium (hence a subgame-perfect equilibrium) in every subgame.

Subgame-perfection eliminates incredible beliefs like the one in my example. A truly rational person would reason that in the subgame that follows his refusal to give the £20, the other person wouldn’t really wish to blow them both up.

With an infinite horizon and a sufficiently high discount factor δ, “grim” trigger strategies (complete contingent plans: “Cooperate until the other player Defects, then Defect forever”) can support the outcome (Cooperate, Cooperate) every period as a Nash equilibrium.

But this equilibrium is not subgame-perfect, because if the other player returns to conditional Cooperation if you Cooperate following his Defection, because his return is conditional it would be better for you to return to Cooperation as well.

Even grimmer trigger strategies (“Cooperate until either of us Defects, then Defect forever”) do support cooperation as a subgame-perfect equilibrium.

	Don't	Confess
Don't	3, 3	0, 5
Confess	5, 0	1, 1

Prisoner's Dilemma

With these payoffs, players can support (Cooperate, Cooperate) in every period in subgame-perfect equilibrium via the grimmer trigger strategies "Cooperate until either of us Defects, then Defect forever" if $3(1 + \delta + \delta^2 + \ldots) = 3/(1 - \delta) \geq 5 + 1(\delta + \delta^2 + \ldots) = 5 + \delta/(1 - \delta)$, which is true if and only if $\delta \geq ½$.

By contrast, when $\delta \leq ½$ the future is not important enough for threats of future defection to support cooperation, and only repeated (Defect, Defect) is consistent with subgame-perfect equilibrium (or Nash equilibrium).

The limit of ½ is dependent on the magnitudes of the payoffs in the example, but fact that higher values of δ never hurt is general.

Infinite-horizon repeated games usually have a huge multiplicity of subgame-perfect equilibria, both of outcomes and the threats used to support them, most asymmetric, e.g. “Row alternates between Cooperate and Defect and Column always Cooperates until either deviates, in which case both Defect from now on”.

	Don’t	**Confess**
Don’t	**3**, **3**	**0**, **5**
Confess	**5**, **0**	**1**, **1**

Prisoner’s Dilemma

In this subgame-perfect equilibrium Column does worse than in the symmetric one above, but Defection yields both the same payoff, so supporting Column’s strategy as part of the equilibrium is harder than supporting Row’s.

In the hypothesized equilibrium, Column gets $3 + 0\delta + 3\delta^2 + \ldots = 3/(1 - \delta^2) \geq 5 + 1(\delta + \delta^2 + \ldots) = 5 + \delta/(1 - \delta)$ if and only if $\delta \geq 0.59$, so it is a subgame-perfect equilibrium if $\delta \geq 0.59$. The limit is higher than for the symmetric equilibrium.

It would be useful to know, more generally, what kinds of implicit contracts can be supported as subgame-perfect equilibria in repeated games. The Folk Theorem will answer this question later in the course.

5. Examples of cooperative games: marriage, college admissions, and matching markets (slides excerpted from Jonathan Levin's)