Lecture Slides: First-Year M. Phil. Micro: Game Theory/Industrial Organization
Vincent P. Crawford (thanks to Chris Wallace, Joel Sobel, and for figures to MWG)
University of Oxford **4 August 2012, corrected 12 November 2012**

There will be 9 or (if needed) 10 lectures, from 9:30-11:00 Monday and *Thursday* week 5; Monday, Tuesday, and Thursday weeks 6 and 7; and Monday and Tuesday week 8.

The lectures will cover the topics of Chapters 7-9 and 12 of Mas-Colell, Whinston, and Green, *Microeconomic Theory*, Oxford 1995 ("MWG") in order, with additional material. MWG and the lectures will be complements, not substitutes: please read both.

Chapter 7. Basic Elements of Non-Cooperative Games
A. Introduction
B. What is a Game?
C. The Extensive Form Representation of a Game
D. Strategies and the Normal Form Representation of a Game
E. Randomized Choices

Chapter 8. Simultaneous-Move Games
A. Introduction
B. Dominant and Dominated Strategies
C. Rationalizable Strategies
D. Nash Equilibrium
E. Games of Incomplete Information: Bayesian Nash Equilibrium
F. The Possibility of Mistakes: Trembling-Hand Perfection
Appendix: Existence of Nash Equilibrium

Chapter 9. Dynamic Games
A. Introduction
B. Sequential Rationality, Backward Induction, and Subgame Perfection
C. Sequential Rationality and Out-of-Equilibrium Beliefs
D. Reasonable Beliefs, Forward Induction, and Normal Form Refinements Appendix
Appendix B: Extensive Form Trembling-Hand Perfection

Chapter 12. Market Power
A. Introduction
B. Monopoly Pricing
C. Static Models of Oligopoly: Bertrand, Cournot, product differentiation
D. Repeated Interaction: Complete-information repeated games
E. Entry
F. The Competitive Limit
G. Strategic Precommitments to Affect Future Competition
Appendix A: Infinitely Repeated Games and the Folk Theorem
Appendix B: Strategic Entry Deterrence and Accommodation

(Plus, time permitting) Cooperative Game Theory

Chapter 7. Basic Elements of Non-Cooperative Games
A. Introduction

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

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

Time permitting, we may make a whirlwind tour of cooperative game theory at the end.

But even if not, you should be aware that cooperative game theory exists, and is better suited to analyzing some economic settings, e.g. where the structure of the game is unclear or unobservable, and it is desired to make predictions that are robust to it.

- *Cooperative* game theory assumes rationality, 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 cooperative agreements that might emerge from rational bargaining.

- It therefore blends normative and positive elements.

- *Noncooperative* game theory also assumes rationality.

But by contrast:

- Noncooperative game theory replaces cooperative game theory's assumptions of unlimited communication and ability to make agreements with a fully detailed model of the situation and a detailed model 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 from the data of the situation.

- As a result, noncooperative game theory is mainly positive, though it is used for normative purposes in some applications, such as mechanism design.

Like the term "game" itself, "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 preferences, played a leading role in the development of the theory—and its public image—most applications combine elements of conflict with elements of coordination.

- Some applications of noncooperative game theory 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 at all) 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 important exceptions) avoids the incentive and coordination issues that are the focus of noncooperative analyses of cooperation.

In game theory, maintaining a clear distinction between the structure of a game and behavioral assumptions about how players respond to it is analytically as important as keeping preferences conceptually separate from feasibility in decision theory.

We will first develop a language to describe the structure of a noncooperative game.

We will then develop a language to describe assumptions about how players behave in games, gradually refining the notion of what it means to make a rational decision.

In the process we will illustrate how game theory can elucidate questions in economics.

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

(This may help you be patient about not yet knowing exactly what it means to be rational.)

B. What is a Game?

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

● the *rules*, which specify the order of players' decisions, their feasible decisions at each point they are called upon to make one, and the information they have at such points

● how players' decisions jointly determine the physical *outcome*

● players' preferences over outcomes (or probability distributions of outcomes)

Assume that the numbers of players, feasible decisions, and time periods are finite.

These can be relaxed, and they will be relaxed here for decisions and time periods.

Preferences over outcomes are modeled just as in decision theory.

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

Assume for now that players face no uncertainty about the structure other than shared uncertainty about how their decisions determine the outcome, that players know that no player faces any other uncertainty, that players know that they know, and so on; i.e. that the structure is *common knowledge*.

Later we will develop a way to model other kinds of uncertainty, shared or not.

● It is essential that a player's decisions 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” can be well-defined if what happens if 2 doesn’t also try is clearly specified.

● It is essential that specifying all of each player’s decisions should completely determine an outcome (or at least a shared probability distribution over outcomes) in the game.

If a specification of the structure of a game does not pass these tests, it must be modified until it does.

E.g. if your model includes a (magical, but useful!) fiction like the Walrasian auctioneer, who always finds prices that balance players’ supplies and demands even though a player’s desired supply cannot be realized without another player’s willing demand, or vice versa, you must replace the auctioneer with an explicit model of how players’ decisions determine realized trades and prices (as called for by Kenneth Arrow in "Toward a Theory of Price Adjustment," Abramovitz et al., eds. *The Allocation of Economic Resources: Essays in Honor of Bernard Francis Haley*, Stanford 1959; and since realized in general equilibrium theory by Shapley and Shubik, "Trade Using One Commodity as a Means of Payment," *Journal of Political Economy* 1977, and in auction theory by everyone).

E.g. if you object to a game analysis on the grounds that players are not really required to participate in the game as modeled, the (only!) remedy is to explicitly add a player’s decision whether to participate to the game, and then to insist that it be explained by the same principles of behavior the analysis uses to explain players’ other decisions.

C. The Extensive Form Representation of a Game

Some games that are important in economics have *simultaneous moves*; examples below.

"Simultaneous" means *strategically simultaneous*, in the sense that players' decisions are made without knowledge of others' decisions.

It need not mean literal synchronicity, although that is sufficient for strategic simultaneity.

But many important games have at least some sequential decisions, with some later decisions made *with* knowledge of others' earlier decisions.

We need a way to describe and analyze both kinds of game.

One way to describe either kind of game is via the *extensive form* or *game tree*, which shows a game's sequence of decisions, information, outcomes, and payoffs.

(The other way is via the *strategic* or *normal form* or *payoff function*, discussed later.)

Figure 7.C.1 shows a version of Matching Pennies with sequential decisions, in which Player 1 moves first and player 2 observes 1's decision before 2 chooses his decision.

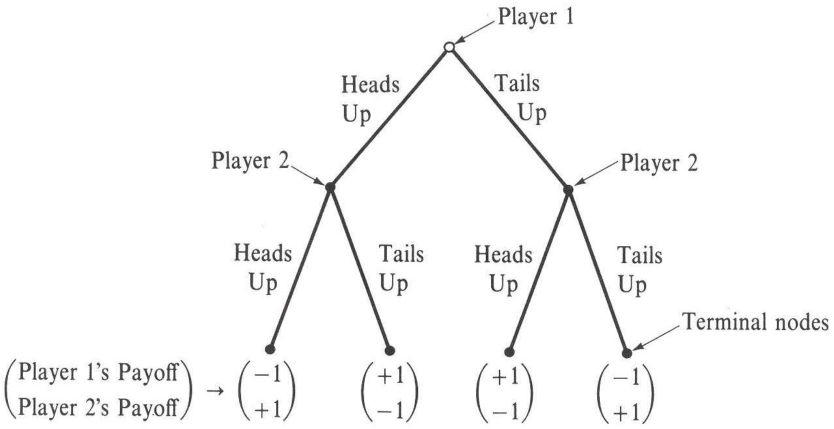

Figure 7.C.1
Extensive form for Matching Pennies Version B.

We can represent the usual Matching Pennies with simultaneous decisions by introducing an *information set*, which includes the decision nodes a player cannot distinguish and at which he must therefore make the same decision, as in the circled nodes in Figure 7.C.3 (or in analogous figures with decision nodes connected by dotted lines as in Kreps).

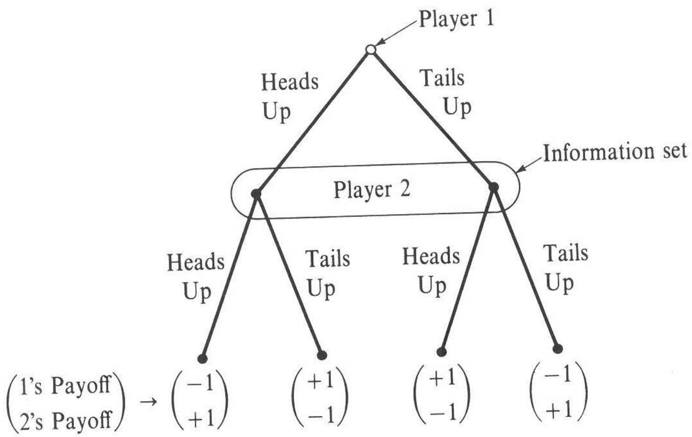

Figure 7.C.3
Extensive form for Matching Pennies Version C.

The order in which simultaneous *decision nodes* are listed has some flexibility, as in Figure 7.C.3 where player 2 could have been at the top; but for sequential decisions the order must respect the timing of information flows. (Information about decisions already made—as opposed to predictions of future decisions—has no reverse gear.)

All decision nodes in an information set must belong to the same player and have the same set of feasible decisions. (Why?)

Figure 7.C.2 gives a partial game tree for Tic-Tac-Toe.

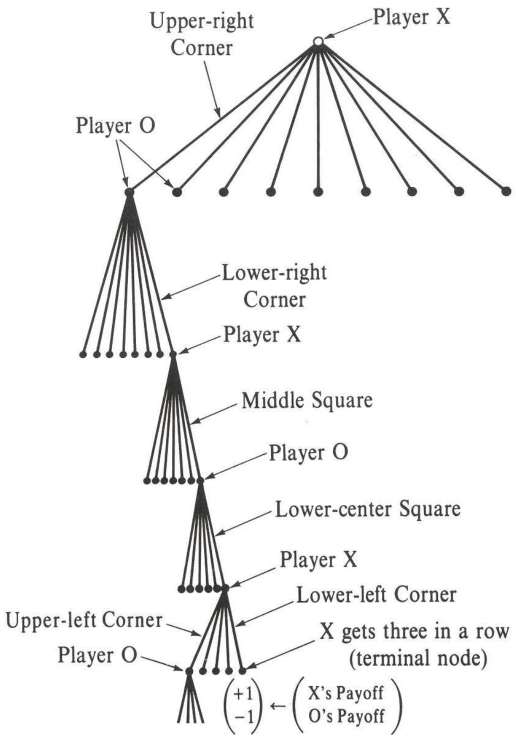

Figure 7.C.2
Part of the extensive form for Tick-Tack-Toe.

Players are normally assumed necessarily to have *perfect recall* of their own past decisions (and other information). If so, the tree must reflect this (as in Figure 7.C.2).

Figure 7.C.4 shows two games *without* perfect recall of players' own past decisions.

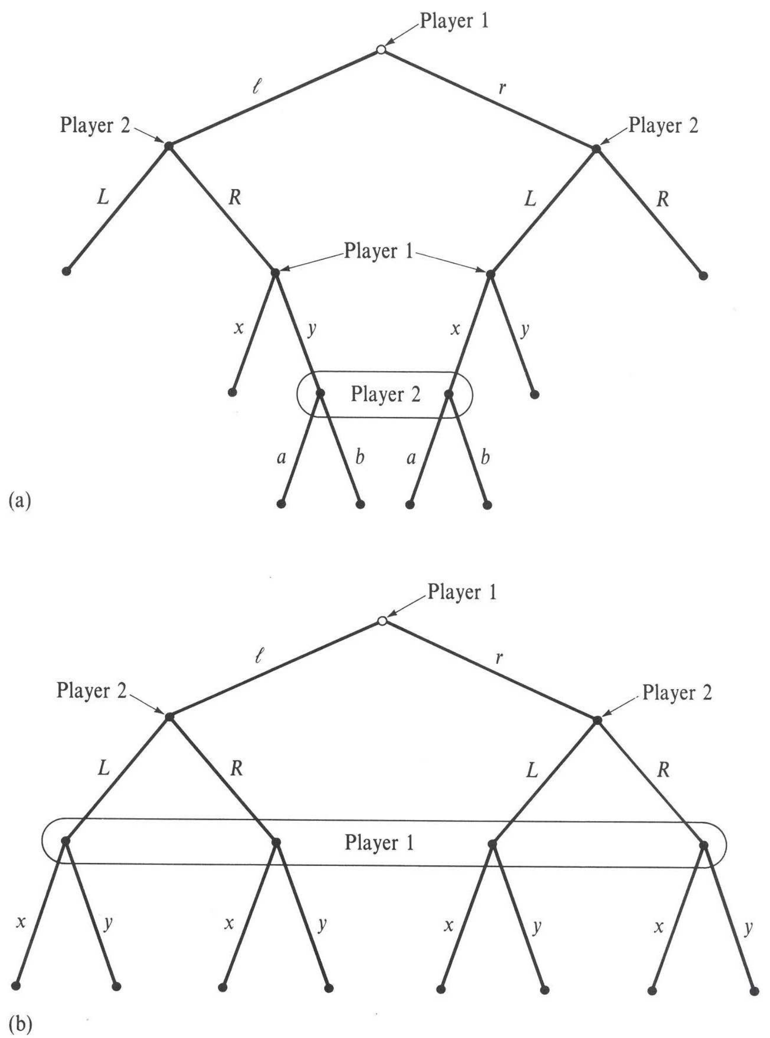

Figure 7.C.4
Two games not satisfying perfect recall.

Shared uncertainty (in economics, roughly "symmetric information") can be modeled by introducing moves by an artificial player (without preferences) called Nature, who chooses the structure of the game randomly, with commonly known probabilities as in Figure 7.C.5.

(It's a good exercise to describe this variant of Matching Pennies in words.)

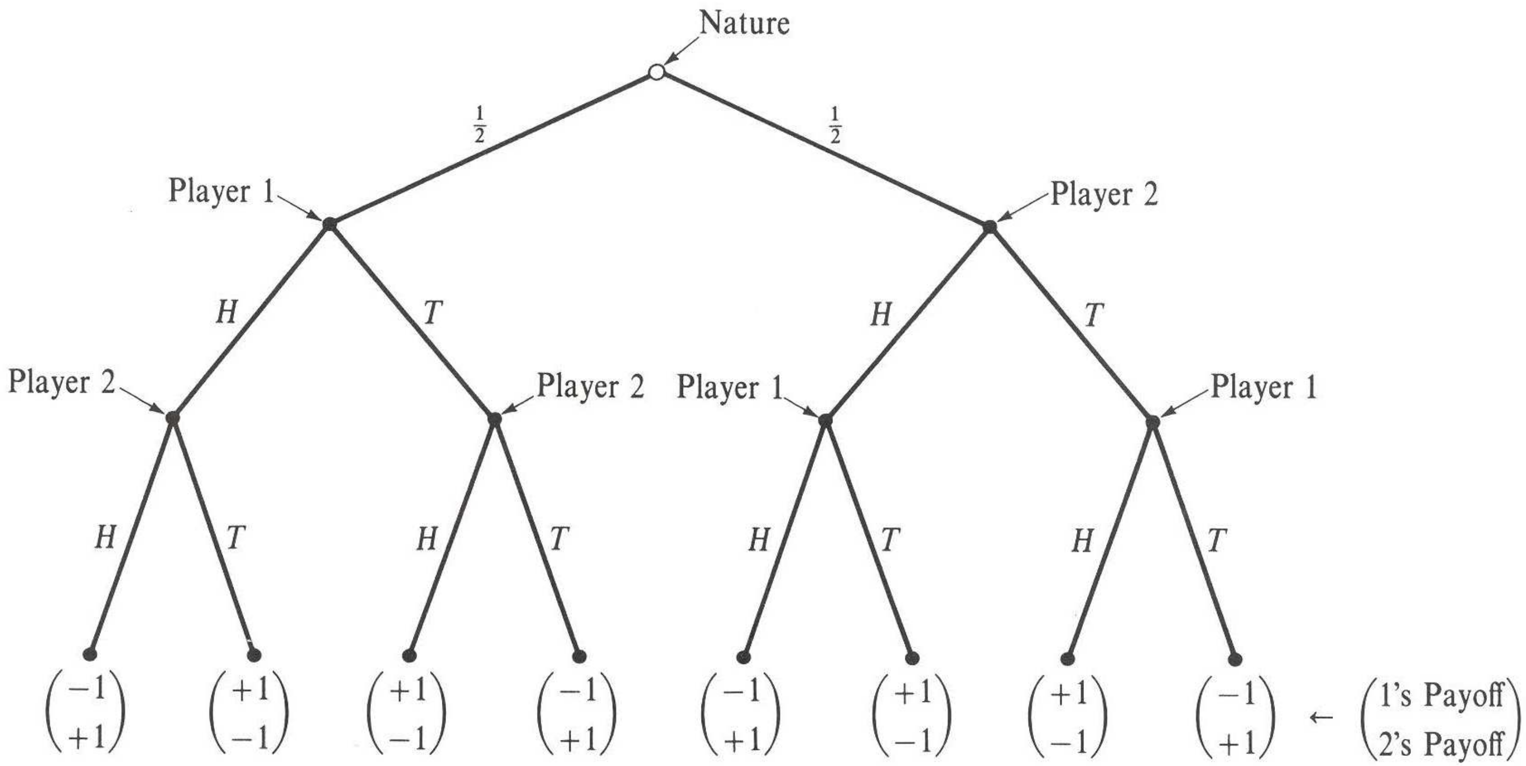

Figure 7.C.5 Extensive form for Matching Pennies Version D.

D. Strategies and the Normal Form Representation of a Game

For sequential games it is important to distinguish *strategies* from *decisions* or *actions*.

A *strategy* is a complete contingent plan for playing the game, which specifies a feasible decision for each of a player's information sets in the game.

(Recall that his decision must be the same for each decision node in an information set.)

Thus a strategy is like a detailed chess textbook, *not* like a single decision or action.

But in a simultaneous-move game a strategy reduces to a single decision or action.

Built into the notions of subgame-perfect, sequential, and perfect Bayesian equilibrium defined below is the assumption that conditional on what a player can observes, he can predict the probability distributions of his own and others' future decisions and their consequences.

If players have this kind of foresight, then their rational sequential decision-making in "real time" should yield exactly the same distribution of decisions as simultaneous choice of fully contingent strategies at the start of play, for reasons essentially like those that justify Bellman's Principle of optimality in dynamic programming.

This allows us to focus, for the purpose of characterizing equilibria, on simultaneous-move games, as including sequential games and blends of simultaneous and sequential games:

"Let us each write our own chess textbook. Then we will give our books to a neutral referee and let him play out the game for us and tell us who won."

(But don't try this at home with commercially available textbooks, which aren't complete.)

Because strategies are complete contingent plans, players *must* be thought of as choosing them *simultaneously* (without observing others' strategies), independently, and irrevocably at the start of play.

Why must a strategy must be a *complete* contingent plan, specifying decisions even for a player’s own nodes that he knows will be ruled out by his own earlier decisions?

Otherwise other players’ strategies would not contain enough information for a player to evaluate the consequences of his own alternative strategies, which in general requires a complete model of other players’ decisions.

We would then be unable to correctly formalize the idea that a strategy choice is rational.

Putting the point in an only seemingly different way, in individual decision theory zero-probability events can be ignored as irrelevant, at least for expected-utility maximizers.

But in games zero-probability events cannot be ignored because what has zero probability is *endogenously* determined by players' strategies.

Some terminology: A *game* maps strategy profiles (one for each player) into payoffs (with outcomes implicit). A *game form* maps profiles into outcomes, without specifying payoffs.

Specifying strategies make it possible to describe an extensive-form game's relationship between strategy profiles and payoffs by its (unique) *normal form* or *payoff matrix* or (usually when strategies are continuously variable) *payoff function*.

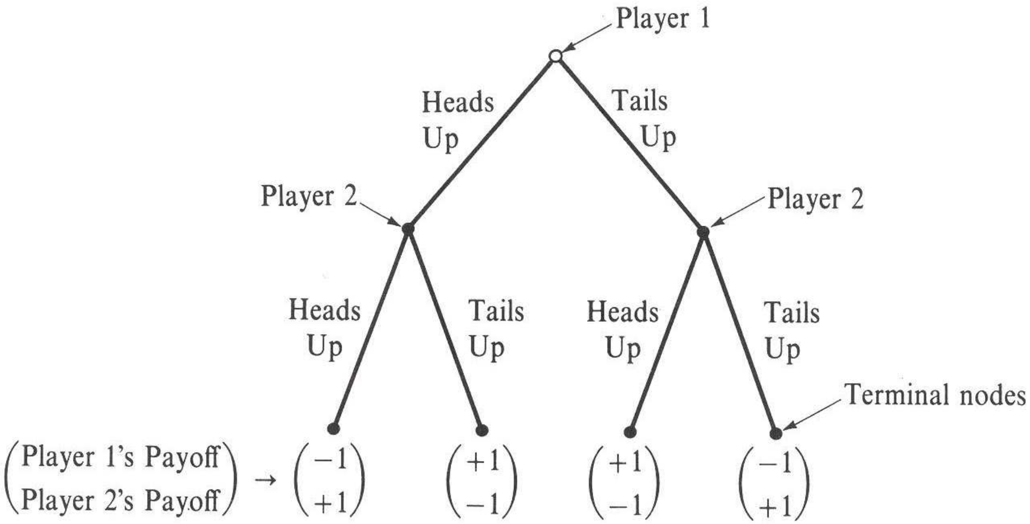

Figure 7.C.1
Extensive form for Matching Pennies Version B.

Player 2

		s_1	s_2	s_3	s_4
Player 1	H	−1, +1	−1, +1	+1, −1	+1, −1
	T	+1, −1	−1, +1	+1, −1	−1, +1

Figure 7.D.1
The normal form of Matching Pennies Version B.

The mapping from the extensive to the normal form isn't univalent in both directions, e.g. the normal form for the sequential version B of Matching Pennies:

		Player 2			
		s_1	s_2	s_3	s_4
Player 1	H	−1, +1	−1, +1	+1, −1	+1, −1
	T	+1, −1	−1, +1	+1, −1	−1, +1

Figure 7.D.1 The normal form of Matching Pennies Version B.

has possible extensive forms other than the one in Figure 7.C.1, such as the canonical:

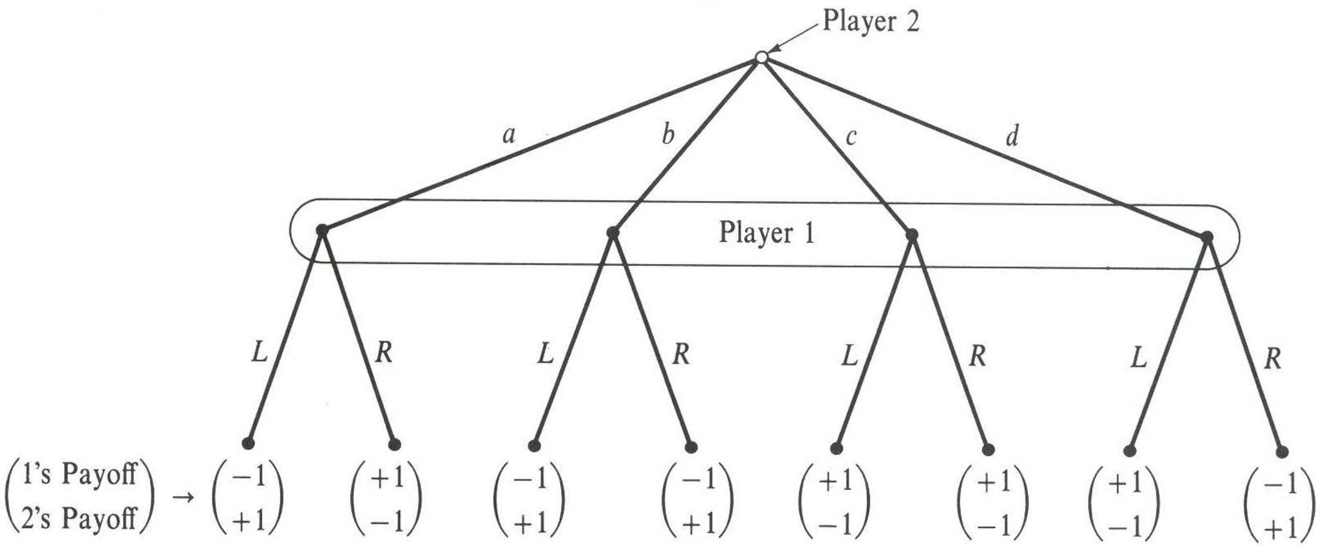

Figure 7.D.2 An extensive form whose normal form is that depicted in Figure 7.D.1.

E. Randomized Choices

In game theory it is useful to extend the idea of strategy from the unrandomized (*pure*) notion we have considered to allow *mixed strategies* (randomized strategy choices).

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

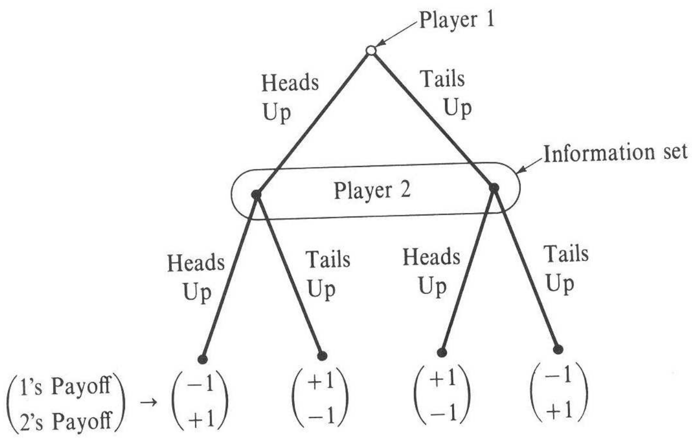

Figure 7.C.3
Extensive form for Matching Pennies Version C.

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

Matching Pennies

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

Mixed strategies will enable us to show that (reasonably well-behaved) games always have rational strategy combinations, i.e. that *Nash equilibria* always exist.

In extensive-form games with perfect recall, mixed strategies are equivalent to *behavior strategies*, probability distributions over pure decisions at each node (Kuhn's Theorem; see MWG problem 7.E.1).

Chapter 8. Simultaneous-Move Games

A. Introduction: Intuition-building Examples
B. Dominant and Dominated Strategies

Define *strictly* or *weakly dominant* and *dominated* strategies, e.g. (for strictly) in

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

Crusoe "v." Crusoe

Crusoe v. Crusoe is not really a game, just two individual decision problems; each player therefore has a strategy that is best independent of the other's strategy, hence dominant.

If there is a (strictly or weakly) dominant strategy, all other strategies must be dominated.

But there can be dominated strategies without there being a dominant strategy, which makes the notion of dominated more useful than the notion of dominant strategy.

The idea is that a rational player would never play a strictly dominated strategy, because there are no beliefs about others' strategies that make it a best response.

A rational player might play a weakly dominated strategy if he has sharply focused beliefs.

Due to the linearity in probabilities of expected payoffs, dominance (strict or weak) for pure strategies implies (strict or at least weak) dominance for mixed strategies with positive probabilities only on those pure strategies.

But there can be dominance by mixed strategies without dominance by pure strategies, e.g. for Column in Domination via Mixed Strategies R is strictly dominated by a 50-50 mix of L and C. Also see Figure 8.B.5.

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

Domination via Mixed Strategies

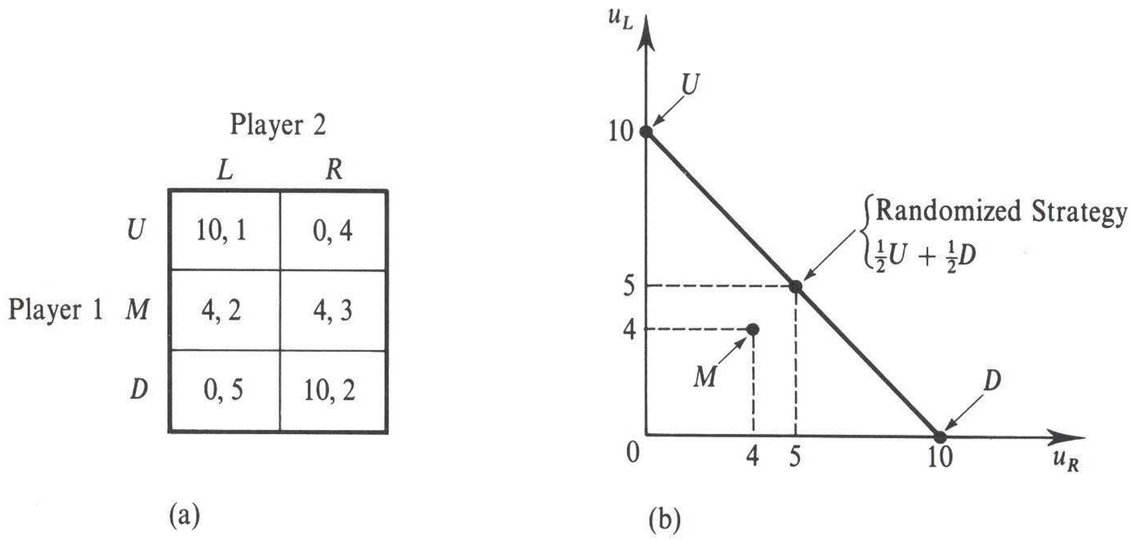

Figure 8.B.5
Domination of a pure strategy by a randomized strategy.

	Confess	Don't
Confess	-5 / -5	-1 / -10
Don't	-10 / -1	-2 / -2

Prisoner's Dilemma

In Prisoner's Dilemma (unlike in Crusoe v. Crusoe) players' decisions do affect each other's payoffs. Even so, each player still has a strictly dominant strategy.

("Don't" = "Cooperate"; "Confess" = "Defect". Why "'s" and not "s'" (which the game's inventor insisted on)? Methodological individualism.)

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

Prisoner's Dilemma is the simplest possible model of incentive problems, which makes it a popular platform for analyses of institutions that overcome such problems.

And the fact that Prisoner's Dilemmas or similar situations abound in real societies alone suffices to show the fatal intellectual flaw in libertarianism.

Yet a Prisoner's Dilemma model is far too simple, because it ignores the difficulty of coordination and possible conflicts of interest between different ways to cooperate.

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

Pigs in a Box

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

There is a lever at one end, which when pushed yields 10 units of grain at the other end.

The story behind the matrix: Pushing costs either pig the equivalent of 2 units of grain.

Identify payoffs with the amount of grain consumed, less the cost (if any) of pushing.

Further, if R (big pig) pushes while C (little pig) waits, C can eat 5 units before R lumbers down and shoves C aside.

But if C pushes while R waits, C cannot push R aside and R gets all but one unit.

If both C and R push, arriving at the grain at the same time, C gets 3 units and R gets 7.

If both C and R wait, both get 0.

When behavior settles down in experiments with pigs, it tends to be at R Push, C Wait.

The little pig (C) does better, even though the big pig (R) can do anything C can do!

This couldn't happen in an individual decision problem: a larger feasible set can never make a rational decision-maker worse off.

It happens here because Wait strictly dominates Push for C, but not for R: the way players' payoffs are determined means that only R has an incentive to Push. (This makes the game what we will call *dominance-solvable* below.)

Thus in games, (the right kind of) weakness can be an advantage! R might get a higher payoff if he could somehow commit himself, say by limiting his ability to shove C aside, to giving C some of the grain to create an incentive for C to Push.

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

If the pigs had studied game theory, they wouldn't have to "settle down": They could just figure out at the start (using "iterated dominance") that they should play (R Push, C Wait).

That they eventually got there anyway suggests that learning and rationality-based arguments yield the same conclusions in the long run. (Why does this happen here?)

Define *iterated deletion* of strictly dominated strategies (*iterated strict dominance*).

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

Pigs in a Box

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

Dominance-solvable

The idea is that a rational player would never play a strictly dominated strategy, because there are no beliefs about others' strategies that make it a best response.

(Even so, a rational player might play a *weakly* dominated strategy for some beliefs.)

Further, a rational player who knows that the other player is rational, knows that the other player knows that he himself is rational, and so on, would never play a strategy that does not survive iterated deletion of strictly dominated strategies.

The result of iterated *strict* dominance is independent of the order of elimination (MWG, problem 8.B.4).

When iterated *strict* dominance reduces the game to a unique strategy profile, the game is called *dominance-solvable*, as in Domination via Mixed Strategies or Dominance-solvable:

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

Domination via Mixed Strategies

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

Dominance-solvable

By contrast, the result of iterated *weak* dominance may *not* be independent of the order of elimination, as in Give Me a Break:

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

Give Me a Break

Thus even when iterated weak dominance reduces the game to a unique strategy profile, the result may not yield a unique profile.

Iterated weak dominance is often useful, but it must be used with care.

C. Rationalizable Strategies

Now we will start to formalize the idea of rational decisions in games.

The idea must be consistent with the idea of rationality for individual decisions, i.e. a player's rational strategy must at least be defined as one that maximizes his expected payoff, given some beliefs.

But that is not the end of the story, because in games the outcome is influenced by other players' decisions as well as the player's own decisions.

Thus a player's beliefs are not only about background uncertainty, as we know how to handle for individual decisions; but also about the strategies chosen by other players.

The problem is that those strategies are chosen by players who are presumed also to be rational, and who recognize the need to predict the player's own rational decision, and who recognize the player's need to predict *their* rational decisions, and so on….

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 ad infinitum.

Focus on the problem of predicting other players' strategies by assuming for now that the structure of the game is common knowledge.

This allows simultaneous decisions as in Figure 7.C.3, and shared uncertainty about how players' decisions determine the outcome with commonly known distributions, modeled as “moves by nature” as in Figure 7.C.5; but they won’t matter for this discussion.

In game theory a game whose structure is common knowledge is called a game of *complete information* (or replacing the old-fashioned game-theory term with a roughly equivalent modern economic term, “symmetric information”).

“Complete” does not imply “perfect” information, e.g. with simultaneous decisions.

A first guess at how to formalize the idea of rational decisions in games is that assuming that players are rational in the decision-theoretic sense of maximizing expected payoffs given some beliefs is enough to yield a useful theory of behavior in games.

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

But that guess fails badly in even slightly more complex games, such as Pigs in a Box or Give Me a Break.

A second guess is that assuming that players are rational *and* that that fact is common knowledge yields a useful theory.

That guess works in some games, such as Domination via Mixed Strategies (using mixed as well as pure-strategies) or Dominance-solvable:

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

Domination via Mixed Strategies

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

Dominance-solvable

In these dominance-solvable games, iterated *strict* dominance reduces the game to a unique strategy profile, which we will see means that common knowledge of players' rationality yields a unique prediction.

But even that guess also fails badly in most economically interesting games.

E.g. *any* strategy is consistent with common knowledge of rationality for some beliefs in Matching Pennies, Gives Me a Break, or Unique Equilibrium without Dominance:

	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 without Dominance

In Unique Equilibrium without Dominance there is a “tower” (my co-authors tell me it should be “helix”) of beliefs, consistent at all levels with common knowledge of rationality, to support any combination of strategies.

But except for the beliefs in the helix that supports the strategy combination (M, C), the beliefs that make other strategies consistent with common knowledge of rationality differ wildly across players and levels.

We will see that (M, C) is Unique Equilibrium without Dominance’s unique Nash equilibrium.

Even though the second guess often fails, it is useful to characterize the implications of the assumption that it is common knowledge that players are rational.

This leads to a notion called *rationalizability*.

MWG define a *rationalizable* strategy as one that survives the iterated removal of strategies that are never a (weak) best response to any beliefs.

As with iterated strict dominance, the order of removal of such strategies doesn't matter (MWG problem 8.C.2).

Player 2

Player 1		b_1	b_2	b_3	b_4
	a_1	0, 7	2, 5	7, 0	0, 1
	a_2	5, 2	3, 3	5, 2	0, 1
	a_3	7, 0	2, 5	0, 7	0, 1
	a_4	0, 0	0, −2	0, 0	10, −1

Figure 8.C.1
$\{a_1, a_2, a_3\}$ are rationalizable strategies for player 1; $\{b_1, b_2, b_3\}$ are rationalizable strategies for player 2

The set of strategies that survive the iterated removal of never-weak-best response strategies cannot be larger than the set that survive iterated strict dominance, because strictly dominated strategies can never be weak best responses to any beliefs.

In *n*–person games, *n* > 2, the set of strategies that survive the iterated removal of never-weak-best responses can be smaller than the set that survives iterated strict dominance.

But any strategy that is not strictly dominated must be a best response to some *correlated* combination of others' strategies (separating hyperplane theorem, MWG Exercise 8.C.4).

In two-person games the two sets are the same because with only one other person, never-weak-best-response strategies are exactly those that are strictly dominated (separating hyperplane theorem again; in general they can differ; MWG Sections 8.D-E).

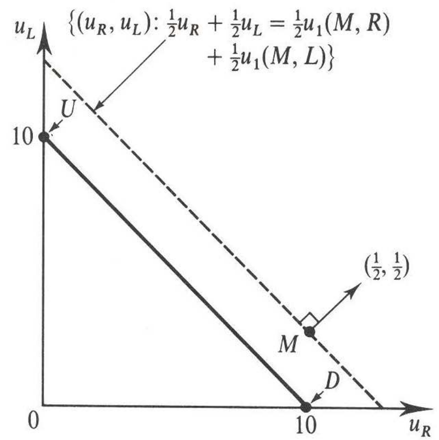

Figure 8.C.2
In a two-player game, a strategy is a best response if it is not strictly dominated.

Focus on two-person games, where the correlation of others' strategies is irrelevant.

Approach the notion of rationalizability via a sequence of notions called *k*-rationalizability, defined to reflect the implications of *k* levels of mutual knowledge of rationality (i.e. all players know that all are rational, know that all know it, and so on, up to *k* levels).

Rationalizability, which reflects the implications of common knowledge of rationality, is then equivalent to *k*-rationalizability for all *k*.

A 1-rationalizable strategy (the sets R1 on the next slide) is one for which there is a profile of others' strategies that makes it a best response.

A 2-rationalizable strategy (the sets R2) is one for which there exists a profile of others' 1-rationalizable strategies that make it a best response.

And so on, recursively….

For two-person games, *k*-rationalizability for all *k* is equivalent to MWG's definition of a rationalizable strategy as one that survives iterated removal of never-weak-best response strategies; and to a definition in terms of surviving iterated strict dominance.

Rationalizability and k-rationalizability generally yield set-valued restrictions on individual players' strategy choices (unlike Nash equilibrium, which restricts their relationship).

In Dominance-solvable, M and C are the only rationalizable strategies; in Unique Equilibrium without Dominance all strategies are rationalizable, for each player. (Each game has a unique Nash equilibrium (M, C).)

		R1,R2	**R1,R2,R3,R4**	
		L	**C**	**R**
R1,R2,R3	**T**	**0** **7**	**5** **0**	**3** **0**
R1,R2,R3,R4	**M**	**0** **5**	**2** **2**	**0** **5**
R1	**B**	**7** **0**	**5** **0**	**3** **7**

Dominance-solvable

		Rk for all k	**Rk for all k**	**Rk for all k**
		L	**C**	**R**
Rk for all k	**T**	**0** **7**	**5** **0**	**7** **0**
Rk for all k	**M**	**0** **5**	**2** **2**	**0** **5**
Rk for all k	**B**	**7** **0**	**5** **0**	**0** **7**

Unique Equilibrium without Dominance

Theorem: Common knowledge of 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 for Dominance-solvable and Unique Equilibrium without Dominance games.

(This theorem needs common knowledge only for indefinitely large games. The number of levels of iterated knowledge of rationality needed is just the number of rounds of iterated dominance, which for finite games is bounded by the size of the payoff matrix.)

Consider Ultimatum Contracting.

There are two players, R(ow) and C(olumn), and 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.

This game depends on whether C can observe R's proposal before deciding whether to accept. We can represent either version by its extensive form or game tree.

Whether or not C can observe R's proposal, R has two pure strategies, X and Y.

If C cannot observe R's proposal, C has two pure strategies, a(ccept) and r(eject).

If C can observe R's proposal he can make his decision depend on it, and therefore has four pure strategies, "a (if X proposed), a (if Y proposed)", "a, r", "r, a", and "r, r."

C's additional information 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: It’s feasible for C to ignore R’s proposal, but both R and C know that is not always optimal.

First suppose R prefers Y to X to Z, while C prefers X to Y to Z.

E.g. R's payoffs u(X) = 1, u(Y) = 2, u(Z) = 0; C's payoffs v(X) = 2, v(Y) = 1, v(Z) = 0.

Then we get the normal forms:

	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

Ultimatum Contracting with Unobservable Proposal is dominance-solvable (C's strategy a strictly dominates r; given that, R's strategy Y dominates X), with unique rationalizable (and equilibrium) outcome (Y, a).

Intuition: a rational R knows that a rational C will accept whichever contract R proposes even if he cannot observe it; and R will therefore propose Y, his most preferred contract.

Ultimatum Contracting with *Observable* Proposal is not (strictly) dominance-solvable: C's strategy (a, a) strictly dominates (r, r), but when (r, r) is eliminated there is no strict dominance, so the remaining strategies are all rationalizable.

(What beliefs support (Y; a, r)? Why are they consistent with common knowledge of rationality?)

Now suppose that 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 payoffs are unchanged).

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

Ultimatum Contracting with Unobservable Proposal

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

Ultimatum Contracting with Observable Proposal

Ultimatum Contracting with Unobservable Proposal is no longer dominance-solvable (there is now no strict dominance), and all strategies are rationalizable.

But there is a unique Nash equilibrium, now (Y, r).

Intuition: C knows that R, knowing that C cannot discriminate between proposals, would propose Y, which is worse for C than Z; C will therefore reject.

In Ultimatum Contracting with *Observable* Proposal all strategies but r, a for C are rationalizable.

But there is an equilibrium (X; a, r) that reflects the intuition that if C can observe R's proposal, C will accept X but not Y, so R will propose X, which he prefers to Z, and C will accept. (There are also unintuitive equibria, which we will learn how to rule out below.)

D. Nash Equilibrium

Most economically important games have multiple rationalizable outcomes, so rational players must base their decisions on interdependent predictions of others' decisions that are not dictated by common knowledge of rationality.

In such games much of game theory's power comes from assuming that players choose strategies in *Nash equilibrium*, a strategy profile for which each player's strategy is a *best response* to other players' strategies (a fixed point of the best-response correspondence).

Any equilibrium strategy survives iterated elimination of strictly dominated strategies, and is *k*-rationalizable for all *k*, hence rationalizable. (Why?)

In games that are dominance-solvable in *k* rounds, the surviving combination of *k*-rationalizable strategies is the unique equilibrium, e.g. in Dominance-solvable.

But in general, not all combinations of rationalizable strategies are in equilibrium.

E.g. in Unique Equilibrium without Dominance any strategy is rationalizable, for either player; but there is a unique equilibrium profile of strategies.

A Nash equilibrium is a kind of rational expectations equilibrium, in that if all players expect the same strategy profile and choose strategies that are best responses given their beliefs, their beliefs will be confirmed if and only if they are in equilibrium.

(This differs from the usual notion of rational expectations in that it is players' strategies that are predicted, and players' predictions interact.)

Nash equilibrium is therefore often unconsciously *identified* with rationality in games; but equilibrium is a much stronger assumption than common knowledge of rationality.

Equilibrium requires that players' strategies are best responses to *correct* beliefs; thus it reflects the implications of common knowledge of rationality *plus* common beliefs.

Compare the belief towers supporting equilibrium and non-equilibrium strategy profiles in Unique Equilibrium without Dominance.

Unlike rationalizability, equilibrium is a property of strategy *profiles*, relationships between strategies. *Equilibrium strategy* often refers to any strategy that's part of an equilibrium.

An equilibrium can be either in pure or mixed strategies.

A mixed strategy profile is an equilibrium if for each player, his mixed strategy maximizes his expected payoff over all feasible mixed strategies (MWG Definition 8.D.2).

Theorem: A mixed strategy profile is an equilibrium if and only if for each player, all pure strategies with positive probability yield the same expected payoff, and all pure strategies he uses with zero probability yield no higher expected payoff (MWG Proposition 8.D.1; the Kuhn-Tucker conditions for maximizing expected payoffs that are linear in probabilities).

Mention Nash's population interpretation and the "beliefs" interpretation of mixing.

Can use this and best-response functions to compute equilibria, e.g. in Matching Pennies.

Notice Row’s equilibrium strategy is determined by Column’s payoffs, and vice versa!

It is a good exercise to show from best-response cycles that Unique Equilibrium without Dominance has no mixed-strategy equilibria.

Equilibrium is surely a *necessary* condition for a common, rational prediction about behavior, but how might players come to have correct beliefs?

Traditional rationale: Players (or hyperintelligent pigs) deduce correct (self-fulfilling) beliefs about each other's strategies when they first play a game from common knowledge (why common?) of a theory of strategic behavior that makes a unique prediction for the game in question, and so (if rational in the decision-theoretic sense) play an equilibrium immediately.

Adaptive rationale: Players (like real pigs playing Pigs in a Box) learn to predict others' strategies in repeated play of analogous games, adjusting their strategies over time in response to observed payoffs; and so (if rational in the decision-theoretic sense) eventually converge to equilibrium.

Traditional game theory focuses on rationality-based reasoning, while adaptive learning models make assumptions directly about how players adjust strategies over time.

Both approaches agree that the possible limiting outcomes are Nash equilibria (in the game that is repeated, not the game that describes the entire process).

But the approaches differ on convergence and the likelihoods of alternative equilibria.

Existence of Nash Equilibrium

Theorem: Every finite game has a mixed-strategy Nash equilibrium (MWG Proposition 8.D.2).

Existence of equilibrium may require mixed strategies, in general, as in Matching Pennies.

The next theorem gives a more abstract and more general existence result.

Theorem: Every game whose strategy spaces are nonempty, convex, and compact subsets of Euclidean space, and whose payoff functions are jointly continuous in all players' strategies and quasiconcave in own strategies has a Nash equilibrium (MWG Proposition 8.D.3).

Compare with the conditions for existence of an optimum in a decision problem: Need to add joint continuity and quasiconcavity in own strategies of payoff functions.

Interpret for mixed strategies in finite games, pure strategies in games with continuously variable pure strategies.

These theorems also trivially imply the existence of rationalizable strategies.

More Intuition-building Examples: Nonuniqueness of Equilibrium and Coordination

These games are worth some attention because economics is about coordination, but competitive markets with Walrasian auctioneers beg some important questions.

	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 decide differently even though there may be no clue to break the symmetry of their roles.

Games like Alphonse and Gaston show that coordination may pose nontrivial difficulties even when players' preferences are identical.

(Mixed strategies can help Alphonse and Gaston learn to coordinate in repeated play; but they then serve a completely different purpose than in Matching Pennies.)

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.

(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 "after you, my dear Alphonse" and thus never getting through a doorway. They are mostly forgotten, but they live on in the Alphonse-Gaston games in the dual-control lighting circuits in our homes.)

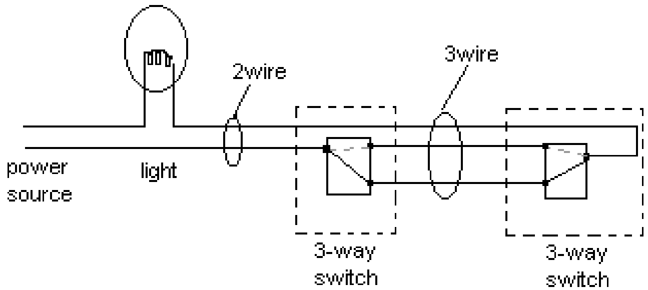

	Other Player: Stag	Other Player: Rabbit
Stag	2, 2	0, 1
Rabbit	1, 0	1, 1

Two-person Stag Hunt

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

***n*-Person Stag Hunt**

In Stag Hunt (Rousseau's story, assembly line, staff 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. Compute it in each case.)

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 autarky and participating in a highly productive but brittle society, which is more rewarding but riskier because dependent on coordination.

Selection among strict equilibria: Harsanyi-Selten's notions of risk- and payoff-dominance.

E. Games of Incomplete Information: Bayesian Nash Equilibrium

Recall that in game theory a game whose structure is common knowledge is called a game of *complete information* (or replacing the old-fashioned game-theory term with a roughly equivalent modern economic term, “symmetric information”).

Recall that a game of complete information need not have *perfect information*: There may still be simultaneous decisions, so that a player making a decision cannot always observe all previous decisions.

A game of *incomplete information* (roughly, “asymmetric information”) allows both simultaneous decisions and players to have private information about the structure.

In 1967-68 Harsanyi, building on von Neumann's 1928 analysis of Poker and Bluffing (*Theory of Games and Economic behavior*, Chapter 19), showed how to adapt the tools of complete-information noncooperative game theory to incomplete-information games.

He argued that all important informational differences across players could be modeled by assigning each player a *type* that parameterizes his preferences; and assuming a player knows his own type when he makes decisions, but other players only have priors about it.

E.g. if player 1 does not know whether a certain decision is feasible for player 2, player 2 might have two types, one of which gives that decision extremely low payoff.

Harsanyi then argued that it is reasonable to assume that any differences in players' beliefs are derived from Bayesian updating of an initially *common* prior within a common model (a view that is now called the *Harsanyi doctrine* or the *common prior assumption*).

Given these arguments, one can reduce a game of incomplete information to an equivalent game of complete information, in which Nature first draws players' types from a common distribution and players observe their own types and then play the game. (Players never directly observe others' types, but might be able to infer them during play.)

The *ex ante*/complete-information view of the game, in which players choose type-contingent strategies before observing their own types (hence with symmetric information), is then analytically equivalent to the *interim*/incomplete-information view of the game, in which players choose their strategies after observing their own types.

A pure-strategy *Bayesian Nash equilibrium* (or *Bayesian equilibrium*) is a profile of *decision rules* (mapping a player's type into a strategy) that are in equilibrium in the ex ante game of complete information (MWG Definition 8.E.1).

A player’s strategy in the ex ante/complete-information game is a complete contingent plan as before, but now it is also contingent on the player’s type.

For reasons already explained, to formalize the idea that a strategy choice is rational, a strategy must be type-contingent even though the player knows his own type before he is called upon to make any decisions.

Thus a profile of decision rules is a Bayesian Nash equilibrium if and only if for all types that have positive prior probability for a player, the player's contingent strategy maximizes his expected payoff given his type, where the expectation is taken over other players' types, conditional on the player's own type (MWG Proposition 8.E.1).

(This definition allows for the possibility that a player’s own type may convey information about other players’ types; but it also allows players’ types to be statistically independent.)

Recall that Ultimatum Contracting has two players, R(ow) and C(olumn).

There are 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.

Ultimatum Contracting depends on whether C can observe R's proposal before deciding whether to accept. We can represent either version by its extensive form or game tree.

Whether or not C can observe R's proposal, R has two pure strategies, X and Y.

If C cannot observe R's proposal, C has two pure strategies, a(ccept) and r(eject).

If C can observe R's proposal he can make his decision depend on it, and C therefore has four pure strategies, "a (if X proposed), a (if Y proposed)", "a, r", "r, a", and "r, r."

Suppose as before that R prefers Y to X to Z, e.g. R's payoffs $u(X) = 1$, $u(Y) = 2$, $u(Z) = 0$.

But now suppose that C has two possible types, C_1 with probability p and C_2 with probability (1-p), with C_1's preferences $v_1(X) = 2$, $v_1(Y) = 0$, and $v_1(Z) = 1$, and C_2's preferences $v_2(X) = 2$, $v_2(Y) = 1$, and $v_2(Z) = 0$.

Thus C_2 but not C_1 will accept R's favorite contract Y; so that R's optimal proposal reflects a tradeoff between its desirability if accepted and the probability of acceptance.

Only C observes his type, but p and the rest of the structure are common knowledge.

R's pure strategies are still X and Y.

C's pure strategies (type-contingent, hence the same for each type) now map his type and R's proposal into an accept or reject decision, so that C has 2x2x2x2 = 16 pure strategies, 4 for each type, chosen independently.

The extensive form has a move by Nature first, then two decision nodes for R in the same information set, then 4 decision nodes for C, each in its own information set. (Why?)

There is a unique (unless $p = ½$) sensible outcome (which will be called a weak perfect Bayesian equilibrium below), in which R proposes Y if $p < ½$, X if $p > ½$, and either if $p = ½$; C_1 accepts X but rejects Y; and C_2 accepts X or Y.

The analysis is easy because the only privately informed player has a passive role.

Now consider Ultimatum Contracting with a continuously divisible pie, say of size 1; and a privately observed, continuously distributed outside option payoff for the responder.

R proposes a division x for R and 1 - x for C, and C accepts or rejects.

If C accepts, R and C get payoffs x and 1 - x.

If C rejects, R gets 0 and C gets outside option payoff y with c.d.f. F(y), where F(0)=0, F(1)=1, and F(·) is continuously differentiable with positive density (e.g. uniform, with F(y) ≡ y when y є [0,1]).

Any proposal risks rejection, with the probability of rejection increasing in x.

Thus R's optimal proposal reflects a tradeoff between desirability and probability of acceptance as before.

For most F(·) there is an essentially unique intuitive outcome (which we will see below is a weak perfect Bayesian equilibrium), in which C accepts iff 1 - x ≥ y (≥ rather than > without loss of generality, because the event 1 - x = y has 0 probability) and R proposes x*, 1 - x*, where x* solves $\max_x xF(1 - x)$. (In the uniform case with F(y) ≡ y, x* = ½.)

For some F(·) the problem $\max_x xF(1 - x)$ has multiple solutions, in which case weak perfect Bayesian equilibrium is essentially nonunique.

The analysis is again easy because the only privately informed player has a passive role.

Now for a harder analysis, where a privately informed player has an active role.
(Thanks to Chris Wallace for many slides in this section and for chapter 12 below.)

Bayesian Battle of the Sexes

"Two M.Phil. students simultaneously decide whether to meet in a pub or a cafe. The first prefers the cafe, whilst the second prefers the pub. Perhaps because of this fact, the first M.Phil. student likes the second; but the feeling may not be mutual. In fact, the second likes the first with probability $\frac{1}{2}$ and *hates* the first with probability $\frac{1}{2}$. Students prefer to spend time with people they like."

Players. M.Phil. student 1, and two *types* of M.Phil. student 2 (2_l and 2_h).

Actions. Each of the players can choose between Cafe and Pub.

Payoffs. The payoffs are given in the below matrices, each of which occurs with probability $\frac{1}{2}$:

	Cafe	Pub
Cafe	4, 3	1, 1
Pub	0, 0	3, 4

or

	Cafe	Pub
Cafe	4, 0	1, 4
Pub	0, 3	3, 1

Information. Player 2 knows which matrix applies. Player 1 doesn't (assigns probability $\frac{1}{2}$ to each).

The Strategic-Form Version

Model this situation as a three-player strategic form game:

Players. Three players $i \in N = \{1, 2_l, 2_h\}$.

Strategies. For each $i \in N$, $s_i \in \{\text{Cafe}, \text{Pub}\}$. Write $s_1 \in \{C, P\}$ and $s_{2_j} \in \{C_j, P_j\}$.

Payoffs. Payoffs are represented in the following matrix:

	C_l, C_h	C_l, P_h	P_l, C_h	P_l, P_h
C	$\underline{4}, \underline{3}, 0$	$\underline{\tfrac{5}{2}}, \underline{3}, \underline{4}$	$\underline{\tfrac{5}{2}}, 1, 0$	$1, 1, \underline{4}$
P	$0, 0, \underline{3}$	$\tfrac{3}{2}, 0, 1$	$\tfrac{3}{2}, \underline{4}, \underline{3}$	$\underline{3}, \underline{4}, 1$

Player 1 gets the payoff in the bottom-left corner, player 2_l gets the payoff in the middle, and player 2_h gets the payoff in the top-right corner. Payoffs from best-responses are underlined.

Nash Equilibria

Having represented the Bayesian (or incomplete information) game as a strategic-form game, a "Bayesian-Nash equilibrium" (BNE) of the former is a Nash equilibrium of the latter.

Pure-Strategy Equilibrium. A single pure-strategy Nash equilibrium exists at $\{C, C_l, P_h\}$.

- The expected payoff to player 1 from playing C is $\frac{1}{2} \times 4 + \frac{1}{2} \times 1 = \frac{5}{2}$...
- Which is more than they would get from deviating to P: $\frac{1}{2} \times 0 + \frac{1}{2} \times 3 = \frac{3}{2}$.
- Player 2_l can do no better (deviating to P_l yields 1 rather than 3).
- Player 2_h can do no better (deviating to C_h yields 0 rather than 4).

Of course, in actuality it is player 2 that chooses C_l and P_h. Could player 2 do any better by varying *both* strategies? No! Player 2_l's decision does not affect player 2_h's payoffs, and vice-versa.

The equilibrium entails student 1 going to the cafe (which is preferred) and student 2 deciding whether or not to go to the cafe depending on whether they like student 1. Seems natural.

Mixed-Strategy Equilibria. There are two: $\{(\frac{1}{3}, \frac{2}{3}), P_l, (\frac{2}{3}, \frac{1}{3})\}$ and $\{(\frac{2}{3}, \frac{1}{3}), (\frac{2}{3}, \frac{1}{3}), P_h\}$.

Mixed Equilibria

To check these mixed equilibria, recall the requirement for indifference.

- Consider $\{(\frac{1}{3}, \frac{2}{3}), P_l, (\frac{2}{3}, \frac{1}{3})\}$. Looking at the payoff matrix on slide 2, what is player 1's payoff?
- From C player 1 receives $\frac{2}{3} \times \frac{5}{2} + \frac{1}{3} \times 1 = 2$. From P player 1 receives $\frac{2}{3} \times \frac{3}{2} + \frac{1}{3} \times 3 = 2$.
- Indifferent, and prepared to mix…what about player 2_l?
- From C_l receives $\frac{1}{3} \times 3 + 0 = 1$. From P_l receives $\frac{1}{3} \times 1 + \frac{2}{3} \times 4 = 3$. So will choose P_l.
- Player 2_h gets $0 + \frac{2}{3} \times 3 = 2$ from C_h, and $\frac{1}{3} \times 4 + \frac{2}{3} \times 1 = 2$ from P_h. Indifferent, so can mix.

It is straightforward to confirm the other mixed equilibrium, and prove that there are no other equilibria, pure or mixed (left as an exercise!)

Put these ideas in a formal framework…

Bayesian Games of Incomplete Information

Definition 15. A Bayesian game of incomplete information consists of:

1. *Players.* A finite set of players labelled $i \in N = \{1, \ldots, n\}$.
2. *Types.* For each $i \in N$, a set of types T_i, with typical member $t_i \in T_i$.
3. *Actions.* For each $i \in N$, a set of actions A_i, with member typical $a_i \in A_i$.
4. *Beliefs.* For each $i \in N$ and $t_i \in T_i$, a probability measure p_i over T_{-i}, written $p_i(t_{-i}|t_i)$.
5. *Payoffs.* For each $i \in N$, a vNM utility function $u_i : A \times T \mapsto \mathcal{R}$.

Anything with these five features can be written as a Bayesian game:

$$\Gamma = \langle N, \{T_i\}_{i \in N}, \{A_i\}_{i \in N}, \{p_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle.$$

Notation. Once again, define $a \in A = \times_{i \in N} A_i$ and $t \in T = \times_{i \in N} T_i$, as action and type profiles.

Formulating the Bayesian Battle of the Sexes

This is a rather formal definition. How does the earlier example fit into it?

1. *Players.* There are two: $N = \{1, 2\}$.
2. *Types.* Player 1 has only one type $T_1 = \{1\}$. Player 2 has two: $t_2 \in T_2 = \{l, h\}$
3. *Actions.* For each player the actions available are $A_1 = A_2 = \{\text{Cafe, Pub}\}$.
4. *Beliefs.* Player 1 has $p_1(l|1) = p_1(h|1) = \frac{1}{2}$. Player 2 has $p_2(1|h) = p_1(1|l) = 1$.
5. *Payoffs.* Are as described in the matrices on the first slide.

Notice that types are *independent*: $p(t_{-i}|t_i) = p(t_{-i})$. Players own types do not reveal information about their opponent's types. This need not be the case — e.g. global games.

Notice that types are *private*: $u_i(a, t) = u_i(a, t_i)$. Payoffs only depend upon own-type draws, and not directly upon opponents' types draws.

Bayesian-Nash Equilibrium

Definition 16. A *Bayesian-Nash Equilibrium* of a Bayesian game Γ is a Nash equilibrium of the strategic-form game $\mathcal{G} = \langle \hat{N}, \{\hat{S}_i\}_{i \in \hat{N}}, \{\hat{u}_i\}_{i \in \hat{N}} \rangle$, where:

1. *Players.* The set of players is $\hat{N} = \times_{j \in N} \{j \times T_j\}$, e.g. player $i = (j, t_j) \in \hat{N}$.
2. *Strategies.* For each $i = (j, t_j) \in \hat{N}$ with $j \in N$ the set of strategies is $\hat{S}_i = A_j$ for all $t_j \in T_j$.
3. *Payoffs.* For each $i = (j, t_j) \in \hat{N}$ with $j \in N$ the vNM utility function $\hat{u}_i : \hat{S} \mapsto \mathcal{R}$ is

$$\hat{u}_i(s_i, s_{-i}) = \sum_{t_{-j} \in T_{-j}} p_j(t_{-j} | t_j) \times u_j(s_i, s_{-j}(t_{-j}); t_j, t_{-j}),$$

   where for each player-type pair's strategy, $s_{(k,\tau)} = s_k(\tau) \in A_k$ for $\tau \in T_k$, and $s_i \in \hat{S}_i = A_j$.

Players are all combinations of players and types: "player-type pairs". **Strategy** sets for each player-type pair is the action set for the associated player. **Payoffs** to strategy profiles are expectations over beliefs of payoffs to the associated action profiles of each player-type pair.

A Bayesian-Nash equilibrium is a situation in which no player-type pair has a profitable deviation.

Bayesian Strategies

Sounds complicated…but isn't difficult to apply.

A *Bayesian Strategy* for a $i \in N$ is a mapping from their types to their actions $s_i : T_i \mapsto A_i$.

It is straightforward to extend this to mixed strategies. A strategy is then a mapping from a given player's type space to the set of probability distributions over their action space:

$$s_i : T_i \mapsto \Delta(A_i).$$

The Bayesian-Nash equilibrium strategy profiles from the earlier game are:

1. Player 1 plays $s_1(1) = C$. Player 2 plays $s_2(t_2)$ where $s_2(l) = C$ and $s_2(h) = P$.
2. Player 1 plays $s_1(1) = (\frac{1}{3}, \frac{2}{3})$. Player 2 plays $s_2(l) = P$ and $s_2(h) = (\frac{2}{3}, \frac{1}{3})$.
3. Player 1 plays $s_1(1) = (\frac{2}{3}, \frac{1}{3})$. Player 2 plays $s_2(l) = (\frac{2}{3}, \frac{1}{3})$ and $s_2(h) = P$.

Another Noisy Battle of the Sexes

Consider the following variation. Player 1 (row) does not know what payoffs player 2 (column) receives from playing Pub. In fact, there is some "noise" on player 2's payoff, δ. Likewise player 1's payoff to Cafe is perturbed by some ε, from the perspective of player 2. Suppose $\delta, \varepsilon \sim U[0, a]$.

	Cafe	Pub
Cafe	$4+\varepsilon$, 3	$1+\varepsilon$, $1+\delta$
Pub	0, 0	3, $4+\delta$

The payoff matrix above illustrates this game. A Bayesian strategy for player 2 maps all their types (of δ) to an action (either Cafe or Pub). Consider the following "cut-off" strategy for player 2:

If $\delta < \bar{\delta}$ play Cafe, if $\delta \geq \bar{\delta}$ play Pub:

A *pure* strategy. The probability that player 2 plays Cafe, from player 1's perspective, is $\bar{\delta}/a$.

A Pure-Strategy Equilibrium

What should player 1 do given player 2 uses this strategy? Play Cafe if

$$[4+\varepsilon]\frac{\bar{\delta}}{a} + [1+\varepsilon]\left(1-\frac{\bar{\delta}}{a}\right) \geq 3\left(1-\frac{\bar{\delta}}{a}\right) \quad \Leftrightarrow \quad \varepsilon \geq 2-6\frac{\bar{\delta}}{a}.$$

Thus player 1 defines their cut-off strategy with $\bar{\varepsilon} = 2 - 6\bar{\delta}/a$ as:

If $\varepsilon \geq \bar{\varepsilon}$ play Cafe, if $\varepsilon < \bar{\varepsilon}$ play Pub.

Given this strategy, an analogous argument yields $\bar{\delta} = 2 - 6\bar{\varepsilon}/a$. Solving for $\bar{\delta}$ and $\bar{\varepsilon}$:

$$\bar{\delta} = \bar{\varepsilon} = \frac{2a}{6+a}$$

Thus, there is a Bayesian-Nash equilibrium in cut-off strategies where player 1 plays Cafe if $\varepsilon \geq 2a/(6+a)$ and Pub otherwise. Player 2 plays Pub if $\delta \geq 2a/(6+a)$ and Cafe otherwise.

Purifying Bayesian-Nash Equilibria

Recall the perfect-information battle-of-the-sexes payoffs with $\varepsilon = \delta = 0$:

	Cafe	Pub
Cafe	3 4	1 1
Pub	0 0	4 3

Recall that the mixed Nash equilibrium involves player 1 playing Cafe with probability $\frac{2}{3}$, and Pub with $\frac{1}{3}$. Player 2 plays Cafe with probability $\frac{1}{3}$ and Pub with $\frac{2}{3}$.

Notice that the probability with which player 1 plays Cafe in the Bayesian game is

$$\Pr\left[\varepsilon \geq \frac{2a}{6+a}\right] = 1 - \frac{2}{6+a} \quad \longrightarrow \quad \frac{2}{3} \quad \text{as} \quad a \to 0.$$

$a \to 0$ collapses the distribution to a point at 0. The *pure*-strategy Bayesian-Nash equilibrium resembles the *mixed*-strategy Nash equilibrium of the unperturbed game: "purification".

Purification

> "The probability distributions over strategies induced by the pure-strategy (Bayesian-Nash) equilibria of the perturbed game converge to the distribution of the (mixed Nash) equilibrium of the unperturbed game."

This is Harsanyian purification. The statement can be made precise, but the idea is simple:

A mixed-strategy Nash equilibrium of an unperturbed game closely resembles a pure-strategy Bayesian equilibrium of a perturbed game when the perturbations become vanishingly small.

Does this fact offer a justification for the prediction of mixed equilibria? Does it blur the distinction between pure and mixed equilibria? Possibly…

- Notice there are other Bayesian-Nash equilibria of the perturbed game…
- Player 1 plays Cafe for all ε, and player 2 plays Cafe for all δ.
- Player 1 plays Pub for all ε, and player 2 plays Pub for all δ.

F. The Possibility of Mistakes: Trembling-Hand Perfection

Equilibrium eliminates strictly dominated strategies. What about weakly dominated strategies?

A Nash equilibrium is (normal-form) *trembling-hand perfect* if there is some sequence of *tremble-perturbed games* (define; a weakly dominated strategy becomes strictly dominated in a tremble-perturbed game) converging to the original game for which there is some sequence of Nash equilibria converging to that equilibrium (MWG Definition 8.F.1).

E.g. (T, L) or (T, R) or mixtures in Give Me a Break; but only (T, L) in Give Us a Break.

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

Give Me a Break

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

Give Us a Break

Theorem: A Nash equilibrium is trembling-hand perfect iff there is a sequence of *totally mixed strategies* converging to the equilibrium such that each player's equilibrium strategy is a best response to every element of the sequence (MWG Proposition 8.F.1).

Theorem: In a trembling-hand perfect equilibrium, no weakly dominated strategy can be played with positive probability (MWG Proposition 8.F.2). But a trembling-hand perfect equilibrium may include strategies that do not survive iterated elimination of weakly dominated strategies, such as (T, R) in Give Me a Break. Any *strict equilibrium* (such as (T, L) in Give Us a Break; define) is trembling-hand perfect. Any *finite game* (define) has a trembling-hand perfect equilibrium in mixed strategies, just as we will see below that any finite game has an equilibrium in mixed strategies.

Informally define extensive-form trembling-hand perfect equilibrium as trembling-hand perfect equilibrium in the *agent normal form* (MWG Definition 9.BB.1).

Chapter 9. Dynamic Games

A. Introduction

B. Sequential Rationality, Backward Induction, and Subgame Perfection

In sequential games some useful ideas depend essentially on the extensive form. Recall Ultimatum Contracting with Observable Proposal.

Two players, R(ow) and C(olumn); 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(X) = 1, u(Y) = 2, u(Z) = 0; C's payoffs: v(X) = 2, v(Y) = 1, v(Z) = 0.

	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

	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

The intuitive outcome (Y; a, a) is an equilibrium. But there are other equilibria, (Y; r, a) and (X; a, r), one with the counterintuitive outcome X, which survive iterated strict dominance.

More generally, whenever play doesn't reach a given node in an equilibrium, equilibrium doesn't restrict the decision at that node at all. (Why?)

In the equilibria (Y; r, a) and (X; a, r), C’s strategy plans to reject one of R's possible proposals, irrationally, and R's anticipation of that keeps R from making that proposal.

As a result, C’s irrationality does not reduce his strategy’s payoff in the entire game.

Such equilibria are said to involve "incredible threats"; misleadingly because the threat is only implicit in the expectations that support the equilibrium, not explicit like a real threat.

Can rule out some such incredible threats via the notion of *subgame-perfect equilibrium*.

The basic idea is that if our "solution concept" for an entire extensive-form game is equilibrium, we ought to be willing to apply it to the games that remain following partial play of the game, i.e. our solution concept ought to be time-consistent (off as well as on the equilibrium path, so stronger than the decision-theoretic notion of time-consistency).

A *subgame* is a subset of a game that starts with an information set with a single node, contains all and only that node's successors in the tree, and contains all or none of the nodes in each information set (MWG Definition 9.B.1, Figure 9.B.5).

Subgames and Imperfect Information

A subgame must always start at a singleton information set. It must also never break an information set. So, for example, there is no subgame starting at player 2's move in the game below:

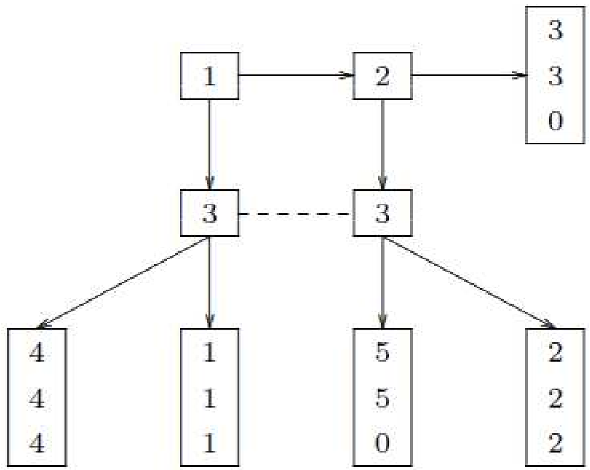

Nor is there any subgame starting from player 3's move. In fact, this game has only one subgame — the whole game. Backward induction and subgame perfection do not help refine the equilibria.

A *subgame-perfect equilibrium* is a strategy profile that induces an equilibrium (hence in fact a subgame-perfect equilibrium) in every subgame (MWG Definition 9.B.2).

In Ultimatum Contracting with Observable Proposal, the intuitive outcome (Y; a, a) is a subgame-perfect equilibrium; but (Y; r, a) and (X; a, r) don't specify equilibria in the subgames in which C rejects.

	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

Strict dominance has no power in normal forms derived from extensive-form games, because contingencies off the equilibrium path create payoff ties between strategies.

As a result, there can be unintuitive rationalizable and even equilibrium strategies.

But subgame-perfect equilibrium mimics iterated weak dominance in a particular order, which the extensive form makes salient.

That order yields a particular, often intuitive equilibrium in undominated strategies.

Theorem: Characterization of subgame-perfect equilibria in finite games of perfect or imperfect information via *backward induction*, plugging in payoffs of equilibria in subgames and folding back (MWG Proposition 9.B.3, Example 9.B.3 predation game).

Entry Deterrence with an Explicit Threat

Modify the entry deterrence game to allow for an explicit threat at a cost of c...

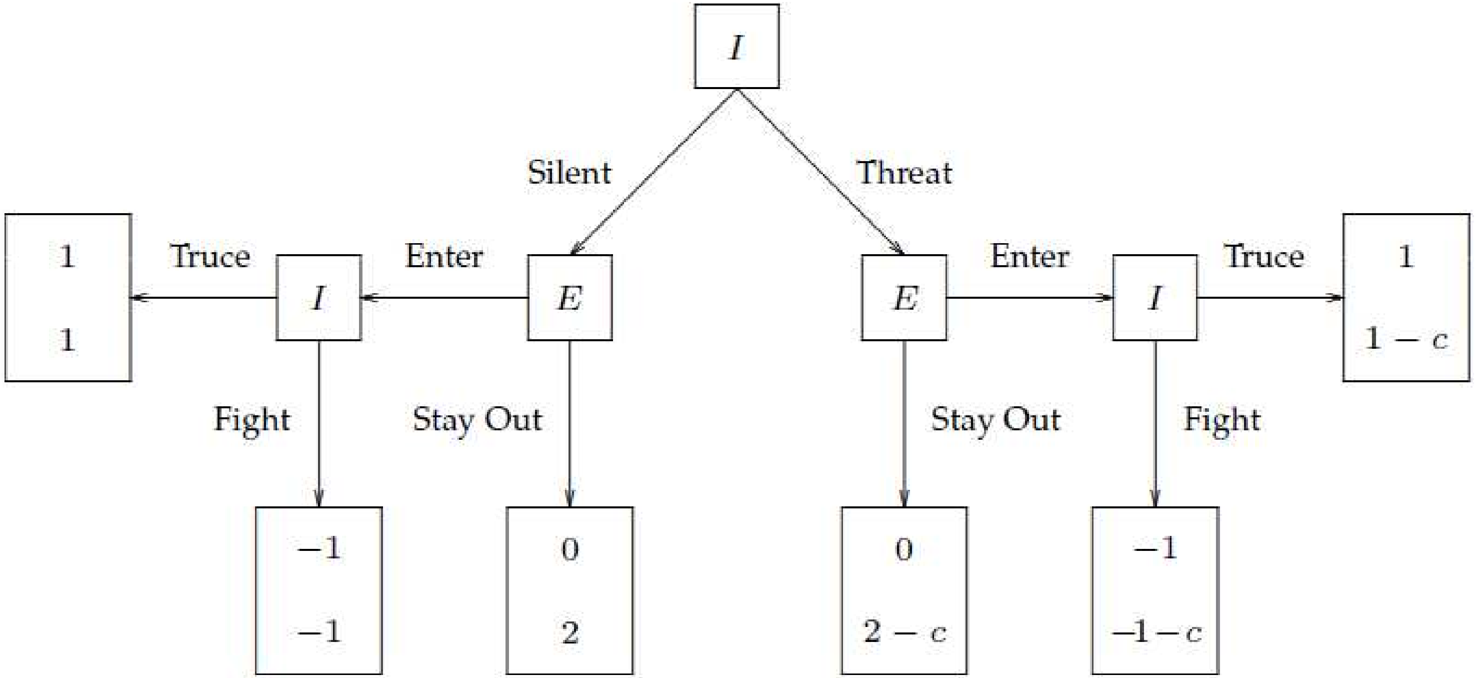

There are 5 subgames. Need to find a Nash equilibrium of each, or use backward-induction... there are many Nash equilibria, but only {(Silent,Truce,Truce), (Enter,Enter)} is subgame perfect.

Subgame-perfect equilibrium in a game with time-sequenced strategy choices yields insight into the roles of observability and irreversibility in commitment.

Without observability decisions are strategically simultaneous, and their temporal order doesn't alter the game’s feasible strategies, payoffs, normal form, equilibrium outcomes, or subgame-perfect equilibrium outcomes.

With observability, but without irreversibility (or at least without costly reversibility, in which case the decision to incur those costs is what is really irreversible), the initial decision inessentially alters the game’s feasible strategies, payoffs, and normal form, and has no effect on the game’s equilibrium outcomes or subgame-perfect equilibrium outcomes.

In particular, a player’s *cheap talk* announcement (one with no *direct* payoff effects) of his *intention* to choose a particular strategy has no effect on the game’s equilibrium or subgame-perfect equilibrium outcomes, although it might help to focus players’ beliefs on a particular equilibrium in games like Alphonse-Gaston or Stag Hunt.

Theorem: Existence of pure-strategy subgame-perfect equilibrium, and (if no ties) uniqueness of subgame-perfect equilibrium in finite games of perfect information. Existence of subgame-perfect equilibrium in games of imperfect or incomplete information (MWG Proposition 9.B.2, generalizes Zermelo's Theorem; MWG Example 9.B.3).

Trembling-hand perfect equilibria are subgame-perfect, but not vice versa. (Why?)

Theorem: **I**n finite-horizon games (even with imperfect information) with unique equilibria and immediate observability of pure strategies each period, and payoffs summed over periods, subgame-perfect equilibrium strategies are the same as the concatenated equilibria of the games played each period. Illustrate proof in finitely repeated Prisoner's Dilemma (MWG Proposition 9.B.4).

Theorem: **I**n finite-horizon games with multiple equilibria and immediate observability of pure strategies each period, and payoffs summed over periods, subgame-perfect equilibrium strategies can differ from the concatenated equilibria of the games played each period (MWG Proposition 9.B.9, illustrated below).

Subgame-perfect equilibrium is not adequate as a formalization of the idea of sequential rationality, because most games that pose nontrivial sequential strategic questions have no proper subgames, so that any equilibrium is subgame-perfect. (MWG Example 9.C.1.)

Subgame Perfection and Imperfect Information

Imperfect-information games may have no "proper" subgames: all equilibria are subgame perfect.

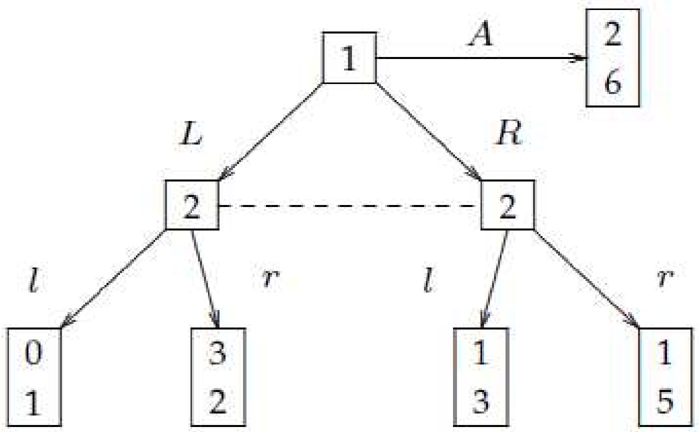

$\Longrightarrow$

	l	r
A	2 6	2 6
L	0 1	3 2
R	1 3	1 5

The above game has no proper subgames. All Nash equilibria are subgame perfect:

- There is a subgame-perfect equilibrium at $\{A, l\}$. This yields payoffs of $(2, 6)$.
- There is a subgame-perfect equilibrium at $\{L, r\}$. This yields payoffs of $(3, 2)$.
- There are also mixed equilibria (A with probability 1 and l with probability $p \geq \frac{1}{3}$).

Surely player 2 will play r if called upon to move? 1 should play L! $\Rightarrow$ sequential rationality.

Modifying the Extensive Form

Rewrite the extensive-form game from the last slide in the following way:

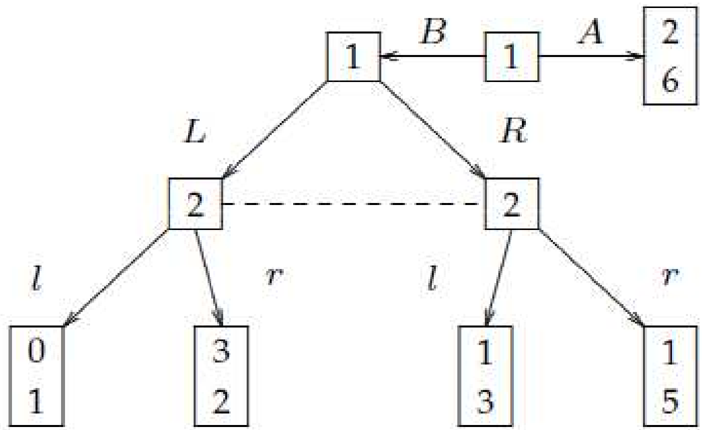

This has (almost) the same strategic form. But there is now a subgame starting from history (B).

- There is a unique Nash equilibrium of this subgame: $\{L, r\}$.
- A subgame-perfect equilibrium must induce a Nash equilibrium in each subgame.
- Hence r must be played in equilibrium. $\{(B, L), r\}$ is the unique subgame-perfect equilibrium.

Extend notion of perfection to games with imperfect/incomplete information next week.

Extensive-Form Games with Imperfect Information

When there is imperfect information, there may be no proper subgames — so subgame perfection may coincide with Nash equilibrium and hence not provide tighter predictions.

- Sometimes it will be possible to modify the extensive form…
- …and subgame perfection may once again rule out some "incredible" threats.
- Often modification of the extensive form will still leave multiple subgame-perfect equilibria.

In this case, further refinements are possible — motivated by players' reasoning.

- Forward Induction may reduce the number of subgame-perfect equilibria…
- …via the iterated deletion of weakly dominated strategies.
- Perfect-Bayesian equilibrium involves the introduction of "beliefs" for players…
- …and can itself be refined: for example, via the intuitive criterion.

Games with no Proper Subgames

Recall the following game. It has no proper subgames. All Nash equilibria are subgame perfect.

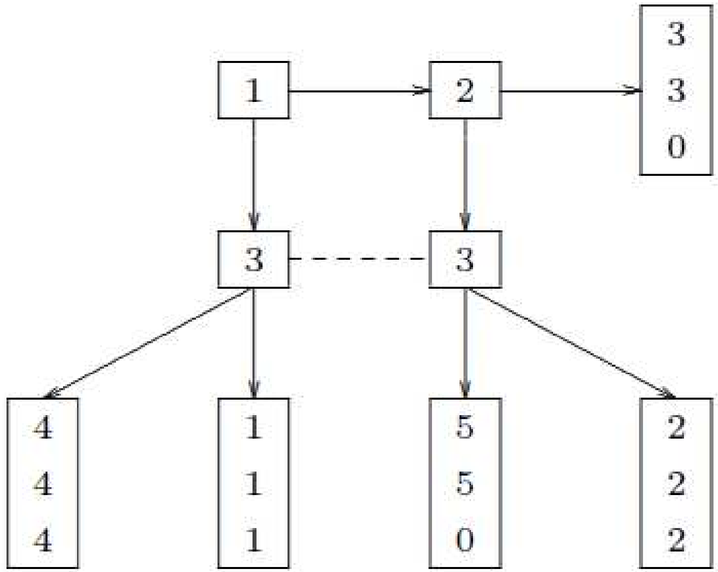

If player 3 is called upon to move, cannot observe what player 1 has done, and hence whether or not player 2 has had a move. But 3's optimal action depends upon what has happened.

- Does modifying the extensive form help make a prediction? If not…
- What does player 3 believe has happened? Will this help predict?

C. Sequential Rationality and Out-of-Equilibrium Beliefs

A *system of beliefs* is a probability distribution over nodes, which gives the relative likelihoods of being at each node in an information set, conditional on having reached it (MWG Definition 9.C.1).

Beliefs

Recall Definition 23: a strategy for player $i \in N$ is a function mapping each history where i is called upon to move (that is, $P(h) = i$) to an action $A(h)$. Recall that $A(h) = A(h')$ if $h, h' \in I_i$.

Add to this the notion of a *belief.* Each player assigns probabilities to each possible history within a particular information set, and does so for every one of their information sets.

Definition 26. Beliefs for $i \in N$ are probability distributions $\mu_{i,I_i} : I_i \mapsto [0, 1]$ for each $I_i \in \mathcal{I}_i$.

An *assessment* or *prospect* is a strategy profile and a set of beliefs for each player (s, μ).

Rather than equilibria being particular strategy profiles, they will now be particular assessments.

Beliefs are part of any equilibrium.

These may include beliefs over information sets that are never reached. It is vitally important, when calculating and writing down equilibria, to remember this.

Applying Beliefs

Defining player 3′s beliefs in the example from earlier (a version of "Selten′s Horse") yields:

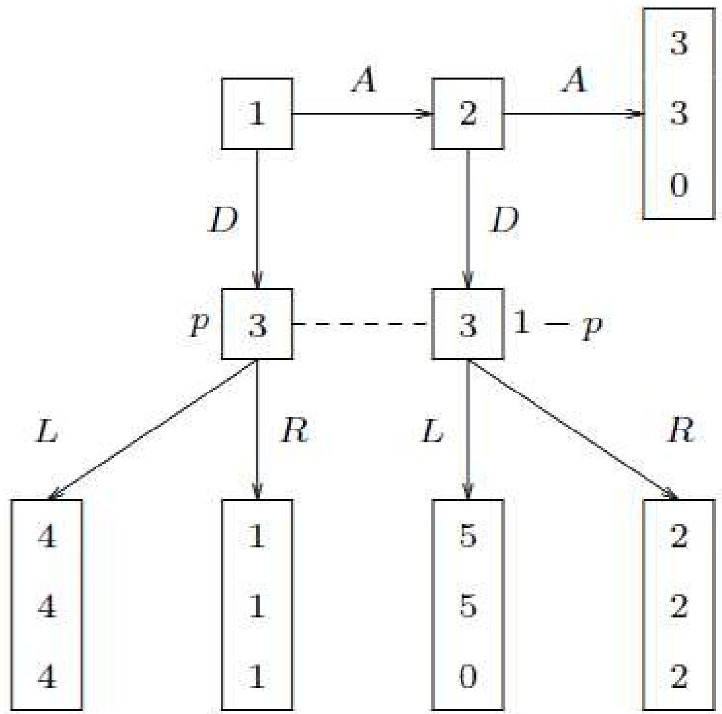

A strategy profile is an action for each of the three players, e.g. $\{A, A, L\}$.

Beliefs are given by the probability distribution $(p, 1 - p)$ over player 3′s information set.

Optimality and Beliefs

Suppose the belief of player 3 that history (D) has occurred is set at p.

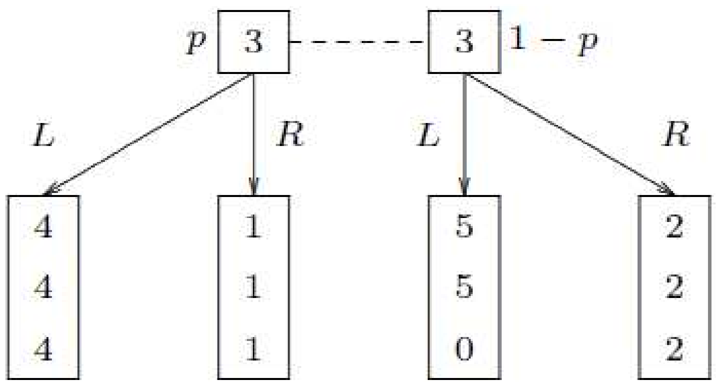

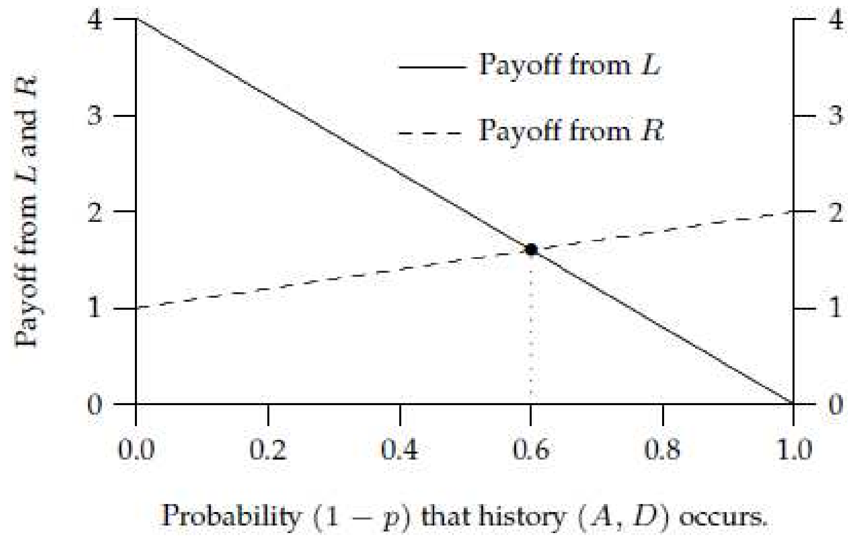

The payoffs for L and R are $L = 4p$ and $R = p + 2(1-p)$. L is a best response whenever:

$$L \geq R \quad \Leftrightarrow \quad 4p \geq p + 2(1-p) \quad \Leftrightarrow \quad p \geq \frac{2}{5} \quad \left[\text{or } 1 - p \leq \frac{3}{5}\right].$$

Inconsistent Beliefs

Set $p > \frac{2}{5}$. Compute best responses, player 3 strictly prefers L. So…

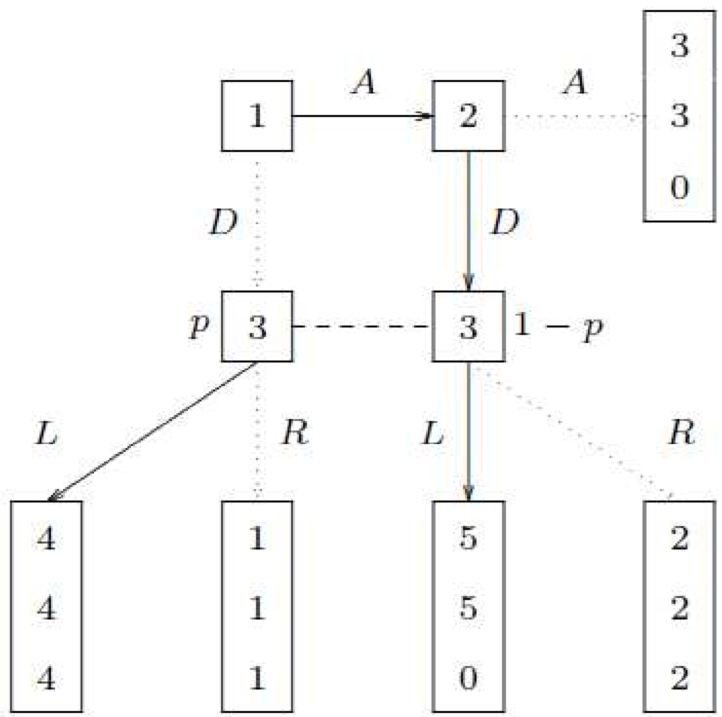

Player 2 will play D, so player 1 will play A. But now the belief p is inconsistent with the play of the game: history (A, D) occurs with probability one, not with probability $1 - p < \frac{3}{5}$.

Consistent Beliefs

Set $p \leq \frac{2}{5}$. Now player 3 (weakly) prefers to play R, so...

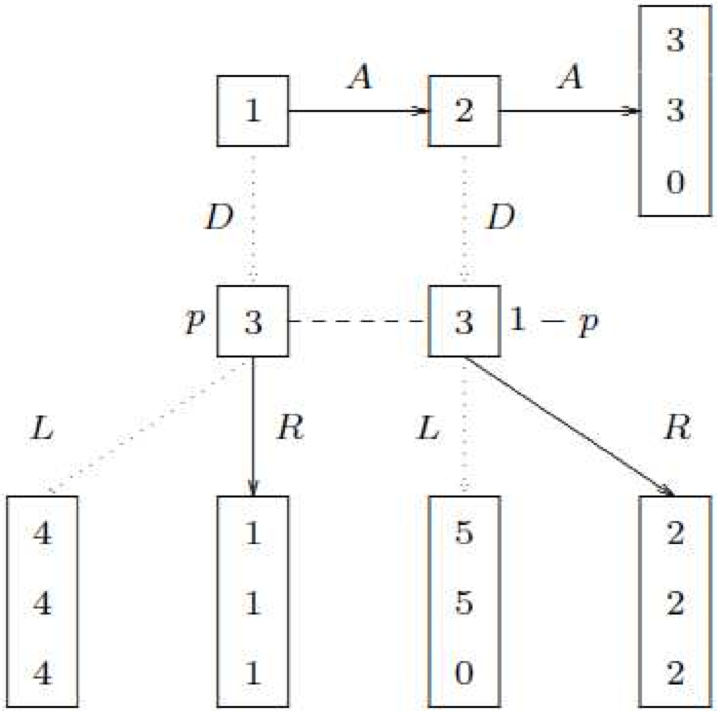

Player 2 will play A, and player 1 will play A. The belief p is *not* inconsistent with the play of the game. It is not however, uniquely defined. For any such p there is a "perfect-Bayesian equilibrium".

A strategy profile is *sequentially rational* at an information set if no player can do better, given his beliefs about what has happened so far, by changing his strategy (MWG Definition 9.C.2). Generalizes notion of sequential rationality to games like MWG Example 9.C.1, where subgame-perfect equilibrium does not capture idea of sequential rationality.

A strategy profile and system of beliefs is a *weak perfect Bayesian equilibrium* if the strategy profile is sequentially rational given the beliefs, and the beliefs are derived from the strategy profile using Bayes' Rule whenever possible (MWG Definition 9.C.3, MWG Example 9.C.1).

("Weak" because the definition is completely agnostic about zero-probability updating.)

Perfect-Bayesian Equilibrium

Bringing together ideas of the optimality of strategies *given* beliefs and the consistency of beliefs *given* equilibrium play motivates the notion of perfect-Bayesian equilibrium.

Definition 26. A (weak) perfect-Bayesian equilibrium is a strategy profile (see definition 22) and a set of beliefs for each player (see definition 25), (s^*, μ^*) such that:

1. At every information set I_i player i's strategy maximises their payoff, given the actions of all the other players, and player i's beliefs.
2. At information sets reached with positive probability when s^* is played, beliefs are formed according to s^* and Bayes' rule when necessary.
3. At information sets that are reached with probability zero when s^* is played, beliefs may be arbitrary but must be formed according to Bayes' rule when possible.

Intuitively: optimal actions given beliefs and consistent beliefs in equilibrium.

Note. A formal definition will not be given here, it takes a little too long. See Fudenberg and Tirole (1991), pp. 331-333.

Theorem: A strategy profile is an equilibrium in an extensive form game if and only if there exists a system of beliefs such that the profile is sequentially rational given the beliefs at all information sets that have positive probability of being reached by the profile; and those beliefs are derived from the profile using Bayes' Rule whenever possible, i.e. except for events that have zero probability in the equilibrium (MWG Proposition 9.C.1).

A strategy profile and system of beliefs is a *sequential equilibrium* if the profile is sequentially rational given the beliefs, and there exists a sequence of completely mixed strategies converging to the profile, such that the beliefs are the limit of beliefs derived using Bayes' Rule from the totally mixed strategies (MWG Definition 9.C.4).

Sequential equilibrium strengthens weak perfect Bayesian equilibrium by requiring more consistency of zero-probability beliefs, adding equilibrium play off equilibrium path.

A sequential equilibrium is trivially a weak perfect Bayesian equilibrium, but not vice versa.

Sequential equilibrium is closely related to *perfect* Bayesian equilibrium (MWG 452).

Theorem: A sequential equilibrium is subgame-perfect, but not vice versa. (MWG Proposition 9.C.2, MWG Example 9.C.1.)

Example: Milgrom and Roberts' (1982 *Econometrica*) Model of Informational Entry Deterrence (Kreps 463-480, Figure 13.2 at Kreps 473).

Two expected-profit maximizing firms, Incumbent and (potential) Entrant, choose Quantities, perfect substitutes, I in both of two periods, E only in second period.

I has two possible unit costs, constant across periods, which only it observes: \$3 with probability ρ and \$1 with probability 1- ρ.

E's unit cost is certain and commonly known by both to be \$3.

Both firms have fixed costs of \$3.

ρ and the rest of the structure are common knowledge.

(Example is typical in having private information only one level below the top; but method can handle more general information structures, which however tend to look contrived.)

In the first period, I observes its unit cost c and chooses Q, which determines P = 9 - Q.

In the second period, E observes the first-period P and chooses whether or not to enter.

If E enters, I and E are Cournot competitors in the second period, taking into account whatever information is revealed in equilibrium by I's first-period P.

If E stays out, I is a monopolist in the second period.

The analysis is hard because the privately informed I plays an active role. I's first-period actions can signal its type to E, and in equilibrium both I and E must weigh the indirect, informational payoff implications of I's first-period decisions against their direct effects.

First analyze the Cournot subgame following entry, given E's beliefs.

If E assesses that c = 3 has probability μ, the Cournot equilibrium is $Q_E = 2(2+\mu)/3$, $Q_I|(c = 1) = (10 - \mu)/3$, $Q_I|(c = 3) = (7 - \mu)/3$, with $\pi_E = 4(2 + \mu)2/9$, not including its fixed cost of 3.

Thus E enters iff $4(2 + \mu)2/9 > 3$, or $\mu > 0.598$. E.g., if E knows c = 3, I and E each set $Q_i = 2$ and get $\pi_i = 1$ (= 4 - 3), so it's profitable to enter. If E knows c = 1, I sets $Q_I = 10/3$ and E sets $Q_E = 4/3$ and gets $\pi_E = -11/9$, so it's not profitable to enter.

Now consider I's first-period decision. The first-period monopoly optimum is Q = 4, P = 5, π = 13 if c = 1; Q = 3, P = 6, π = 6 if c = 3.

However, there is no weak perfect Bayesian equilibrium in which each type of I chooses its monopoly optimum in the first period.

For in such an equilibrium, E could infer I's type by observing P, and would enter if P = 6, believing that c = 3. But then the high-cost type of I would get π = 6 in the first period and π = 1 in the second, less over the two periods than the π = 5 and π = 6 it could get (in the hypothesized equilibrium) by switching to P = 5 and thereby preventing E from entering.

The conclusion that there is no equilibrium of this kind does not depend on zero-probability inferences, and therefore holds for weak perfect Bayesian equilibrium or any stronger notion. Only one type needs to want to defect to break the equilibrium, and this is enough to invalidate it as a prediction even if that type is not realized. (Why?)

Now consider whether there can be a weak perfect Bayesian pooling equilibrium, in which both types of I charge the same price with probability one, and are therefore not distinguishable in equilibrium.

(Looking for each possible kind of equilibrium like this is a characteristic form of analysis.)

If $\rho < 0.598$, there is a sequential (and weak perfect Bayesian) equilibrium in which:

(i) each type of I sets $P = 5$ in the first period;

(ii) E sticks with its prior belief $\rho < 0.598$ and therefore stays out if $P \leq 5$ (in any weak perfect Bayesian pooling equilibrium, E must stick with its prior on the equilibrium path);

(iii) E infers that I's costs are high and enters if $P > 5$; and

(iv) entry leads to the Cournot equilibrium with E believing (as common knowledge) that I's costs are high.

In this pooling equilibrium, the high-cost I "hides behind" the low-cost I by giving up some first-period profit to mimic a low-cost I; and both types of I successfully forestall entry.

To see that these strategies and beliefs are consistent with sequential equilibrium, note that:

(i) E's strategy is sequentially rational, given its beliefs;

(ii) the beliefs are consistent with Bayes' Rule on the equilibrium path;

(iii) when c = 1, I charges its favorite first-period price and prevents entry, the best of all possible worlds for I; and

(iv) when c = 3, the only way I could do better is by raising P above 5, but this would cause E to enter and thereby lower total profits.

(Assuming the most pessimistic conjectures about consequences of deviations from equilibrium is a characteristic strategy for identifying the largest possible set of candidates for a weak perfect Bayesian equilibria.)

Note that the beliefs used here also satisfy a natural monotonicity restriction, in that a higher P never lowers E's estimate that I's costs are high.

If $\rho > 0.598$, there is no weak perfect Bayesian pooling equilibrium.

For such an equilibrium would always lead to entry, making a high-cost I unwilling to charge other than its first-period optimal monopoly price.

A low-cost I would prefer a different price, even if it didn't prevent entry.

However, if $\rho > 0.598$ (or in fact for any ρ) there is a separating (screening, sorting) sequential (hence weak perfect Bayesian) equilibrium in which:

(i) a high-cost I charges its optimal monopoly price, 6, in the first period;

(ii) a low-cost I charges 3.76 in the first period;

(iii) E infers that costs are high if $P > 3.76$ and therefore enters;

(iv) E infers that costs are low if $P \leq 3.76$ and therefore stays out;

(v) both types of I charge their monopoly price in the second period if there is no entry; and

(vi) entry leads to the Cournot equilibrium with E believing (as common knowledge) that I's costs are high.

In this separating equilibrium, a low-cost I successfully distinguishes itself from a high-cost I by distorting its first-period price enough to prevent a high-cost I from mimicking it.

Entry occurs exactly when it would with complete information, and the only effect of incomplete information is the distortion of the low-cost I's first-period price, which benefits consumers and hurts the low-cost I.

That the presence of alternative "bad" types hurts "good" types is typical.

To see that these strategies and beliefs are consistent with sequential equilibrium, note that:
(i) E's strategy is again sequentially rational, given the hypothesized beliefs;

(ii) the beliefs are (trivially) consistent with Bayes' Rule on the equilibrium path (and again monotonic);

(iii) a low-cost I would like to set $P > 3.76$ in the first-period, but that would lead to entry and reduce total profits (easy to check); and

(iv) a high-cost I gets $\pi = 6$ in the first period and $\pi = 1$ following entry in the second, just above what it would get by setting $P \leq 3.76$ and forestalling entry (3.76 was chosen to make it just too costly for the high-cost I to mimic the low-cost I in this equilibrium).

This didn’t depend on ρ, so this is a weak perfect Bayesian equilibrium for any ρ.

D. Reasonable Beliefs, Forward Induction, and Normal Form Refinements

Forward induction is a refinement that restricts beliefs to those that reflect plausible inferences from players' past decisions, which often corresponds to a particular kind of iterated *weak* dominance in the normal form (MWG Figure 9.D.1).

Battle of the Sexes Revisited

"Two M.Phil. students need to meet up (again) to discuss their love for economics. They can meet in either the pub or the cafe. The first likes coffee, and prefers the cafe. The other is a big fan of beer — and prefers the pub. They would both rather meet (wherever it may be) than miss each other. The first student, however, always has the option not to go at all. A deep passion for childish PlayStation-type activities leads them to obtain a bigger payoff staying at home than they would receive from going to the pub."

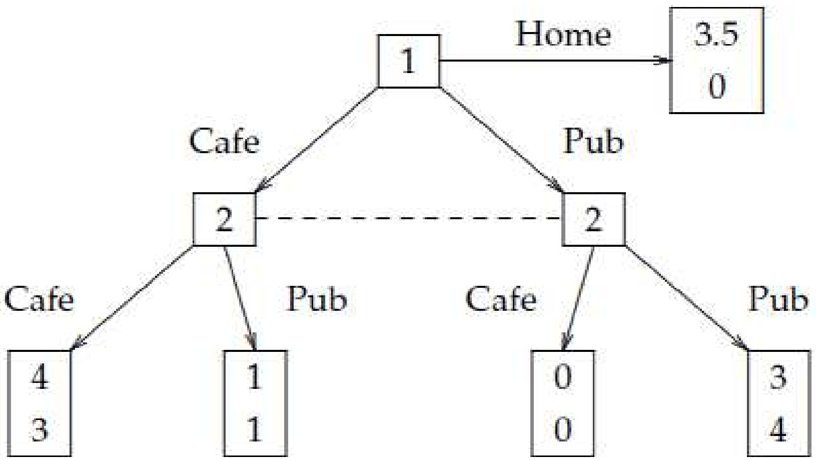

Multiple Subgame-Perfect Equilibria

This game has no proper subgames, so all Nash are subgame perfect. Modify extensive form:

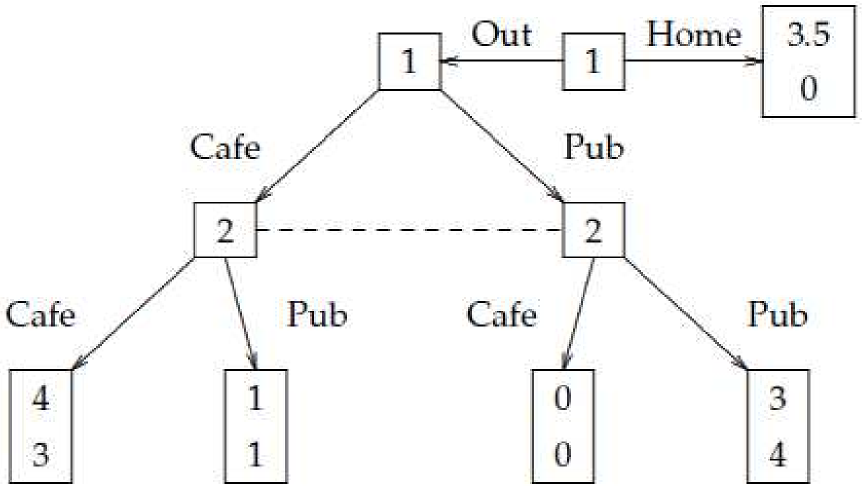

$\Longrightarrow$

	Cafe	Pub
Home, Pub	3.5 0	3.5 0
Home, Cafe	3.5 0	3.5 0
Out, Cafe	4 3	1 1
Out, Pub	0 0	3 4

This modification still has multiple pure-strategy subgame-perfect equilibria...

- Two (pure) Nash equilibria of the subgame starting at (Out): {Cafe, Cafe} and {Pub, Pub}.
- Every strategy combination of the subgame at (Home) is a Nash equilibrium.
- Thus {(Out, Cafe), Cafe} and {(Home, Pub), Pub} are subgame perfect.

Forward Induction

But is there something fishy about the equilibrium $\{$(Home, Pub), Pub$\}$? Why would player 1 ever choose to go to the Pub? They could do better by staying at Home.

- Hence, given that they choose Out, surely they intend to choose Cafe?
- $\{$(Out, Cafe), Cafe$\}$ is the only equilibrium that survives *forward induction*.
- This requires a consideration of player 2's *beliefs* about player 1's action.

Equilibria that survive forward induction will survive the iterated deletion of weakly dominated strategies (again, the order of deletion can matter).

- Notice that (Out, Pub) is *strictly* dominated by (Home, Pub) for player 1.
- In the reduced game (Pub) is *weakly* dominated by (Cafe) for player 2.
- (Home, Pub) and (Home, Cafe) are now strictly dominated for player 1.

More formally: forward induction requires an equilibrium to remain an equilibrium even when strategies dominated in that equilibrium are removed from the game, and this procedure is iterated.

Chapter 12. Market Power
A. Introduction
B. Monopoly Pricing

Monopoly

Consider a single firm in an industry. It faces a demand curve $x(p)$, and so will choose p via...

$$\max_p \pi(p) = \max_p \Big\{ px(p) - c(x(p)) \Big\} \quad \Leftrightarrow \quad \max_q \pi(q) = \max_q \Big\{ p(q)q - c(q) \Big\}$$

At an optimal quantity, $q^* > 0$ therefore, a first-order condition holds: $p'(q^*)q^* + p(q^*) = c'(q^*)$.

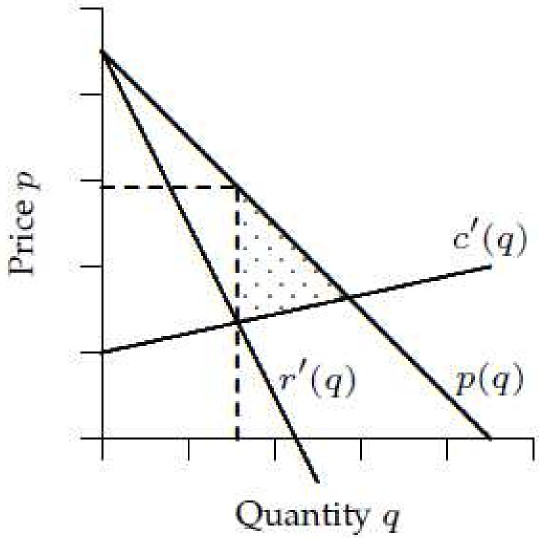

- Marginal revenue is $r'(q) = p'(q)q + p(q)$.
- Marginal cost is $c'(q)$. Optimum $\Rightarrow c'(q^*) = r'(q^*)$.
- Demand is downward sloping $\Rightarrow p'(\cdot) < 0$. So

$$p(q^*) > c'(q^*)$$

- Recall perfectly competitive price $p^+ = c'(q^+)$.
- So $q^* < q^+$ and $p^* > p^+ \Rightarrow$ deadweight loss.

C. Static Models of Oligopoly: Bertrand, Cournot, product differentiation

Importance of separating assumptions about structure and behavior. Different "solution concepts" as equilibrium or subgame-perfect equilibrium in different games.

Static Oligopoly

In the remainder of this lecture, the game-theoretic ideas and concepts already introduced are applied to the standard industrial organisation models of static oligopoly. For example:

Pricing. Two firms selling identical products must choose their prices. The firm with the lower price gains the entire market, but firms would rather charge high prices.

Production. Two profit-maximising firms must choose the scale of their output. Increasing output increases sales, but depresses the market price (which affects both firms).

Investment. Two firms choose investment levels. The firm with the higher investment wins the market, but firms would like to invest as little as is necessary to do so.

The concepts from the previous two lectures are applicable to these games…

1. Represent such interactions as strategic-form games.
2. Iteratively delete strictly dominated strategies where possible.
3. Construct best-reply functions for the players (and plot them).
4. Find pure-strategy and mixed-strategy Nash equilibria.

Theorem: Bertrand duopoly with constant returns to scale, perfectly substitutable goods: Simultaneous price choices by firms yields competitive outcome as unique equilibrium (MWG Proposition 12.C.1).

Bertrand Competition

"Two firms selling identical products must simultaneously choose what price to charge. The firm that charges the lower price gains the entire market, but firms would rather charge high prices. A group of consumers will only buy if the price is less than $\bar{p}$. For simplicity, and without loss of generality, the marginal cost of production is zero."

Players. Two firms labelled $i \in N = \{1, 2\}$.

Strategies. Player i chooses price $p_i \in [0, \infty)$.

Payoffs. Payoffs are profits. There is a unit mass of consumers. If $p_1 = p_2$ the market is split 50:50.

$$\pi_i = \begin{cases} p_i & p_i < \min\{\bar{p}, p_j\}, \\ p_i/2 & p_i = p_j < \bar{p}, \\ 0 & p_i \geq \bar{p} \text{ and/or } p_i > p_j. \end{cases}$$

Bertrand-Nash Equilibrium

- There is a unique pure-strategy Nash equilibrium at $p_1 = p_2 = 0$:
  - If the lowest price were negative, then that firm will make a loss.
  - If the lowest price were strictly positive, then opponent should undercut.
  - If one price is zero, e.g. $0 = p_i < p_j$ then firm i should raise its price.
  - Hence only possibility is $p_1 = p_2 = 0$, where there is no better reply.
- Notice that best-reply functions are not well-defined everywhere:
  - Suppose, for example, that $0 < p_j < \overline{p}$.
  - Always a best-reply for player i to undercut player j.
  - But if player i undercuts by ε, then ε as small as possible without $\varepsilon = 0$.
  - Mathematically, the set of feasible payoffs is open above, cannot attain a maximum.
- The Bertrand specification is *degenerate* — owing to the discontinuity in payoffs.
- Be careful when using continuous action sets!

Theorem: Cournot duopoly with constant returns to scale, perfectly substitutable goods: Simultaneous quantity choices by firms yields equilibrium (not necessarily unique) with prices between competitive and monopoly prices (MWG Proposition 12.C.2, Example 12.C.1).

Cournot Competition with Linear Demand

"Two profit-maximising firms simultaneously choose production quantities of a homogeneous good. Market price is decreasing in total quantity Q, with linear demand, so that $p = a - bQ$. There are constant unit production costs of c for each firm."

Players. Two firms labelled $i \in \{1, 2\}$.

Strategies. Player 1 chooses quantity $x \in [0, \infty)$ and player 2 chooses quantity $y \in [0, \infty)$.

Payoffs. Payoffs are profits. That is, for players 1 and 2 respectively:

$$\pi_1 = x\,[a - b(x + y) - c] \quad \text{and} \quad \pi_2 = y\,[a - b(x + y) - c]\,.$$

1. Fix firm 2's strategy. Calculate a best reply for firm 1, yielding a best-reply function.
2. Fix firm 1's strategy. Calculate a best reply for firm 2: a second best-reply function.
3. Combine the two best-reply functions. Solve to find a Nash equilibrium.

Cournot Best-Reply Functions

Fixing y, profits for player 1 are $\pi_1 = x\,[a - b(x+y) - c]$. This is strictly concave in x, so can calculate first-order conditions for a solution:

$$\frac{\partial \pi_1}{\partial x} = [a - b(x+y) - c] - bx = a - 2bx - by - c = 0.$$

Rearrange this to obtain: $2bx = a - by - c$ which implies $B_1(y) = (a - by - c)/2b$.

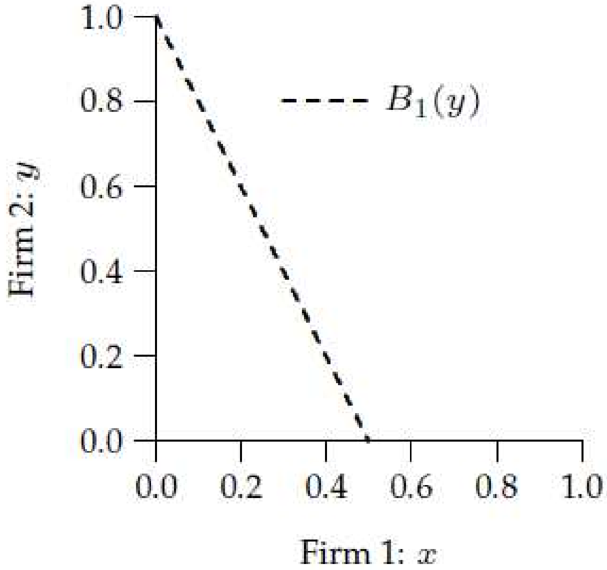

- Plot of reaction function for a=b=1, c=0.
- This is downward sloping:

$$\frac{\partial B_1(y)}{\partial y} = -\frac{1}{2} < 0$$

- Quantities are *strategic substitutes*.
- This is a *submodular* game.

Cournot-Nash Equilibrium

At a Nash equilibrium, players mutually best-respond: $y = B_2(x)$ and $x = B_1(y)$. So,

$$x = \frac{a - by - c}{2b} \quad \text{and} \quad y = \frac{a - bx - c}{2b}.$$

Solve these two equations simultaneously. The solution will be symmetric since the first-order conditions (e.g. $bx = a - b(x + y) - c = a - bQ - c$) depend only upon Q.

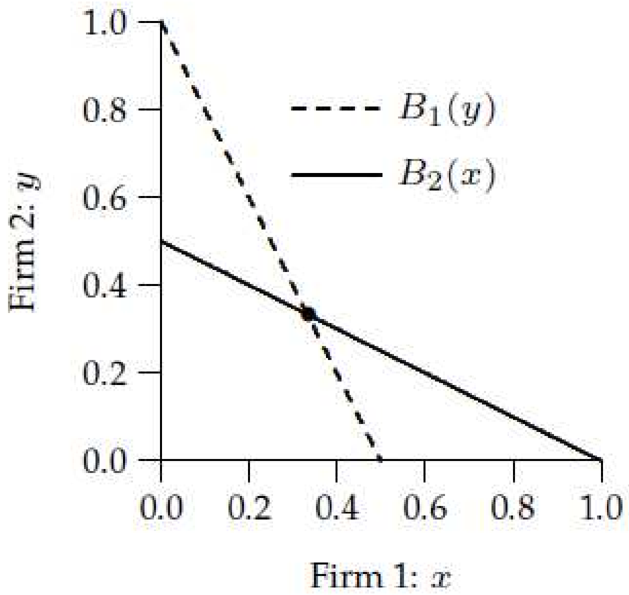

- From symmetry: $x = (a - bx - c)/2b$.
- Multiplying up: $2bx = a - bx - c$.
- Adding a term to both sides: $3bx = a - c$.
- So, finally: $x^* = y^* = (a - c)/3b$.
- Plot the B_i functions and equilibrium.

Iterative Deletion of Strictly Dominated Strategies

The Cournot-Nash equilibrium strategies are the only survivors from the iterated deletion of strictly dominated strategies. To see this, simplify the game $a = b = 1$ and $c = 0$, so $\pi_1 = x(1 - x - y)$.

Consider the strategies $x \in (\frac{1}{2}, \infty)$. These are strictly dominated by $x = \frac{1}{2}$. The profits from playing these strategies are respectively, $x(1 - x - y)$ and $\frac{1}{2}(\frac{1}{2} - y)$. Suppose to the contrary

$$x(1 - x - y) > (1/2 - y)/2 \quad \Leftrightarrow \quad (1/2 - x)y > 1/4 - x + x^2.$$

Since $y \geq 0$, and $x \in (\frac{1}{2}, \infty)$, the left-hand side is less-than-or-equal-to zero. The right-hand side is minimised at zero when $x = \frac{1}{2}$, and therefore is positive, a contradiction. Similarly for $y > \frac{1}{2}$.

Now consider $x \in [0, \frac{1}{4})$. These strategies are strictly dominated by $x = \frac{1}{4}$. The payoffs are $x(1 - x - y)$ and $\frac{1}{4}(\frac{3}{4} - y)$. Suppose again, to the contrary, that

$$x(1 - x - y) > (3/4 - y)/4 \quad \Leftrightarrow \quad (1/4 - x)y > 3/16 - x + x^2.$$

Now $y \leq \frac{1}{2}$, so this is true $\forall y$ if and only if $(\frac{1}{4} - x)\frac{1}{2} > \frac{3}{16} - x + x^2 \Leftrightarrow 0 > \frac{1}{16} - \frac{1}{2}x + x^2$. Which is a contradiction. This process continues until $x = y = \frac{1}{3}$ remains.

Equilibrium in General Cournot Games

- General: n firms, firm i has constant marginal cost c_i, inverse demand $P(Q)$.
- Maximise profits for firm i. If $P(Q) < c_i$ then $q_i = 0$. Otherwise:

$$\pi_i = q_i[P(Q) - c_i] \quad \Rightarrow \quad \frac{\partial \pi_i}{\partial q_i} = P(Q) - c_i + q_i P'(Q) = 0 \quad \Leftrightarrow \quad q_i = -\frac{P(Q) - c_i}{P'(Q)}.$$

- Individual quantities are uniquely defined by industry supply Q.
- Thus, if $c_i = c$ for all i, then any equilibrium is symmetric.
- Sum the first-order conditions for all n firms, divide by n to obtain:

$$\frac{nP(Q) - [\sum_{i=1}^{n} c_i]}{P(Q)} + \frac{QP'(Q)}{P(Q)} = 0 \quad \Leftrightarrow \quad \frac{P(Q) - \frac{1}{n}\sum_{i=1}^{n} c_i}{P(Q)} = \frac{1}{n\varepsilon}$$

- Hence outcome determined by industry-average of marginal cost.
- In games where there is a single "state variable" (here, Q) determining equilibria...
- ...the solution boils down to a single fixed-point equation.

Cournot versus Monopoly and Perfect Competition

Let marginal costs be constant and equal to c for every firm. In n-firm Cournot, from previous slide

$$P(Q) + q_i P'(Q) = c \quad \Rightarrow \quad nP(Q) + P'(Q)\sum_i^n q_i = nc \quad \Rightarrow \quad P(Q) + P'(Q)\frac{Q}{n} = c$$

- Since $P'(\cdot) < 0$ (demand is downward sloping), $P(Q) > c$, so $Q < q^+$ (where $P(q^+) = c$).

Competitive industries produce more, and at a lower price.

- Now, suppose the monopoly optimal quantity is q^*. Suppose $q^* > Q$.
- Take a particular firm i, and let firm i increase q_i so that the new industry quantity $= q^*$.
- Joint profits must increase (they are maximised at q^* by definition).
- But aggregate quantity has risen, so price has fallen, so the other firms (who didn't alter their quantities) are worse off. As joint profits have risen, i must be better off $\Rightarrow$ a profitable deviation.
- So $q^* \leq Q$. But $q^* \neq Q$ since above equation can't be satisfied by same Q at $n > 1$ and $n = 1$.

Monopolies produce less, and at a higher price.

The Hotelling Line

Two firms are located at either end of a unit interval $[0, 1]$. A unit mass of consumers is distributed uniformly on the interval. The firms charge p_i for a good produced with constant marginal cost c.

- The cost of buying from firm i is $p_i + td$. t is a transport cost and d is the distance from firm i.
- A particular consumer $z \in [0, 1]$ will buy from i (positioned at 0) if $p_i + tz < p_j + t(1 - z)$.
- The indifferent consumer $\hat{z}$ satisfies $p_i + t\hat{z} = p_j + t(1 - \hat{z})$. Thus

$$\hat{z} = \frac{t + p_j - p_i}{2t}$$

- Assuming $\hat{z} \in [0, 1]$, if i charges p_i and j charges p_j then firm i's demand is given by

$$q_i = \hat{z} = \frac{1}{2} + \frac{p_j - p_i}{2t}$$

- If $\hat{z} > 1$ then $q_i = 1$, and if $\hat{z} < 0$ then $q_i = 0$. (Note this assumes consumers *always* buy one of the products, this can be guaranteed if their valuation for the good is sufficiently high).

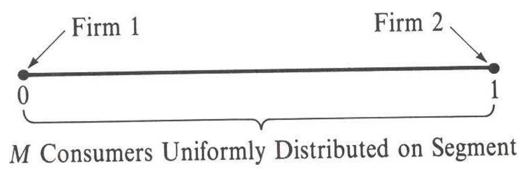

Figure 12.C.5
The linear city.

Differentiated Products

An alternative interpretation: "Two firms selling differentiated products simultaneously choose prices. Total market size is a single unit mass. Suppose that each consumer is willing to pay a large amount to obtain a product. They do not necessarily buy from the cheapest firm, however."

- If $p_j - p_i > t$ then firm i captures the whole market: $q_i = 1$ and $q_j = 0$.
- If $|p_j - p_i| \leq t$ then the split depends on the price difference:

$$q_i = \frac{1}{2} + \frac{p_j - p_i}{2t}.$$

Players. Two firms $N = \{i, j\}$.

Strategies. Player i chooses $p_i \in [0, \infty)$.

Payoffs. If $p_i - p_j > t$, then $\pi_i = 0$. If $p_j - p_i > t$, then $\pi_i = p_i - c$. Otherwise

$$\pi_i = (p_i - c)\left(\frac{1}{2} + \frac{p_j - p_i}{2t}\right).$$

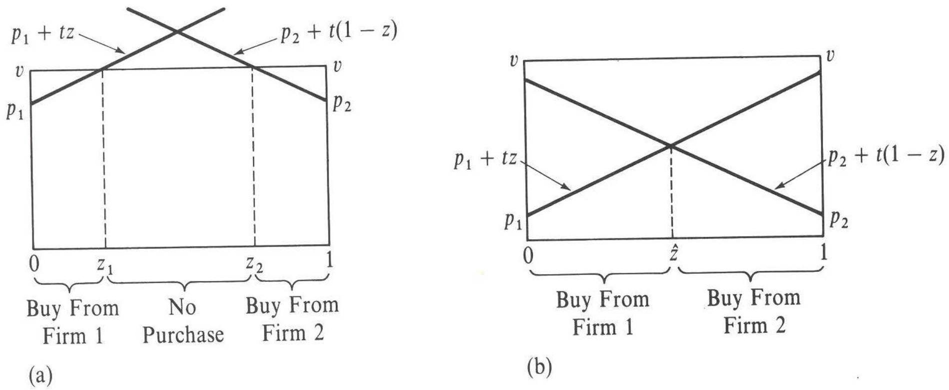

Figure 12.C.6
Consumer purchase decisions given p_1 and p_2. (a) Some consumers do not buy (b) All consumers buy

Differentiated Products Best-Reply Functions

Profit is concave in price. Differentiate to obtain the first-order condition

$$\frac{\partial \pi_i}{\partial p_i} = \frac{t + p_j - 2p_i + c}{2t} = 0.$$

Solving for p_i yields the best-reply function $B_i(p_j) = (t + c + p_j)/2$. This is upward sloping since $\partial B_i/\partial p_j > 0$. Prices are *strategic complements* — the game is *supermodular*.

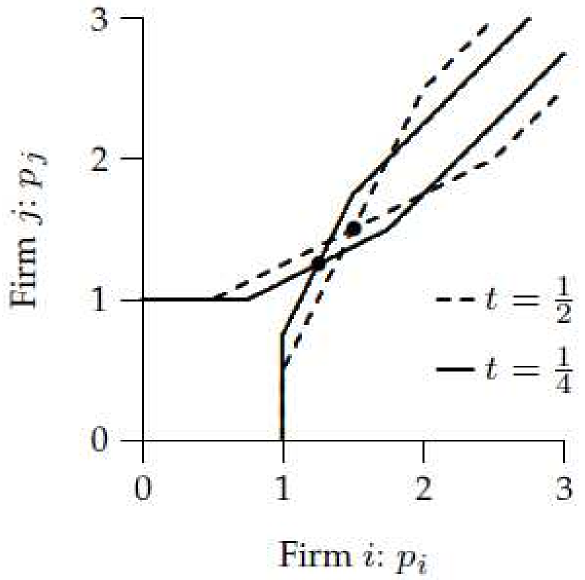

- This solution applies when $|p_j - p_i| \leq t$. In fact,

$$B_i(p_j) = \begin{cases} c & p_j < c - t, \\ (t + c + p_j)/2 & c - t \leq p_j \leq 3t + c, \\ p_j - t & 3t + c < p_j. \end{cases}$$

- For an interior equilibrium, $p_i = (t + c + p_j)/2$.
- Symmetry ensures $p_i = p_j = p^*$.
- So $p^* = (t + c + p^*)/2$, and so $p^* = t + c$.

D. Repeated Interaction

Complete-information repeated games

Define a *repeated game* as dynamic game in which same *stage game* is played over and over again each period by the same players. The stage game could be anything, even another repeated game.

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

(More realistic than assuming an arbitrarily specified endpoint is common knowledge?)

The repeated Prisoner's Dilemma is the canonical (but overworked, not representative) model of using repeated interaction to overcome short-run incentive problems.

The Prisoners' Dilemma Again (and Again)

Recall (a version of) the Prisoners' Dilemma game. Wouldn't it be nice if players could cooperate?

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

Perhaps if the game is played repeatedly, cooperation would be possible in early periods by threatening to defect later if cooperation was not observed? Suppose the game is repeated T times:

- In the last time period (T), D is a dominant strategy for both players.
- In period $T - 1$ neither player can influence future decisions so each will play D.
- This logic continues: iterating back to the first period, both players play D throughout.

Thus the unique subgame-perfect equilibrium involves defection in every period, regardless of previous play. This is also the unique Nash equilibrium! Why?

…and Again

Consider a strategy that called upon a player to play C in some period $1 \leq t \leq T$. This is dominated by a strategy that is identical everywhere except that it calls for D at period t.

Continuing (and tightening) this iterated dominance argument yields a unique Nash equilibrium.

However, this depends upon the *uniqueness* of the equilibrium in the one-shot game… repeating a game in this manner will usually expand the set of achievable outcomes in the *stage game*.

	L	M	R
T	0 0	3 4	6 0
M	4 3	0 0	0 0
B	0 6	0 0	5 5

In the above game $\{B, R\}$ is not a Nash equilibrium of the stage game, but it can be played in the first period, as part of a subgame-perfect equilibrium, if the game is played twice.

Repeated Games with Discounting

A stage game is a (mixed extension of the) strategic-form game $\mathcal{G} = \langle N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N}\rangle$.

Definition 27. The repeated game with discounting of the stage-game $\mathcal{G}$ is the extensive-form game $\Gamma = \langle N, H, P, \{U_i\}_{i\in N}\rangle$ with,

1. $H = \cup_{t=0}^{T} S^t$ (where $S^0 = \emptyset$ is the initial history and T is the number of *stages*).
2. $P(h) = N$ for each non-terminal history $h \in H$.
3. Payoffs involve a discount factor, $\delta \in (0, 1)$, and are such that

$$U_i = \sum_{t=1}^{T} \delta^{t-1} u_i(s^t),$$

where $s^t \in S$ is the strategy profile of the game $\mathcal{G}$ played in stage $t \in \{1, \ldots, T\}$.

For infinitely repeated games let $T = \infty$ and (normalising to per-period payoffs)

$$U_i = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(s^t).$$

The Stage Game

So: how do players play $\{B, R\}$ in equilibrium in the following game when repeated twice?

	L	M	R
T	0 0	3 4	6 0
M	4 3	0 0	0 0
B	0 6	0 0	5 5

The pure-strategy Nash equilibria of the stage game are $\{T, M\}$ and $\{M, L\}$.

Consider a mixed-strategy σ for row player placing probability p on T and $1-p$ on M. Row gets:

	L	M	R
σ	$4(1-p)$	$3p$	$6p$
B	0	0	5

This shows row's payoffs from the mixed-strategy σ and from the pure-strategy B. Thus σ dominates B, with $1 > p > \frac{5}{6}$. By symmetry, R is dominated for column player.

The Reduced Stage Game

Deleting strictly dominated strategies yields the following reduced game:

	$L\ (y)$	$M\ (1-y)$	Expected
$T\ (x)$	0 0	3 4	$3(1-y)$
$M\ (1-x)$	4 3	0 0	$4y$
Expected	$3(1-x)$	$4x$	

Note that there are three Nash equilibria of this reduced game, two pure and one mixed:

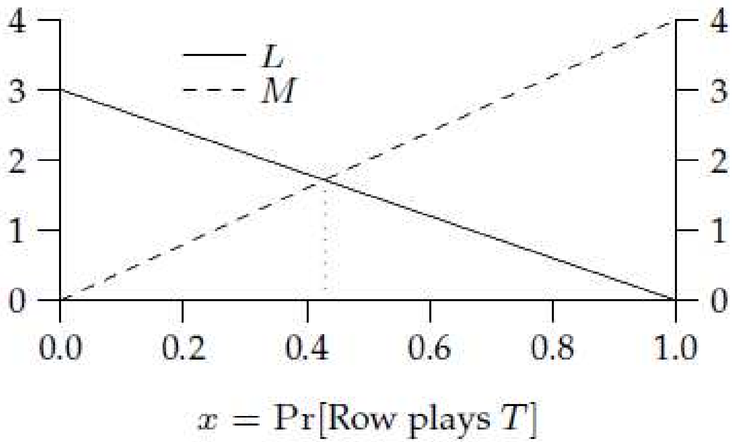

$$\Rightarrow \quad Z = \begin{cases} \{T, M\} & \Rightarrow u = (3, 4) \\ \{M, L\} & \Rightarrow u = (4, 3) \\ \{\frac{3}{7}, \frac{3}{7}\} & \Rightarrow u = (\frac{12}{7}, \frac{12}{7}) \end{cases}$$

The strategy profile $\{B, R\}$ Pareto dominates each of these equilibria. How can this be obtained?

Repeating the Stage Game: Conditional Strategies

The game is played twice, with discount rate $\delta < 1$. Suppose $\{B, R\}$ is played in the first period.

	L		M		R	
T	0	0	3	4	6	0
M	4	3	0	0	0	0
B	0	6	0	0	5	5

$$\Rightarrow \quad \text{with } Z = \begin{cases} \{T, M\} & \Rightarrow u = (3, 4) \\ \{M, L\} & \Rightarrow u = (4, 3) \\ \{\frac{3}{7}, \frac{3}{7}\} & \Rightarrow u = (\frac{12}{7}, \frac{12}{7}) \end{cases}$$

Conditional Strategy: if $\{B, R\}$ is observed in the first period, row plays M and column plays L.

If anything else is observed, row player places probability $\frac{3}{7}$ on T and $\frac{4}{7}$ on M. Column player places probability $\frac{3}{7}$ on L and $\frac{4}{7}$ on M. These strategies constitute a subgame-perfect equilibrium.

- In the second period (nine subgames) play corresponds to a Nash equilibrium.
- In the first period (a single subgame), by following the above strategy, row gets $5 + 4\delta$.
- Column gets $5 + 3\delta$. By deviating, the greatest column could get would be $6 + \frac{12}{7}\delta$.

This is a (subgame-perfect) equilibrium so long as: $\quad 5 + 3\delta \geq 6 + \frac{12}{7}\delta \quad$ or: $\quad \delta \geq \frac{7}{9}$.

Another Example

Consider the below game, played twice with a discount factor $\delta < 1$.

	L		M		R	
T	2	2	0	0	6	0
M	0	0	4	4	0	0
B	0	6	0	0	5	5

Note that $\{B, R\}$ is not a Nash equilibrium of the stage game. Nevertheless, consider:

- Play $\{B, R\}$ in the first period. If $\{B, R\}$ observed, play $\{M, M\}$ in the second.
- If $\{B, R\}$ is not observed in the first period, play $\{T, L\}$.
- Payoff from playing strategy is $5 + 4\delta$.
- Payoff from deviating is at most $6 + 2\delta$.

If $\delta \geq \frac{1}{2}$ these strategy profiles are a subgame-perfect equilibrium.

Equilibria in Repeated Games

Suppose this game is repeated three times. How many subgames are there? Showing that a given strategy profile is a subgame-perfect equilibrium might seem like a daunting task. However…

The **One-Deviation Principle** will simplify this task.

Additionally in repeated games, the number of subgame-perfect equilibria might be very large indeed. Characterising them would once again seem like a *very* daunting task. However…

The **Folk Theorems** will simplify this task.

The rest of the lecture introduces these ideas and applies them to some examples.

The One-Deviation Principle

Definition 28. A strategy for player i satisfies the one-deviation principle (or property) if for any history $h \in H$ such that $i \in P(h)$, there is no deviation that i could make to increase their payoff whilst leaving all the other players' strategies fixed, *and the rest of their own strategy*.

> "A strategy profile in a finite-horizon extensive-form game or in an infinitely repeated game with discount factor $\delta < 1$ is a subgame-perfect equilibrium if and only each player's strategy satisfies the one-deviation principle."

- In other words, only need to check one-deviation-at-a-time at every stage for each player.
- Ignore multiple contemporaneous deviations or multiple sequential deviations.

	L	M	R
T	2 2	0 0	6 0
M	0 0	4 4	0 0
B	0 6	0 0	5 5

$(\times\, 3)\ldots$

Thrice-Repeated Example

1. Consider the strategy: Play B, then B if $\{B, R\}$ is observed and T if anything else, then M if $(\{B, R\}, \{B, R\})$ is observed and T if anything else (for row player).

2. Suppose column player is playing R, then R if $\{B, R\}$ is observed and L if anything else, then M if $(\{B, R\}, \{B, R\})$ is observed and L if anything else.

3. Need only check single deviations for each player at each stage. No profitable deviations in the last stage — if $\{B, R\}$, $\{B, R\}$ is observed M is a best response, and if not, T is a best response.

4. At penultimate stage, B yields a higher payoff if $5 + 4\delta \geq 6 + 2\delta$ (or $\delta \geq \frac{1}{2}$). This second payoff arises from the most profitable single deviation possible at this stage.

5. At initial stage, B yields a higher payoff if $5 + 5\delta + 4\delta^2 \geq 6 + 2\delta + 2\delta^2$ (or $\delta \geq 0.28$ish). This second payoff arises from the most profitable single deviation possible at this stage.

6. Symmetry implies column player's strategy is also subgame perfect for $\delta \geq \frac{1}{2}$.

No need to check any other deviant strategies: e.g. Play B *then* $T \ldots$

Multiple Equilibria in Repeated Games

There are many such subgame-perfect equilibria in this repeated game. In fact, when games are repeated, many many different outcomes can be supported as equilibria. e.g.

- Playing the same Nash equilibrium in every stage is always subgame perfect.
- Playing any sequential combination of Nash equilibria is subgame perfect.
- Conditioning the future Nash equilibrium to be played on current choices…
- …allows non-Nash strategies to be part of subgame-perfect equilibria.

Rather than characterise all these equilibria, characterise the (normalised one-period) *payoffs* that are achievable as part of a Nash (and subgame-perfect) equilibrium.

- In the finitely repeated Prisoners' Dilemma, this argument did not work…
- …because there is a *unique* Nash equilibrium of the stage game and hence…
- …there is no choice of equilibria with which to condition behaviour.
- In the infinitely repeated game, however, there is a multiplicity of equilibria.

Characterising a Game in Payoff Space

Consider the Prisoners' Dilemma. Plotting row player's payoffs against column player's payoffs, and allowing players to mix, which payoffs are achievable in the one-shot game?

	C	D
C	3 3	5 0
D	0 5	1 1

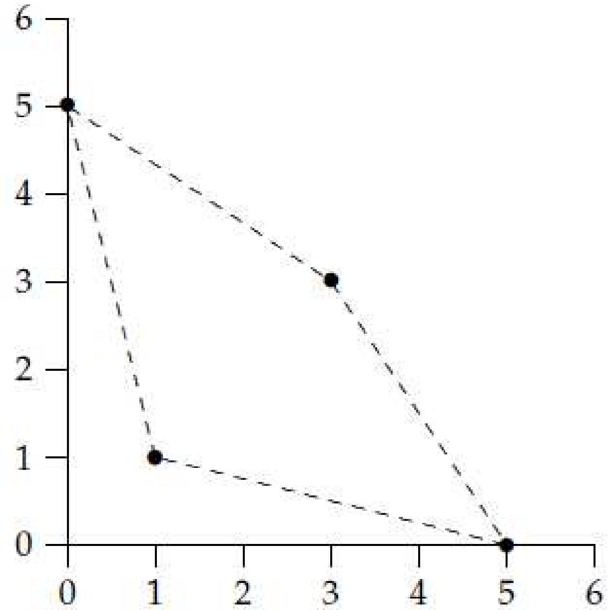

Appropriate mixtures over the two strategies generate all the payoffs inside the diamond. This is the *convex hull* of the payoffs to pure strategies. All these payoffs are *feasible*.

Which of these are supportable as part of an equilibrium in an infinitely-repeated game?

The Nash-Threats Folk Theorem

"Every feasible payoff profile above the Nash equilibrium payoff profile can be achieved by a subgame-perfect equilibrium of the infinitely-repeated game for δ large enough."

	C	D
C	3, 3	0, $\underline{5}$
D	$\underline{5}$, 0	$\underline{1}$, $\underline{1}$

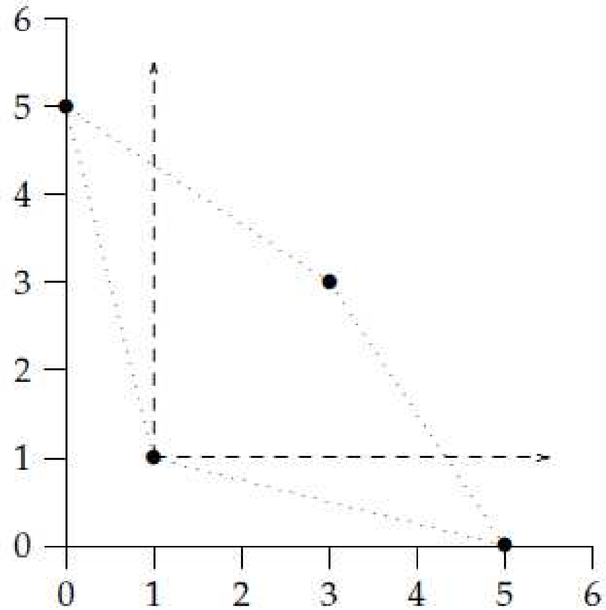

How? Play strategies that generate required payoff combination. If anyone deviates, play the stage Nash equilibrium forever. Consider payoff profile $(3, 3)$ in the Prisoners' Dilemma. . .

Cooperation and Collusion

Compare the payoff stream from $\{C, C\}$ forever with the payoff stream from a single deviation in any subgame (using one-deviation principle and noting all subgames look the same):

$$3 + 3\delta + 3\delta^2 + \ldots = \frac{3}{1-\delta} \geq 4 + \frac{1}{1-\delta} = 5 + \delta + \delta^2 + \ldots$$

This requires $\delta \geq \frac{1}{2}$. The "grim" strategies (play C forever unless D is ever observed, in which case play D forever) constitute a subgame-perfect equilibrium. Logic carries over to collusion...

Consider an infinitely-repeated n-firm Bertrand pricing game. Charge the monopoly price with profits π_M/n. Any deviation prompts marginal-cost pricing for T periods.

$$\frac{\pi_M}{n}\sum_{t=1}^{\infty}\delta^{t-1} = \frac{1}{1-\delta}\frac{\pi_M}{n} \quad \text{versus} \quad \pi_M + \frac{\delta^{T+1}}{1-\delta}\frac{\pi_M}{n}.$$

The former is greater than the latter if $1 - \delta^{T+1} \geq n(1-\delta)$, which holds certainly if δ and T are large enough (in the grim strategy, T is infinite, and $\delta \geq 1 - \frac{1}{n}$ is the appropriate condition).

There are also many asymmetric subgame-perfect equilibria:

E.g. suppose the implicit contract is: Row alternates between C and D and Column always chooses C. This continues until either deviates, after which both choose D from then on.

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 $\delta \geq 0.59$ (approximately), so the asymmetric implicit contract is consistent with subgame-perfect equilibrium as long as $\delta \geq 0.59$.

Column does worse than Row but the threat is symmetric, so supporting Column's strategy as part of a subgame-perfect equilibrium is harder than supporting Row's.

The limit is higher than for the symmetric implicit contract because the asymmetry makes it harder to keep *both* players willing to stay with the implicit contract.

Infinite-horizon repeated games have an enormous multiplicity of equilibria, both of equilibrium outcomes and the threats that can be used to support them (which in this noiseless version of the game never need to be carried out on the equilibrium path).

We've seen one symmetric and one asymmetric efficient equilibrium of the repeated Prisoner's Dilemma. Folk Theorems are useful because they give limits on what kinds of implicit contracts can be supported as subgame-perfect equilibria in repeated games.

Even More Equilibria

Consider the following game, with its associated payoff representation.

	C	D
C	$\underline{4}$, $\underline{4}$	$\underline{3}$, 0
D	0, $\underline{3}$	1, 1

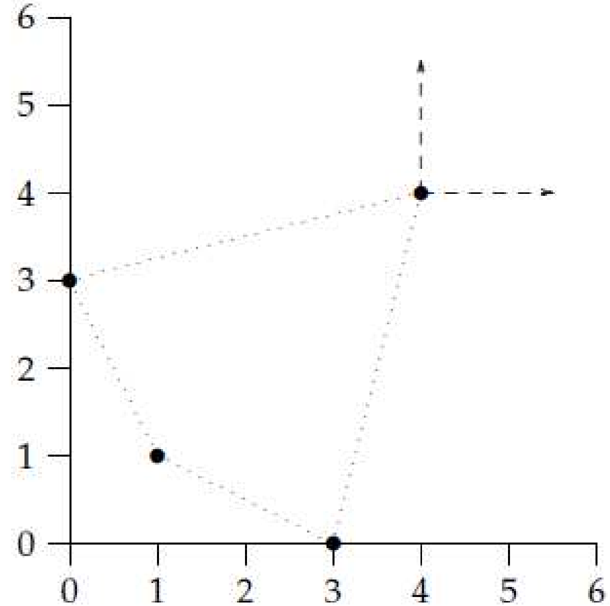

The Nash-threat Folk Theorem only indicates that $(4, 4)$ can be achieved. But many other payoffs can be achieved by subgame-perfect equilibria also. To do this, need to define "minmax" payoffs:

The lowest payoff one player can force the other player to, *given* the other player best-responds.

Minmax Payoffs and another Folk Theorem

Player i's minmax payoff is given by: $u_i^m = \min_{s_{-i}} \left\{ \max_{s_i} u_i(s) \right\}$. The folk theorem says:

> "Every feasible payoff profile above the minmax payoff profile can be achieved by a subgame-perfect equilibrium of the infinitely-repeated game for δ large enough." (There is a technical *full dimensionality* condition.)

	C	D
C	$\underline{4}$, $\underline{4}$	$\underline{3}$, 0
D	0, $\underline{3}$	1, 1

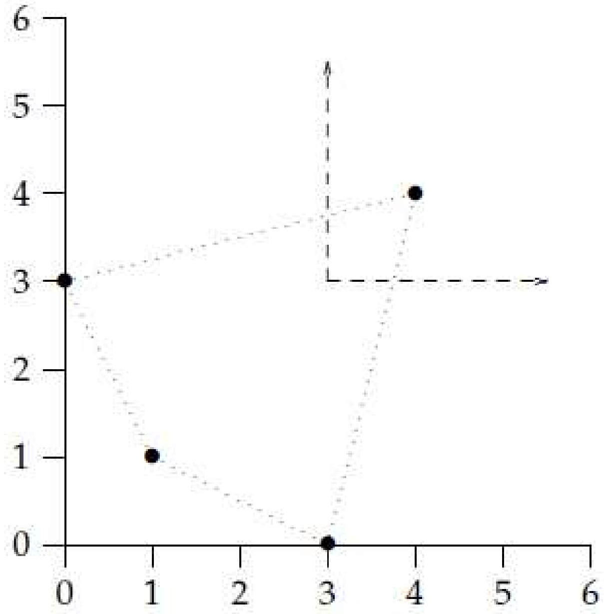

Folk Theorem: In an infinitely repeated game with complete information and observable strategies, for any feasible pair of payoffs strictly greater than those that follow from repeating players' minimax payoffs in the stage game, there is a discount factor such that for all greater discount factors, those payoffs arise in a subgame-perfect equilibrium of the repeated game (MWG Proposition 12.AA.5, Example 12.AA.1).

Easy to prove for stage games like Prisoner's Dilemma, where Nash reversion is minimax.

Harder to prove for other stage games. Temporal convexification with high discount factors. See MWG Proposition 12.AA.5.

Another Example

Consider the following game. What are the minmax payoffs? What are the subgame-perfect equilibrium payoffs achievable in the infinitely-repeated version?

	L	R
T	-1 3	2 -1
M	2 -1	-1 3
B	1 2	1 2

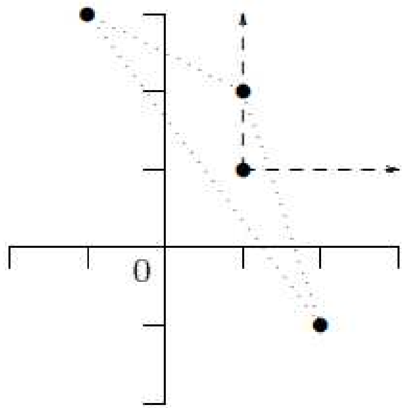

Minmax payoffs are the worst thing that others can do to a player, whilst that player best responds.

- Minmax payoff for the row player is 1, e.g. by column playing $L\&R$ with probability $\frac{1}{2}$ each.
- Minmax for the column player is 1, e.g. play T with probability $\frac{1}{2}$ and M with probability $\frac{1}{2}$.

So: every feasible payoff combination above $(1,1)$ is the outcome of a subgame-perfect equilibrium of the infinitely-repeated game. Requires complex "punishment and reward" strategies.

Applications to oligopoly

Theorem: With a sufficiently high discount factor, the monopoly price can be supported as a subgame-perfect equilibrium outcome in an infinitely repeated Bertrand duopoly by threats to revert forever to the competitive price if anyone deviates (MWG Proposition 12.D.1).

Theorem: With a sufficiently high discount factor, any price from the competitive to the monopoly price can be supported as a subgame-perfect equilibrium outcome in an infinitely repeated Bertrand duopoly by threats to revert forever to the competitive price if anyone deviates. For low discount factors, only the competitive price can be supported (MWG Proposition 12.D.2).

Expanding the number of firms in the Bertrand model shrinks the set of implicit contracts supportable via the Folk Theorem by making the limit on δ more stringent (MWG 405).

Implicit collusion in infinitely repeated Cournot duopoly: Supporting zero payoffs via strategies that yield zero-profit quantities followed forever by the monopoly output until someone deviates (Kreps 524-526, MWG Example12.AA.1).

A Model of Collusion with Incomplete Information

Even when an opponent's actions are not directly observable, similar results are available. The following example illustrates (formal statements are beyond the scope of the course):

- Recall the repeated Bertrand model introduced briefly earlier.
- Suppose now that the price set by one firm is not observable by the other.
- Strategies are conditioned on some publicly-observed state variable, e.g. market demand.

Suppose that with probability $(1-\alpha)$, demand follows from the standard Bertrand model, but with probability α, demand falls to zero (owing to conditions outside the firms' control). Strategies:

- Play p_M (monopoly price) and split the market as long demand is non-zero.
- Play p_C (marginal cost) for T periods if zero demand is observed by either firm.
- Note the firm that cheated (set a price $p < p_M$) will know that the other firm has observed zero demand in that period! Hence both firms know when a "price war" will start.

This last feature is critical to the ensuing analysis. First, calculate payoffs. . .

Calculating Payoffs

Start by considering the payoffs to a player that accrue when players collude and charge p_M. Denote expected payoffs in the collusive phase and in a price war phase as V_C and V_W respectively. Then,

$$V_C = \underbrace{(1-\alpha)\left[\frac{\pi_M}{2} + \delta V_C\right]}_{\text{High Demand}\Rightarrow\text{Retain C Phase}} + \underbrace{\alpha\delta V_W}_{\text{Enter War}} \qquad \text{and} \qquad V_W = \underbrace{\delta^T V_C}_{\text{Wait T Periods for C Phase}},$$

where π_M is monopoly profit. A deviating player would obtain at most V_D, where

$$V_D = (1-\alpha)\left[\pi_M + \delta V_W\right] + \alpha\delta V_W.$$

So if the following inequality obtains, these strategies will form an equilibrium:

$$(1-\alpha)\left[\pi_M + \delta V_W\right] + \alpha\delta V_W \leq (1-\alpha)\left[\frac{\pi_M}{2} + \delta V_C\right] + \alpha\delta V_W.$$

Supporting Collusion

Subtracting $\alpha\delta V_W$ from both sides, and dividing by $(1-\alpha)$, this condition becomes

$$\pi_M + \delta V_W \leq \frac{\pi_M}{2} + \delta V_C \qquad \text{or} \qquad \frac{\pi_M}{2} \leq \delta(V_C - V_W).$$

Solving simultaneously for V_C and V_W from the initial equalities on the previous slide yields

$$V_C = \frac{(1-\alpha)\pi_M/2}{1-(1-\alpha)\delta - \alpha\delta^{T+1}} \qquad \text{and} \qquad V_W = \frac{(1-\alpha)\delta^T \pi_M/2}{1-(1-\alpha)\delta - \alpha\delta^{T+1}}.$$

Finally, substituting into the inequality above gives the condition

$$2(1-\alpha)\delta - (1-2\alpha)\delta^{T+1} \geq 1.$$

This is satisfied for appropriate values of δ and α given a large enough T, (i.e. for α small, δ large and T large). So Folk Theorems are available under incomplete information.

E. Entry and Dynamic Oligopoly

Dynamic Oligopoly

Apply the notion of subgame-perfect equilibrium to a variety of dynamic oligopoly models.

- **Stackelberg Leadership**. The simplest possible application: one firm (the *leader*) chooses a quantity to produce first, and a second (the *follower*) chooses a quantity to produce next.
- **Simple Models of Entry**. Firms simultaneously choose whether to pay a fixed cost of entry, and afterwards they simultaneously decide how much to produce, or what price to charge.
- **Strategic Pre-commitment**. One firm (the *incumbent*) has the option to make a pre-competition strategic investment, after which it competes (in prices or in quantities) with another firm.
- **Location Choice**. Two firms simultaneously choose the extent to which their products are differentiated (a position on the Hotelling line) and then they simultaneously compete in prices.
- **Entry Deterrence**. The incumbent chooses a level of pre-competition investment. A second firm (the *entrant*) chooses whether to enter the industry, and then (quantity) competition takes place.

irreversible decisions that affect future interaction is probably the most important of the three views of entry deterrence in the literature. The others two are reputation in repeated games, and informational as in the Milgrom-Weber entry deterrence model.

Stackelberg Leadership

"A firm (the leader) chooses a quantity, observed by a second (the follower), which then sets its quantity. Price is determined by the total quantity produced, and profits accrue accordingly."

Players. Two firms labelled $i \in \{l, f\}$.

Histories. Firm l chooses a quantity $x \in \mathcal{R}^+$ and then player 2 chooses quantity $y \in \mathcal{R}^+$. The set of histories is $H = \{(\emptyset), (0), \ldots, (x), \ldots, (0,0), \ldots, (x,y), \ldots\} = \emptyset \cup \mathcal{R}^+ \cup (\mathcal{R}^+)^2$.

Player Function. $P(\emptyset) = l$ and $P(x) = f$ for all $x \in \mathcal{R}^+$.

Payoffs. Payoffs are profits. Suppose demand is linear. That is, for players l and f respectively:

$$\pi_l = x\left[a - b(x+y) - c\right] \quad \text{and} \quad \pi_f = y\left[a - b(x+y) - c\right].$$

- Representing this game in strategic form game yields strategies $s_l = x$ and $s_f = y(x)$.
- The follower's strategy maps every choice that the leader could make onto a quantity for f.

Nash and Subgame-Perfect Equilibria

There are *many* Nash equilibria: any strategy profile $s_l = \hat{x}$ and $s_f = \hat{y}(x)$ where $\hat{y}(\hat{x})$ is a best reply to $\hat{x}$ can be supported as a Nash equilibrium by suitable choice of $\hat{y}(x)$ for $x \neq \hat{x}$.

Only one of these is "credible" (subgame perfect): this requires a best-reply in *every* subgame…

Following every x, firm f must play a best reply: thus $y(x) = B_f(x)$. Recall

$$\frac{\partial \pi_f}{\partial y} = 0 \quad \Rightarrow \quad B_f(x) = \frac{a - bx - c}{2b}$$

Firm l plays a best reply in the first period: chooses the optimal x given firm f plays $B_f(x)$. So

$$\max_x x\,[a - b(x + B_f(x)) - c] \quad \Rightarrow \quad x^* = \frac{a-c}{2} \quad \text{and} \quad s_{\text{SPE}} = (x^*, B_f(x))$$

Note that firm l's profits *must* rise relative to Cournot… l could always choose the Cournot quantity. In this linear and strategic-substitutes example above, f's profits fall (not generally true).

Entry

"A large number of symmetric firms simultaneously decide whether to pay an entry fee F to enter the market. In the second period, those that entered compete in quantities (or prices)."

- A subgame-perfect equilibrium requires Nash play in every subgame.
- In the final period, suppose n firms have entered, they must play a Nash equilibrium.
- Suppose profits (excluding the sunk entry fee) from the second period are given by π_n.
- Consider an equilibrium value of n, call it n^*. There are firms that enter, and some that don't.
  - If firm i is an entrant then this is only profitable if $\pi_{n^*} \geq F$.
  - If another firm j did not enter it must be the case that $\pi_{n^*+1} < F$.
- It is reasonable to assume that π_n is decreasing in n. In this case there is a unique n^* that satisfies these two inequalities. Note that the equilibrium is not unique—which firms enter?

Suppose that firms compete in prices (Bertrand) in the second period. Then $\pi_n = 0$ whenever $n \geq 2$. So, if $F > 0$, the equilibrium must involve $n^* = 1$ (so long as monopoly profits, $\pi_1 = \pi^m > F$)!

Now consider a Cournot example of this game…

Cournot Competition and Entry

Suppose in the second period, the n firms compete in quantities with demand given by $P(Q)$. Assume the firms have identical and constant marginal costs c. Recall

$$\pi = q[P(Q) - c] \quad \Rightarrow \quad P(Q) + P'(Q)\frac{Q}{n} = c$$

Now for the linear case, $P(Q) = a - bQ$ and $P'(Q) = -b$, so, substituting in,

$$a - b\left[1 + \frac{1}{n}\right] Q = c \quad \Rightarrow \quad q = \frac{Q}{n} = \frac{1}{b}\left(\frac{a-c}{n+1}\right) \quad \text{so} \quad \pi_n = \frac{1}{b}\left(\frac{a-c}{n+1}\right)^2$$

Firms will enter in the first period so long as $\pi_n \geq F$, and won't if $\pi_{n+1} < F$. Therefore

$$n^* = \lfloor \hat{n} \rfloor \quad \text{where} \quad \hat{n} = \frac{a-c}{\sqrt{bF}} - 1$$

As F falls, the equilibrium number of entrants increases; and Q approaches competitive output.

Notation. $\lfloor n \rfloor$ is the largest integer less or equal to n.

Entry and Welfare

Welfare is the sum of consumer surplus and total firm profits (including the entry fee F). Consider n symmetric firms producing a total quantity of Q at a price $P(Q)$. Then

$$W(n) = \underbrace{\int_0^Q P(X)dX - QP(Q)}_{\text{consumer surplus}} + \underbrace{n\pi_n - nF}_{\text{total profits}} \quad \Rightarrow \quad W(n) = \int_0^Q P(X)dX - Qc - nF$$

The problem is to choose n to maximise $W(n)$. Now total quantity is given by $Q = nq(n)$ where $q(n)$ is each firm's optimal quantity when there are n entrants. Thus

$$\frac{dQ}{dn} = nq'(n) + q(n) \quad \text{so} \quad \frac{dW(n)}{dn} = 0 \quad \Leftrightarrow \quad \left[nq'(n) + q(n)\right](P(Q) - c) = F$$

The welfare-maximising n is different from the number choosing to enter the industry: ignoring integer problems, n^* was determined by $\pi_n = q(n)[P(Q) - c] = F$ — the first term is absent!

In Bertrand there is clearly (in general) under-entry, in Cournot...

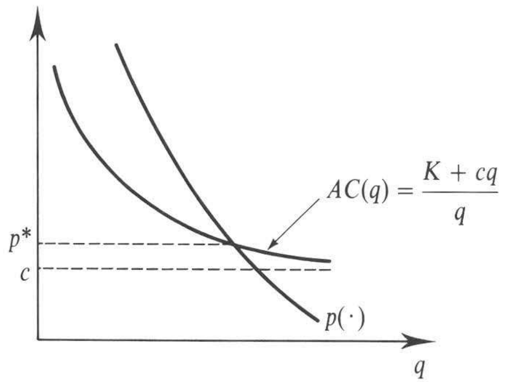

Figure 12.E.3
Equilibrium in the one-stage entry game discussed in Example 12.E.4.

Cournot Competition and Over-Entry

Using the simple linear specification of the earlier slides and lecture, $q(n)$ is given by

$$q(n) = \frac{1}{b}\left(\frac{a-c}{n+1}\right) \quad \text{so} \quad q'(n) = -\frac{1}{b}\frac{a-c}{(n+1)^2}$$

Again, ignore integer issues. Welfare $W(n)$ is maximised where n satisfies

$$\left[nq'(n) + q(n)\right](P(Q) - c) = \frac{1}{b}\frac{a-c}{(n+1)^2}\left(\frac{a-c}{n+1}\right) = \frac{1}{b}\frac{(a-c)^2}{(n+1)^3} = F$$

But recall that the equilibrium number of entrants satisfies $(n^* + 1)^2 = (a-c)^2/bF$, so

$$(n^* + 1) = (n+1)^{3/2} \quad \Rightarrow \quad n^* > n$$

The equilibrium number of entrants exceeds the socially optimal number, there is over-entry.

This is a result of the *business-stealing* effect: entering firms do not care that some of their profits come from taking sales away from other firms, but no social welfare is generated by such activity.

G. Strategic Precommitments to Affect Future Competition

Strategic Pre-commitment

Consider the following general two-period game between two firms (1 and 2).

Period 1. Firm 1 chooses $k \in \mathcal{R}$. This is a strategic investment which alters its own best-reply function in the second period, as well as its profits directly.

Period 2. Firm 1 and firm 2 engage in competition (e.g. price or quantity) by simultaneously choosing strategies $s_1 \in \mathcal{R}^+$ and $s_2 \in \mathcal{R}^+$ respectively. Their profits are given by

$$\pi_1(s_1, s_2, k) \quad \text{and} \quad \pi_2(s_1, s_2) \quad \text{respectively}$$

Equilibrium in the subgame following a choice k is where the best-reply functions cross, that is

$$(s_1^*, s_2^*) \quad \text{where} \quad s_1^* = B_1(s_2^*, k) \quad \text{and} \quad s_2^* = B_2(s_1^*)$$

In the subgame-perfect equilibrium, firm 1's chooses k to maximise its profits:

$$\max_k \pi_1(s_1, s_2, k) \quad \text{subject to} \quad (s_1, s_2) = (s_1^*, s_2^*)$$

Strategic Cost Reduction in Cournot

As an example, consider a situation in which k is an investment in (marginal) cost reduction. The firms then compete in quantities (Cournot). With a very simple form for the cost of investment:

$$\pi_1(q_1, q_2, k) = q_1[P(Q) - c(k)] - k \quad \text{and} \quad \pi_2(q_1, q_2) = q_2[P(Q) - c]$$

Where firm 1′s marginal cost is $c(k)$ with $c'(k) < 0$, and firm 2′s marginal cost is c.

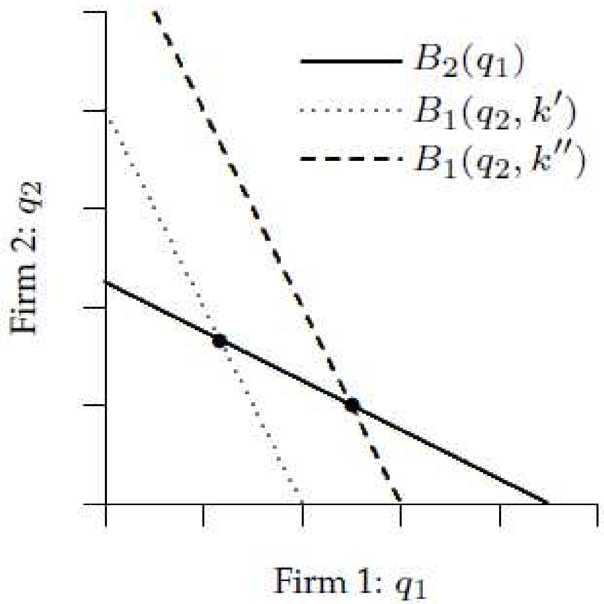

- If price is $P(Q) = 1 - Q$ then firm 1′s best reply is

$$B_1(q_2, k) = \frac{1 - q_2 - c(k)}{2}$$

- Increasing k in the first period reduces $c(k)$ and so…
- …pushes the best-reply function out ($k'' > k'$).
- Equilibrium in second period involves higher q_1…
- …and lower q_2, which is beneficial for firm 1.

Price Competition and Strategic Pre-commitment

The simple example on the previous slide is one of strategic substitutes ($\partial B_i/\partial q_j < 0$). In differentiated-product price competition (for example) $\partial B_i/\partial q_j > 0$: i.e. strategic complements.

$$\pi_1 = (p_1 - c(k))\left(\tfrac{1}{2} + p_2 - p_1\right) - k \quad \text{and} \quad \pi_2 = (p_2 - c)\left(\tfrac{1}{2} + p_1 - p_2\right).$$

(Hotelling profits from lecture 3 with $t = \frac{1}{2}$). Again $c'(k) < 0$ and investment costs are simply k.

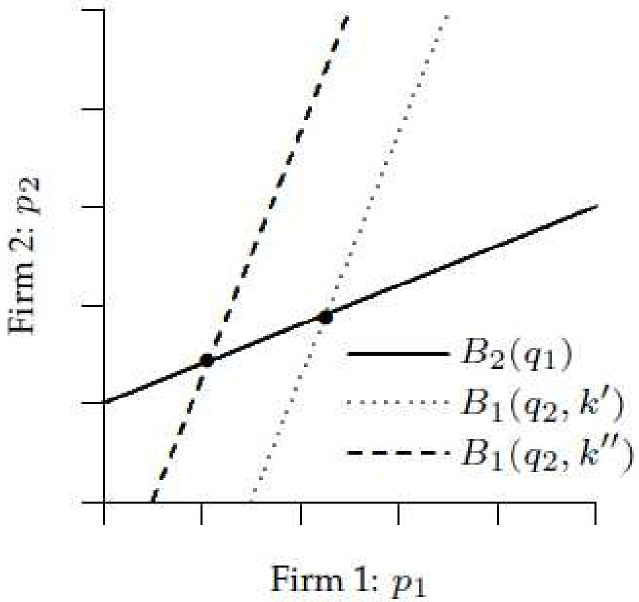

- Now the best-reply function for firm 1 is

$$B_1(p_2, k) = \frac{1}{4} + \frac{p_2 + c(k)}{2}$$

- (For $p_2 \in [c(k) - \frac{1}{2}, c(k) + \frac{3}{2}]$, ignore other p_2).
- Increasing k (from k' to k'') reduces $c(k)$ and…
- …shifts B_1 in, lowering equilibrium p_1 and…
- …also lowering p_2, which is bad for firm 1.

Direct and Strategic Effects

In general, firm 1 must consider the strategic effect on equilibrium prices/quantities in the post-investment game (as well as the direct effect on its e.g. costs) of pre-competition investments.

- Consider the effect on equilibrium profits of a change in investment k:

$$\frac{d\pi_1(s_1^*(k), s_2^*(k), k)}{dk} = \frac{\partial\pi_1(s_1^*, s_2^*, k)}{\partial k} + \frac{\partial\pi_1(s_1^*, s_2^*, k)}{\partial s_1}\frac{ds_1^*(k)}{dk} + \frac{\partial\pi_1(s_1^*, s_2^*, k)}{\partial s_2}\frac{ds_2^*(k)}{dk}$$

- But the second term's first element $\partial\pi_1/\partial s_1$ is zero in equilibrium (firm 1 maximises), so

$$\frac{d\pi_1(s_1^*(k), s_2^*(k), k)}{dk} = \underbrace{\frac{\partial\pi_1(s_1^*, s_2^*, k)}{\partial k}}_{\text{direct effect}} + \underbrace{\frac{\partial\pi_1(s_1^*, s_2^*, k)}{\partial s_2}\frac{ds_2^*(k)}{dk}}_{\text{strategic effect}}$$

- The sign of the strategic effect is determined by whether firm 1's profits increase with firm 2's choice (e.g. in the price competition example, yes; in the quantity competition example, no)...
- ...and the sign of $ds_2^*(k)/dk$, which is determined by the slopes of the best-reply functions...

The Strategic Effect

- To determine the sign of $ds_2^*(k)/dk$, note that in equilibrium $s_2^*(k) = B_2(B_1(s_2^*(k), k))$. Now

$$\frac{ds_2^*(k)}{dk} = \frac{dB_2}{ds_1}\left[\frac{\partial B_1}{\partial s_2}\frac{ds_2^*(k)}{dk} + \frac{\partial B_1}{\partial k}\right]$$

$$\Rightarrow \quad \frac{ds_2^*(k)}{dk} = \left(\frac{dB_2}{ds_1} \times \frac{\partial B_1}{\partial k}\right) \Big/ \left(1 - \frac{dB_2}{ds_1}\frac{\partial B_1}{\partial s_2}\right)$$

- In "stable" games, the denominator is +ve: sequences of best-replies converge to equilibrium, so the product of best-reply slopes is less than 1 (the product is $\frac{1}{4}$ in both previous examples).

$$\text{sign[Strategic Effect]} = \text{sign}\left[\frac{\partial \pi_1}{\partial s_2} \times \frac{dB_2}{ds_1} \times \frac{\partial B_1}{\partial k}\right]$$

  - Would firm 1 like to lower or raise firm 2's strategic choice?
  - How does strategic investment affect firm 1's best-replies?
  - Is firm 2 competing in strategic substitutes or complements?

Note. The firms's best-reply functions can, in general, have different signs—the above argument is unaffected.

Submodular and Supermodular Games

Players. Two players labelled $i \in N = \{1, 2\}$.

Strategies. Player 1 chooses $x \in X \subseteq \mathcal{R}$, player 2 chooses $y \in Y \subseteq \mathcal{R}$.

Payoffs. Payoffs are $u_1(x, y)$ and $u_2(y, x)$, with symmetry $u_1 = u_2 = u$.

Calculate the slope of a player's best-reply function:

$$\begin{aligned} x = B_1(y) \qquad & \Rightarrow \qquad \frac{\partial u(x,y)}{\partial x} = 0, \\ & \Rightarrow \qquad \frac{\partial^2 u(x,y)}{\partial x \partial y} + \frac{\partial^2 u(x,y)}{\partial x^2}\frac{dx}{dy} = 0, \\ & \Rightarrow \qquad \frac{dx}{dy} = -\frac{\partial^2 u(x,y)/\partial x \partial y}{\partial^2 u(x,y)/\partial x^2}. \end{aligned}$$

Denominator is negative from second-order conditions. Sign determined by numerator, i.e.

$$\frac{\partial^2 u(x,y)}{\partial x^2} < 0 \quad \Rightarrow \quad \text{sign}\left\{\frac{dx}{dy}\right\} = \text{sign}\left\{\frac{\partial^2 u(x,y)}{\partial x \partial y}\right\}.$$

Informal Definitions

Definition 13. The game $\mathcal{G} = \langle \{1, 2\}, \{X, Y\}, \{u, u\} \rangle$ is *supermodular* if $\partial^2 u(x, y)/\partial x \partial y > 0$, and so $B_1(y)$ is upward sloping. X and Y are said to be *strategic complements*.

Definition 14. The game $\mathcal{G} = \langle \{1, 2\}, \{X, Y\}, \{u, u\} \rangle$ is *submodular* if $\partial^2 u(x, y)/\partial x \partial y < 0$, and so $B_1(y)$ is downward sloping. X and Y are said to be *strategic substitutes*.

These are only informal definitions. True definitions are a little bit more general. Some examples:

- Cournot competition is typically submodular. Quantities are strategic substitutes: an increase in an opponent's quantity reduces the incentive a firm has to raise quantity.
- Bertrand competition is typically supermodular. Prices are strategic complements: an increase in an opponent's price increases the incentive a firm has to raise price.
- Sometimes best-reply functions can be non-monotonic...

	Strategic Substitutes: $\dfrac{db_2(\cdot)}{ds_1} < 0$	Strategic Complements $\dfrac{db_2(\cdot)}{ds_1} > 0$
$\dfrac{\partial b_1(\cdot)}{\partial k} > 0$	$\dfrac{ds_2^*(k)}{dk} < 0$	$\dfrac{ds_2^*(k)}{dk} > 0$
$\dfrac{\partial b_1(\cdot)}{\partial k} < 0$	$\dfrac{ds_2^*(k)}{dk} > 0$	$\dfrac{ds_2^*(k)}{dk} < 0$

Figure 12.G.1
Determinants of the sign of $ds_2^*(k)/dk$.

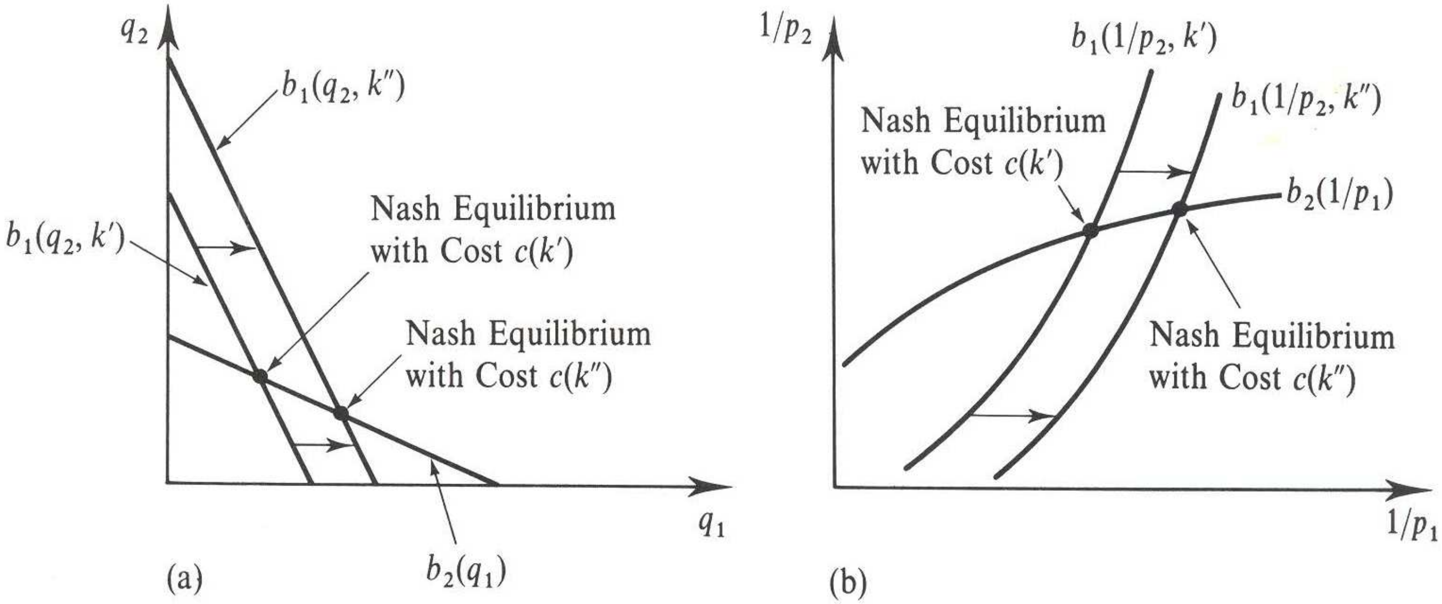

Figure 12.G.2
Strategic effects of a reduction in marginal cost from $c(k')$ to $c(k'') < c(k')$.
(a) Quantity model.
(b) Price model.

An Advertising Game

"Two firms sell a product in a market of fixed size. Suppose that prices are fixed (at 1), but that each firm must choose an advertising budget, denoted by x and y respectively. Advertising is costly, but firms want to obtain a high market share. Advertising is the sole determinant of market share, yielding sales of $x/(x+y)$ and $y/(x+y)$ respectively."

Players. The two firms ($N = \{1, 2\}$).
Strategies. Firm 1 chooses $x \in [0, \infty)$ and Firm 2 chooses $y \in [0, \infty)$.
Payoffs. Profits are given by

$$\pi_1(x, y) = \frac{x}{x+y} - x \quad \text{and} \quad \pi_2(y, x) = \frac{y}{x+y} - y.$$

To calculate best-reply functions, consider Firm 1's maximisation problem:

$$\begin{aligned}
\frac{\partial \pi_1}{\partial x} = \frac{1}{x+y} - \frac{x}{(x+y)^2} - 1 = 0 \qquad &\Leftrightarrow \qquad \frac{1}{x+y} = \frac{x}{(x+y)^2} + 1, \\
&\Leftrightarrow \qquad x + y = x + (x+y)^2, \\
&\Leftrightarrow \qquad \sqrt{y} = x + y.
\end{aligned}$$

Non-Monotonic Best-Replies

This yields the best-reply functions $B_1(y) = \sqrt{y} - y$ and $B_2(x) = \sqrt{x} - x$.

At a Nash equilibrium $y = B_2(x)$ and $x = B_1(y)$. Hence $x^* = y^* = \frac{1}{4}$.

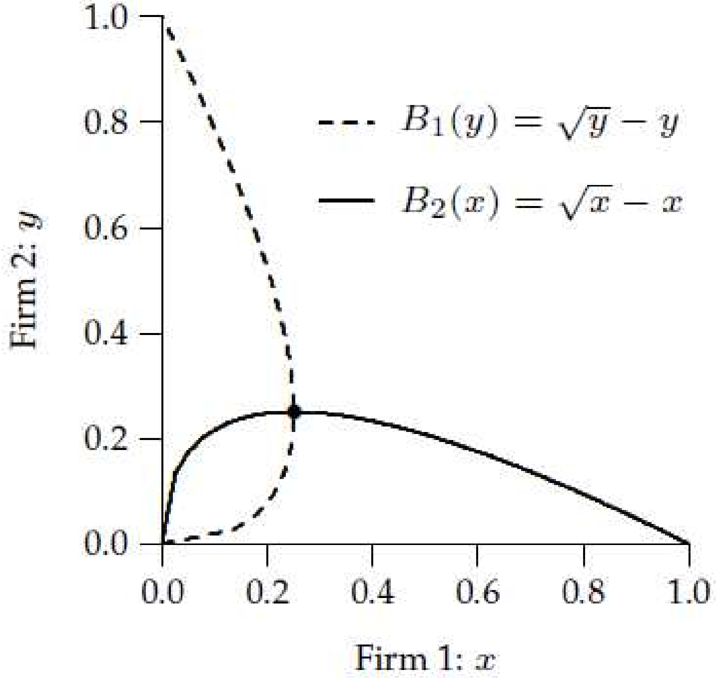

- The best-reply functions for this game slope upward initially, then downward.
- The game is neither sub- nor supermodular.
- Variables are both strategic complements and substitutes, depending on the region.
- Equilibrium at $(0, 0)$? It would seem not…
- Since $\pi_1(0, 0) = \pi_2(0, 0) = \frac{1}{2}$ seems apt.

Note. B_1 and B_2 only make sense for $y > 0$ and $x > 0$. Setting $\pi_1(0, 0) = \pi_2(0, 0) = 1$ generates an equilibrium at $(0, 0)$, but doesn't make much sense. In any case it would not be a very "stable" equilibrium, whereas $(\frac{1}{4}, \frac{1}{4})$ is.

Strategic Entry Deterrence

An important example of a strategic pre-commitment is that of a firm using e.g. capacity as a mechanism to deter (or accommodate) entry by a rival firm. Consider the following 3-stage game:

1. **Pre-Entry Commitment.** Firm 1 (the incumbent) chooses a level of capacity (k). Capacity costs r per unit purchased. The firm can now produce up to k units of the good at marginal cost c.
2. **Entry Decision.** In the second period firm 2 (the entrant) decides whether or not to enter the industry. The entrant must pay an entry fee of f if they decide to enter.
3. **Quantity Competition.** In the final period, if firm 2 has entered, the firms engage in quantity competition; firm 2's marginal cost is $c + r$. If firm 2 didn't enter, firm 1 is a monopolist.

Find the subgame-perfect equilibrium(s) by backward induction in the usual way…

$$B_2(q_1) = B(q_1; c + r) \quad \text{and} \quad B_1(q_2, k) = \min\{B(q_2; c), k\}$$

$B(q; x)$ is the Cournot best-reply function when marginal cost is x and the opponent produces q.

Nash Equilibrium in Stage 3

If the entrant has not entered, the incumbent, firm 1, is a monopolist. Equilibrium is simply to play the monopoly quantity, or (given the capacity constraint k) as close as possible $q_1 = B_1(0, k)$.

In the subgame following entry, the firms compete à la Cournot. To find the equilibrium...

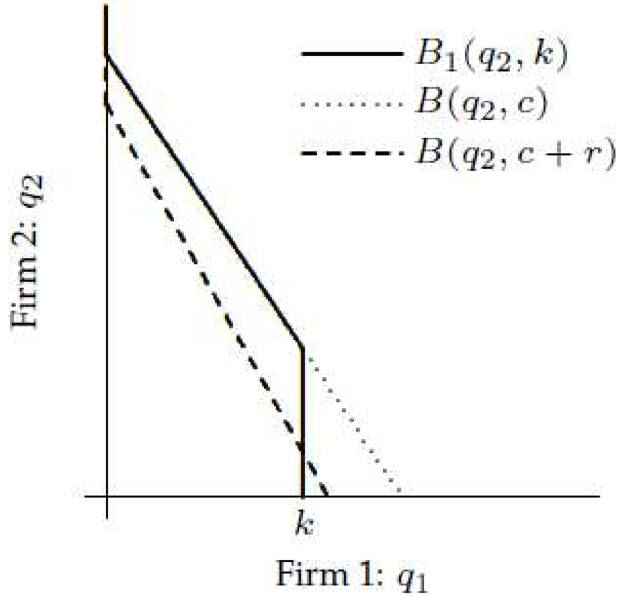

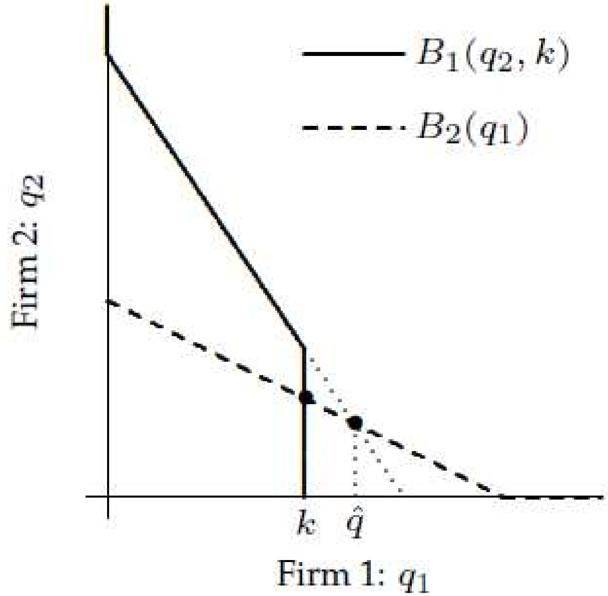

Any quantity q_1 up to $\hat{q}$ may be induced as part of the final-period Nash equilibrium, by appropriate choice of k: if $k > \hat{q}$ it would not be credible to produce at capacity. So $q_1^*(k) \leq \hat{q}$.

$$b_I(q_E|k_I) = \text{Min}\{b(q_E|w), k_I\}.$$

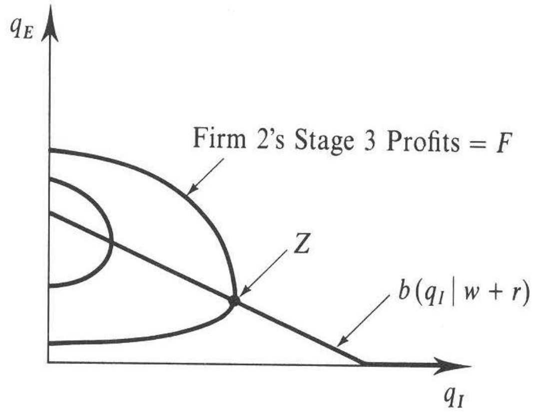

Figure 12.BB.1

Firm E's stage 3 best-response function after entry.

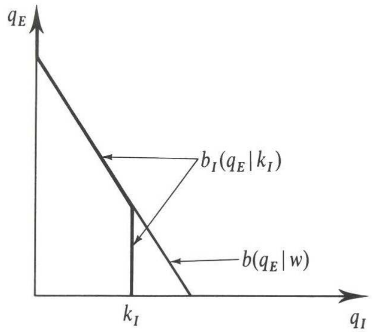

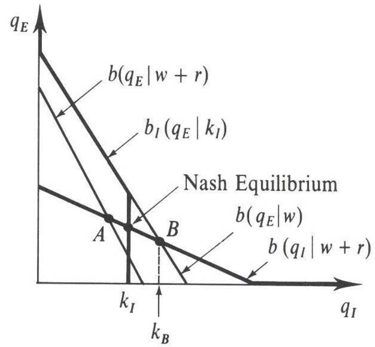

Figure 12.BB.2 (left)

Firm I's stage 3 best-response function after entry.

Figure 12.BB.3 (right)

Stage 3 Nash equilibrium after entry.

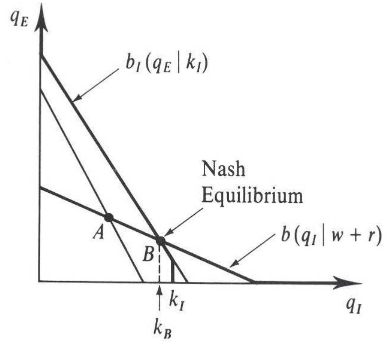

Figure 12.BB.4

A stage 3 equilibrium in which firm I does not use all of its capacity.

Blockaded Entry and Accommodation

In stage 2, firm 2 will choose to enter the industry only if, given the equilibrium induced by k it will make profits of at least f. This will happen only if q_1 is sufficiently small (recall $\partial\pi_2/\partial q_1 < 0$).

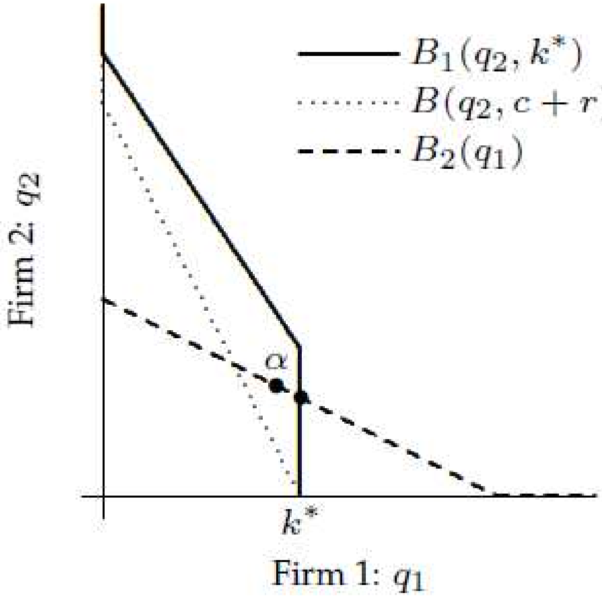

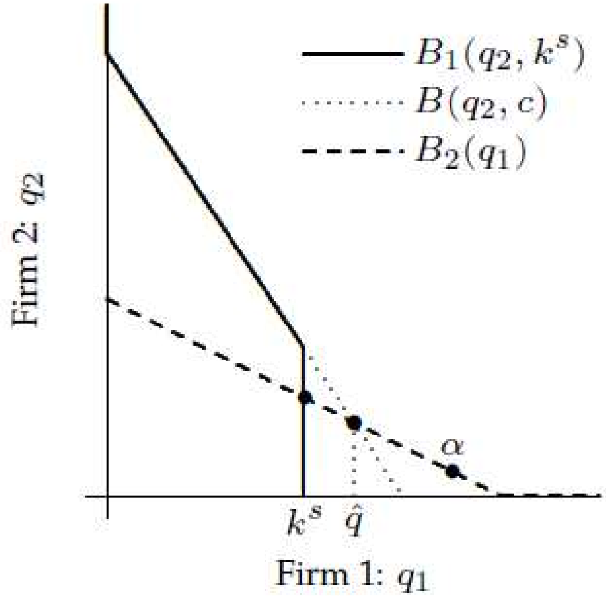

Entry is *blockaded* if point α (where $\pi_2 = f$) is left of $B(0, c + r)$, the optimal monopoly output. Firm 1 builds capacity $k^* = B(0, c + r)$, firm 2 does not enter as f is too high, and $q_1 = k^*$.

If α is to the right of $\hat{q}$, *accommodation* $\Rightarrow$ choose Stackelberg leader's quantity k^s (or $\hat{q}$ if $\hat{q} < k^s$).

Entry Deterrence

Finally, firm 1 may be *able* to deter entry, but may or may not wish to. . .

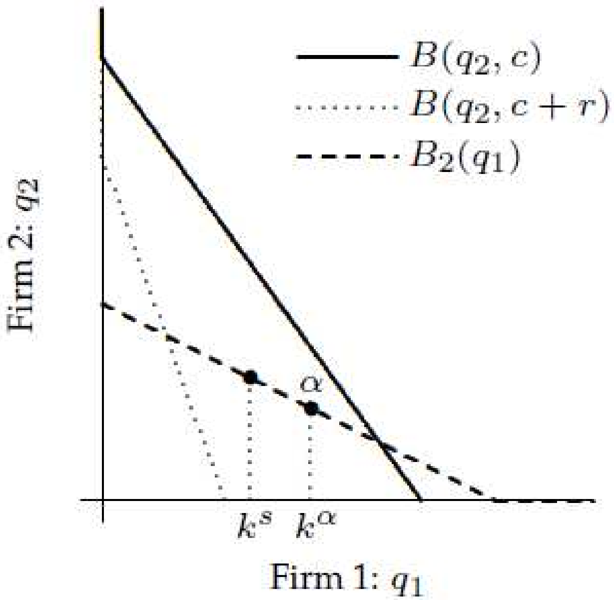

- Again, α is where $\pi_2 = f$: to the left of $B(q_2, c)$.
- Here, entry is neither blockaded nor inevitable.
- Firm 1 has a choice: deter entry by setting $k = k^\alpha$. . .
- . . . or accommodate entry optimally with $k = k^s$.
- Which is better? Compare firm 1's profit $\pi_1(q_1, q_2; k)$,

$$\pi_1(k^\alpha, 0; k^\alpha) \quad \text{with} \quad \pi_1(k^s, B_2(k^s); k^s)$$

- If the former is larger, deter. If not, accommodate.

Note that firm 1 chooses k to be *just* large enough to deter entry in the case that entry is deterred.

Note also that k^α exceeds optimal monopoly output $k^* = B(0, c + r)$. The incumbent firm overinvests in capacity to deter entry, raising output: the threat of entry alone increases welfare.

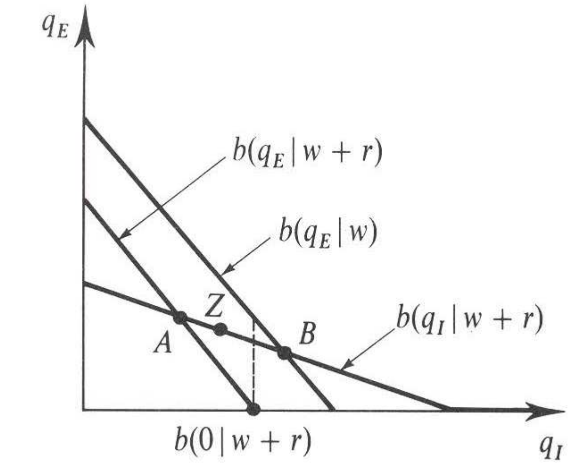

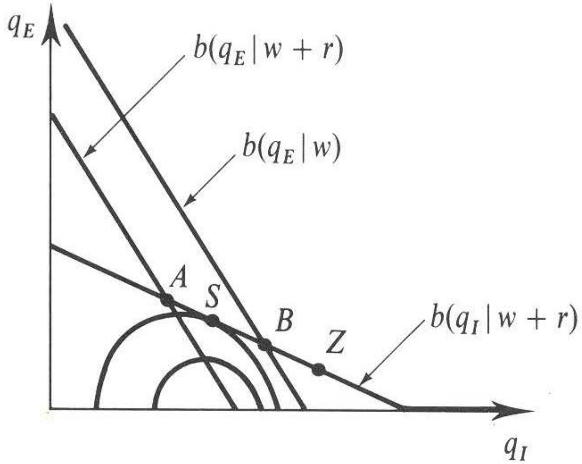

Figure 12.BB.5 (left)
Blockaded entry.

Figure 12.BB.6 (right)
Strategic entry accommodation when entry is inevitable.

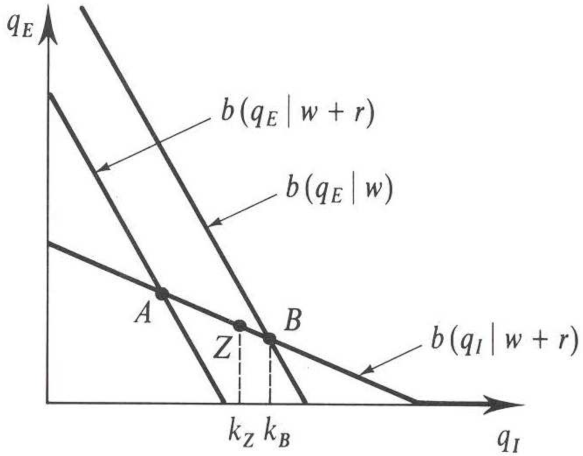

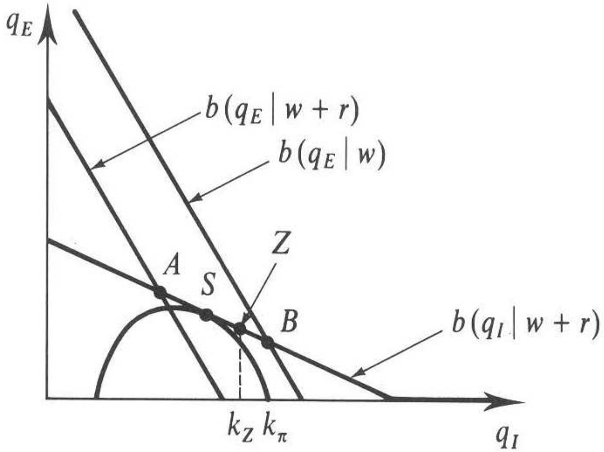

Figure 12.BB.7 (left)
Entry deterrence is possible but not inevitable.

Figure 12.BB.8 (right)
Entry deterrence versus entry accommodation.

(Time permitting) Cooperative Game Theory

The noncooperative alternating-offers model of structured bargaining (MWG Chapter 9, Appendix A, covered in this course by Dr. Meyer) is by far the most popular bargaining model among economic theorists.

But its main theoretical results are fragile: They don't generalize to n players, discrete offers, incomplete information, almost-common knowledge of rationality (Kreps 552-565; Kreps, *Game Theory and Economic Modelling*).

Further, its predictions don't do well when alternating-offers bargaining games are played in the laboratory, partly because Responders punish “unfair” offers as in the Ultimatum Game, and partly because the longer the horizon the more complex the backward induction/iterated dominance argument required to identify the subgame-perfect equilibrium; so complex that people don't believe that others will follow it.

When strategic uncertainty and risk of coordination failure are more important than delay costs, there's a fixed horizon, but there is no fixed pattern of alternating offers (unstructured bargaining), the analysis is very different.

Discuss Nash's (1953 *Econometrica*) demand game model with strategies viewed as the least surplus each player can be induced to accept. Discuss the role of expectations, culture, focal points, strategic moves in determining bargaining outcomes. Discuss Nash's axiomatic (1950 *EMT*) solution.