Notes for a Course in Game Theory

Maxwell B. Stinchcombe

Fall Semester, 2008. Unique #34215

Contents

CHAPTER 1

Organizational Stuff

First, a tentative schedule of topics and dates on which we'll talk about them, this will grow over the course of the semester.

Week, date	Topic(s)
#1, Wed. Aug. 27	Some stories with Nash eq'a in static games Ch. 2
#2, Wed. Sep. 3	Correlated eq'a, optimality of eq'a, equivalence of games Ch. 2 and 3
#3, Mon. Sep. 8	Optimality of eq'a, equiv. of games, genericity Ch. 3
#4, Wed. Sep. 10	Monote comparative statics for decision theory Ch. 4
#5, Mon. Sep. 15	Problems from Ch's 2-3 due.
#5, Mon. Sep. 15	Monote comparative statics for game theory Ch. 4

Meeting Time: We meet Mondays and Wednesday's, 12:30-2 p.m. in BRB 1.118. My e-mail is `maxwell@eco.utexas.edu`, my phone number is 475-8515, but it's a horrible way to reach me. For office hours, I'll hold a weekly problem session, at a time to be determined in class. You can also make appointments, or just drop in.

Texts: These lecture notes, articles referenced and linked in the on-line version of these notes. Much of what is here is also covered in the following sources: Drew Fudenberg and Jean Tirole, *Game Theory* (required reading), Robert Gibbons, *Game Theory for Applied Economists* (very strongly recommended reading), John McMillan, *Games, Strategies, and Managers* (a wonderful, early book on strategic thinking for managers), Eric Rasmussen, *Games and information : an introduction to game theory* (captures a great deal about how economists use game theoretic reasoning, almost used it as the text for the course), Herbert Gintis, *Game Theory Evolving* (as much fun as one can have reading and doing game theory), Brian Skyrms, *Evolution of the Social Contract* (fascinating), Klaus Ritzberger, *Foundations of Non-Cooperative Game Theory* (being careful about what we mean is a difficult but necessary step, this book does it wonderfully).

Problems: The lecture notes contain several Problem Sets, usually due every other week. Your combined grade on the Problem Sets will count for 50% of your total grade, a take-home midterm, given toward the end of October, will be worth 15%, the final exam will be worth 35%. If you hand in an incorrect answer to a problem from the Problem Sets, you can try the problem again. If your second attempt is wrong, you can try one more time.

It will be tempting to look for answers to copy. This is a mistake for two related reasons.

(1) Pedagogical: What you want to learn in this course is how to solve game theory models of your own. Just as it is rather difficult to learn to ride a bicycle by watching other people ride, it is difficult to learn to solve game theory problems if you do not practice solving them.

(2) Strategic: The final exam will consist of game models you have not previously seen. If you have not learned how to solve game models you have never seen before on your own, you will be unhappy at the end of the exam.

On the other hand, I encourage you to work together to solve hard problems, and/or to come talk to me. The point is to sit down, on your own, after any consultation you feel you need, and write out the answer yourself as a way of making sure that you can reproduce the logic.

Background: It is quite possible to take this course without having had a graduate course in microeconomics, one taught at the level of Mas-Colell, Whinston and Green' (MWG) *Microeconomic Theory*. However, some explanations will make reference to a number of consequences of the basic economic assumption that people pick so as to maximize their preferences. These consequences and this perspective are what one should learn in microeconomics. Simultaneously learning these and the game theory will be a bit harder.

In general, I will assume that everyone has a good working knowledge of calculus, a familiarity with simple probability arguments, and a great deal of comfort making arguments in mathematical terms. At various points in the semester, I will use the basic real analysis, perhaps even a bit of measure theory, that should be covered in last year's math for economists course. This material is at the level needed for most, but not all, theoretical research in micro, macro, and econometrics. However, most of this course does not assume any background in real analysis, and, if you talk with me, we will make special arrangements so that the lack of this background does not count against you.

CHAPTER 2

Some Hints About the Scope of Game Theory

In this Chapter, we're going to survey a number of games and indicate what their equilibria are as a way of seeing the scope of game theory. Eric Rasmussen [**16**] calls what we are doing "exemplary theory," meaning "theory by example." I believe that a better name is "theory by suggestion, using toy models." Essentially, we will use game theory models to tell stories about how society works. As well as Rasmusen's first chapter, you should read Chapter 1 in [**8**].

Story telling is an old and honored tradition. If one takes a functionalist approach to social institutions, it is a tradition that is meant to inculcate[1] the values of a society in its members. In slightly less grand words, this is part of your indoctrination into thinking as economists think.

Enough generalities, let us begin.

1. Defining a Game

A game is a collection $\Gamma = (A_i, u_i)_{i \in I}$. This has three pieces:

(1) I is the (usually finite) set of agents/people/players,
(2) for each $i \in I$, A_i is the set of actions or strategies available to i, and, setting $A = \times_{i \in I} A_i$,
(3) for each $i \in I$, $u_i : A \to \mathbb{R}$ represents i's preferences (usually von Neumann-Morgenstern preferences) over the actions chosen by others.

Having described who is involved in a strategic situation, the set I, and having described their available choices, the sets A_i, and their preferences over their own and everybody else's choices, we try to figure out what is going to happen. We have settled on a notion of equilibrium, due to John Nash, as our answer to the question of what will happen.

Definition 2.1. *A* ***Nash equilibrium*** *is a vector $a^* \in A$, $a^* = (a_i^*, a_{-i}^*)$, of actions with the property that each a_i^* is a best response to a_{-i}^*.*

Through examples, we will start to see what is involved.

2. The Advantage of Being Small

Here $I = \{\text{Big, Little}\}$, $A_1 = A_2 = \{\text{Push, Wait}\}$. This is called a 2×2 game because there are two players with two strategies apiece. The utilities are given in the following table.

[1]From the OED, to inculcate is to endeavour to force (a thing) into or impress (it) on the mind of another by emphatic admonition, or by persistent repetition; to urge on the mind, esp. as a principle, an opinion, or a matter of belief; to teach forcibly.

Rational Pigs

	Push	Wait
Push	$(-c+e, b-e-c)$	$(-c, b)$
Wait	$(\alpha b, (1-\alpha)b-c)$	$(0, 0)$

Some conventions: The representation of the choices has player 1, listed first, in this case Big, choosing which row occurs and player 2 choosing which column; If common usage gives the same name to actions taken by different players, then we do not distinguish between the actions with the same name; each entry in the matrix is uniquely identified by the actions a_1 and a_2 of the two players, each has two numbers, (x, y), these are $(u_1(a_1, a_2), u_2(a_1, a_2))$, so that x is the utility of player 1 and y the utility of player 2 when the vector $a = (a_1, a_2)$ is chosen.

There is a story behind the game: there are two pigs, one Big and one Little, and each has two actions.Little pig is player 1, Big pig player 2, the convention has 1's options being the rows, 2's the columms, payoffs (x, y) mean "x to 1, y to 2."

The story is of two pigs in a long room, a lever at one end, when pushed, gives food at the other end, the Big pig can move the Little pig out of the way and take all the food if they are both at the food output together, the two pigs are equally fast getting across the room, but when they both rush, some of the food, e, is pushed out of the trough and onto the floor where the Little Pig can eat it, and during the time that it takes the Big pig to cross the room, the Little pig can eat α of the food. This story is interesting when $b > c - e > 0$, $c > e > 0$, $0 < \alpha < 1$, $(1-\alpha)b - c > 0$. We think of b as the benefit of eating, c as the cost of pushing the lever and crossing the room. With $b = 6$, $c = 1$, $e = 0.1$, and $\alpha = \frac{1}{2}$, this gives

	Push	Wait
Push	$(-0.9, 4.9)$	$(-1, 6)$
Wait	$(3, 2)$	$(0, 0)$

Solve the game, note that the Little pig is getting a really good deal. There are situations where the largest person/firm has the most incentive to provide a public good, and the littler ones have an incentive to free ride. This game gives that in a pure form.

3. The Need for Coordinated Actions

Here is another 2×2 game,

Stag Hunt

	Stag	Rabbit
Stag	(S, S)	$(0, R)$
Rabbit	$(R, 0)$	(R, R)

As before, there is a story for this game: there are two hunters who live in villages at some distance from each other in the era before telephones; they need to decide whether to hunt for Stag or for Rabbit; hunting a stag requires that both hunters have their stag equipment with them, and one hunter with stag equipment will not catch anything; hunting for rabbits requires only one hunter with rabbit hunting equipment. The payoffs have $S > R > 0$, e.g. $S = 20$, $R = 1$, which gives

	Stag	Rabbit
Stag	$(20,20)$	$(0,1)$
Rabbit	$(1,0)$	$(1,1)$

This is a **coordination game**, if the players' coordinate their actions they can both achieve higher payoffs. There are two obvious Nash equilibria for this game. There is a role then, for some agent to act as a coordinator.

It is tempting to look for social roles and institutions that coordinate actions: matchmakers; advertisers; publishers of schedules e.g. of trains and planes. Sometimes we might imagine a tradition that serves as coordinator — something like we hunt stags on days following full moons except during the spring time.

Macroeconomists, well, some macroeconomists anyway, tell stories like this but use the code word "sunspots" to talk about coordination. This may be because overt reference to our intellectual debt to Keynes is out of fashion. In any case, any signals that are correlated and observed by the agents can serve to coordinate the peoples' actions.

4. Correlated Hunting

One version of this is that on sunny days, which happen γ of the time, the hunters go for stag, and on the other days, they go for rabbit. If both hunters always observe the same "signal," that is, the weather is the same at both villages, this gives the following distribution over outcomes:

	Stag	Rabbit
Stag	γ	0
Rabbit	0	$(1-\gamma)$

This is our first **correlated equilibrium**.

It is possible that the weather is different at the hunters' villages. Suppose the joint distribution of sun/rain at the two villages is

	Stag	Rabbit
Stag	a	b
Rabbit	c	d

Suppose both follow the strategy "Stag if it's sunny at my village, rabbit else." If we find conditions for these strategies to be mutual best responses, we've found another correlated equilibrium. If all row and column sums are positive, the conditions for player 1 are

$$20\frac{a}{a+b} + 0\frac{b}{a+b} \geq 1\frac{a}{a+b} + 1\frac{b}{a+b} \text{ if } (a+b) > 0, \tag{1}$$

and

$$1\frac{c}{c+d} + 1\frac{d}{c+d} \geq 20\frac{c}{c+d} + 0\frac{d}{c+d}. \text{ if } (c+d) > 0, \tag{2}$$

These are sensible if you think about conditional probabilities and suppose that the players maximize the expected value of the utility numbers we write down.

Problem 2.1. *Write down the inequalities for player* 2 *that correspond to* (1) *and* (2)*. To avoid the potential embarassment of dividing by* 0*, show that the conditional inequalities in* (1) *and* (2) *are satisfied iff*

$$(3) \qquad 20a + 0b \geq 1a + 1b, \text{ and } 1c + 1d \geq 20c + 0d.$$

Notation 2.2. *For any finite set S, $\Delta(S) := \{p \in \mathbb{R}^S_+ : \sum_a p_a = 1\}$.*

Definition 2.3. *A distribution, $\nu \in \Delta(A)$, is a* ***correlated equilibrium*** *if, for all i, for all $a, b \in A_i$, $\sum_{a_{-i}} u(a, a_{-i})\nu(a, a_{-i}) \geq \sum_{a_{-i}} u(b, a_{-i})\nu(a, a_{-i})$, equivalently, if $\sum_{a_{-i}}[u(a, a_{-i}) - u(b, a_{-i})]\nu(a, a_{-i}) \geq 0$.*

If you want to think in terms of conditional probabilities, then the inequalities in the Definition are $\sum_{a_{-i}} u(a, a_{-i})\frac{\nu(a,a_{-i})}{\sum_{c \in A} \nu(c,a_{-i})} \geq \sum_{a_{-i}} u(b, a_{-i})\frac{\nu(a,a_{-i})}{\sum_{c \in A} \nu(c,a_{-i})}$ because $\frac{\nu(a,a_{-i})}{\sum_{c \in A} \nu(c,a_{-i})}$ is the conditional distribution over A_{-i} given that a was drawn for player i.

You should convince yourself that the only correlated equilibrium of the Rational Pigs game is the Nash equilibrium.

Problem 2.2. *In the Stag Hunt game, what is the maximum correlated equilibrium probability of mis-coordination? If, instead of 1 being the payoff to R, it is $x > 0$, how does the maximum depend on x, and how does the dependence vary across different ranges of x, e.g. $x > 20$?*

5. Mixed Nash Equilibria

A correlated equilibrium that: 1) is not point mass on some action, and 2) has the actions of the players stochastically independent is a **mixed strategy Nash equilibrium**, a mixed eq'm for short.

Let $\alpha \in [0, 1]$ be the probability that 1 goes for Stag, $\beta \in [0, 1]$ be the probability that 2 goes for Stag, the independence gives the distribution

	Stag	Rabbit
Stag	$a = \alpha\beta$	$b = \alpha(1 - \beta)$
Rabbit	$c = (1 - \alpha)\beta$	$d = (1 - \alpha)(1 - \beta)$

Observation: Given the independence, the only way for those numbers to be a correlated equilibrium and have $0 < \alpha < 1$ is to have

$$20\beta + 0(1 - \beta) = 1\beta + 1(1 - \beta), \text{ i.e. } \beta = 1/20.$$

By symmetry, $\alpha = 1/20$ is the only way to have $0 < \beta < 1$ in equilibrium. Combining, this game has 3 Nash equilibria, the two pure strategy eq'a, (S, S) and (R, R), and the mixed equilibrium, $((1/20, 19/20), (1/20, 19/20))$.

Coordination problems often turn out to be deeply tied to complimentarities in the players' strategies. 1's expected utility from play of $((\alpha, (1 - \alpha)), (\beta, (1 - \beta)))$ is

$$U_1(\alpha, \beta) := E\, u_1 = 20\alpha\beta + 1(1 - \alpha),$$

and

$$\frac{\partial^2 U_1}{\partial\alpha\partial\beta} = 20 > 0.$$

This means that increases in β increase 1's marginal utility of increases in α. By symmetry, both players' best responses have this property. This opens the possibility of multiple eq'a.

6. Inspection Games and Mixed Strategies

If you have played Hide and Seek with very young children, you may have noticed that they will always hide in the same place, and that you need to search, while loudly explaining your actions, in other places while they giggle helplessly. Once they actually understand hiding, they begin to *vary* where they hide, they *mix* it up, they *randomize*.[2] Randomizing where one hides is the only sensible strategy in games of hide and seek. One can either understand the randomizing as people picking according to some internal random number generator, or as observing some random phenomenom outside of themselves and conditioning what they do on that. We will discuss a third understanding below. Whatever the understanding,

> ... the crucial assumption that we make for Nash equilibria is that when people randomize, they do it independently of each other.

Another game in which randomization is the only sensible way to play, at least, the only sensible way to play if you play at all often, is Matching Pennies, which is

	H	T
H	$(10, -10)$	$(-10, 10)$
T	$(-10, 10)$	$(10, -10)$

The unique Nash equilibrium for this game is for both players to independently randomize with probability $(\frac{1}{2}, \frac{1}{2})$. Another way to understand this is to imagine you're

Since we have in mind applications from economics, we consider inspection games, which have the same essential structure. The idea in inspection games is that keeping someone honest is costly, so you don't want to spend effort monitoring their behavior. But if you don't monitor their behavior, they'll want to slack off. The mixed strategy equilibria that we find balance these forces. We'll give two versions of this game, a very basic one, and a more complicated one.

The basic version: In this game, there is a worker who can either Shirk, or put in an Effort. The boss can either Inspect or Not. Inspecting someone who is working has an opportunity cost, $c > 0$, finding a Shirker has a benefit b. The worker receives w if they Shirk and are not found out, 0 if they Shirk and are Inspected, and $w - e$ if they put in the effort, whether or not they are Inspected. We assume that $w > e > 0$ so that $w - e > 0$. In matrix form, the game is

	Inspect	Don't inspect
Shirk	$(0, b)$	$(w, 0)$
Effort	$(w - e, -c)$	$(w - e, 0)$

Just as in the childrens' game of hide-and-seek, there cannot be an equilibrium in which the two players always choose one strategy. For there to be an equilibrium, there must be randomization. An equilibrium involving randomization is called a **mixed (strategy) equilibrium**, one not involving randomization a **pure (strategy) equilibrium**.

[2]Once they are teenagers, you do not want to play this game with them.

Solving for the mixed equilibrium, note that there is only one eq'm, then see that β and α must satisfy

$$0\beta + w(1-\beta) = w - e, \text{ i.e. } \beta = e/w$$
$$b\alpha - c(1-\alpha) = 0, \text{ i.e. } \alpha = c/(b+c).$$

Notice what has happened, the probability of shirking goes down as the cost of monitoring goes down, this is sensible, but the probability of monitoring is independent of the monitors costs and benefits, which is peculiar.

Another, more heavily parametrized version: If a chicken packing firm leaves the fire escape doors operable, they will lose c in chickens that disappear to the families and friends of the chicken packers. If they nail or bolt the doors shut, which is highly illegal, they will no longer lose the c, but, if they are inspected (by say OSHA), they will be fined f. Further, if the firedoor is locked, there is a risk, ρ, that they will face civil fines or criminal worth F if there is a fire in the plant that kills many of the workers because they cannot escape.[3] Inspecting a plant costs the inspectors k, not inspecting an unsafe plant costs B in terms of damage done to the inspectors' reputations and careers. Filling in the other terms, we get the game

		Inspectors	
		Inspect	Not inspect
Imperial	unlocked	$(\pi - c, -k)$	$(\pi - c, 0)$
	locked	$(\pi - f - \rho F, f - k)$	$(\pi - \rho F, -B)$

If f and ρF are too low, specifically, if $c > f + \rho F$, then Imperial has a dominant strategy, and the game is, strategically, another version of Rational Pigs.

If $f + \rho F > c > \rho F$ and $f - k > -B$, neither player has a dominant strategy, and there is only a mixed Nash equilibrium. In this case, we have another instance of a game like the inspection game.

Problem 2.3. *Assume that $f + \rho F > c > \rho F$ and $f - k > -B$ in the inspection game. Show that the equilibrium is unchanged as π grows. How does it change as a function of c?*

7. The Use of Deadly Force and Evolutionary Arguments

One of the puzzles about competition between animals over resources is the rarity of the use of deadly force — poisonous snakes wrestle for mates but rarely bite each other, wolves typically end their fights with both alive and almost uninjured after the beaten wolf has exposed their throat, male caribou struggle for mates by locking antlers and trying to push each other around but rarely gouge each other. The puzzle is that the use of deadly force

[3]White collar decisions that kill blue collar workers rarely result in criminal prosecutions, and much more rarely in criminal convictions. See [**14**] for some rather depressing statistics. Emmett Roe, the owner of the Imperial chicken processing plant that locked the doors killed 25 workers and injured 56 more on September 3, 1991. He plea-bargained to 25 counts of involuntary manslaughter, was sentenced to 20 years in prison, and was eligible for early release after 3 years, and was released after 4 and a half years, that is, 65 days for each of the dead. The surviving injured workers and families of the dead only won the right to sue the state for failure to enforce safety codes on February 4, 1997, after a five-year battle that went to the state Court of Appeals. Damage claims will be limited to $100,000 per victim.

against a rival who is not using it rather thoroughly removes that rival from the gene pool, thereby giving an immediate selective advantage to the user of deadly force. The problem with this argument is that it is only selectively advantageous to use deadly force in a population full of non-users of deadly force.

Suppose that there is a prize at stake worth $50x$ (in utility terms) to each of two contestants, where $x \in (0, 2)$. The contestants have two options, aggressive Display or deadly Force. Display immediately runs away from Force. The utility of being seriously injured, which happens $\frac{1}{2}$ of the time if both use deadly Force, is -100. Finally, the loss of time in a long mutual display of aggression has a utility of -10. With payoffs, the game, $\Gamma(x)$, is

	Display	Force
Display	$(-10, -10)$	$(0, 50x)$
Force	$(50x, 0)$	$(50(\frac{1}{2}x - 1), 50(\frac{1}{2}x - 1))$

The mixed strategy equilibrium is interesting from an evolutionary point of view — if we thought of behavior as being genetically determined, then a population consisting of mostly Display (resp. Force) users would give selective advantage to Force (resp. Display) users, and this balancing pattern would be at work unless the population proportions matched the mixed equilibrium probabilities.

Let $\alpha(x)$ be the probability that Display is used in the mixed equilibrium (or, in the evolutionary story, the proportion of the population using Display). The formula is

$$\alpha(x) = \frac{2 - x}{2.4 - x}, \text{ so that } \frac{d\alpha(x)}{dx} = -\frac{0.4}{(2.4 - x)^2} < 0.$$

The derivative being less than 0 means that the more valuable the prize is, the higher the proportion of the time one would expect to see the use of deadly Force.

Note that $\alpha(2) = 0$ so that Force is used all of the time. If $x > 2$, the formula for $\alpha(x)$ breaks down (negative probabilities are no-no's). For these values of x, Force is a dominant strategy for both players.

Notice how the focus has gone from the behavior of a given pair of individuals to a representative pair of individuals, from micro to macro if you will.

8. Games of chicken

Think testoserone poisoning for this one — two adolescent boys run toward each other along a slippery pier over the ocean, at pre-determined points, they jump onto their boogie boards, each has the option of "chickening out" and ending up in the ocean, or going through, since they are going headfirst at full running speed, if they both decide to go through, they both end up with concussions, since chickening out loses face (the other boys on the beach laugh at them), the payoffs are

	Chicken	Thru
Chicken	$(0, 0)$	$(-2, 10)$
Thru	$(10, -2)$	$(-9, -9)$

Sometimes, one thinks about lockout/agreement and strike/agreement problems using a game of chicken.

There are no dominant strategies for this game. There are two pure strategy equilibria, and one, rather disastrous (for the parents of the sense-challenged young men) mixed strategy equilibrium. Here, the complementarities are negative, in terms of α, β, the probabilities that the two play Chicken,

$$U_1(\alpha, \beta) := E\, u_1 = \text{mess} + \alpha\beta[2 - 9 - 10], \text{so } \partial^2 U_1/\partial\alpha\partial\beta < 0.$$

Increases in β decrease 1's utility to α, a decreasing best response set.

9. In Best Response Terms

When $a = (a_i, a_{-i})$, we interchangeably use (b_i, a_{-i}) and $a\backslash b_i$. $\Delta_i = \Delta(A_i)$, $\Delta = \times_{i\in I}\Delta$, $\sigma = (\sigma_i, \sigma_{-i})$, we interchangeably use (ν_i, σ_{-i}) and $\sigma\backslash\nu_i$. For any $\sigma \in \Delta$ and $i \in I$, $U_i(\sigma) := E^\sigma\, u_i = \sum_{a\in A} u_i(a)\sigma(a) = \int_A u(a)\, d\sigma(a)$.

Definition 2.4. *$Br_i(\sigma) = \{\nu_i \in \Delta_i : U_i(\sigma\backslash\nu_i) \geq U_i(\sigma\backslash\Delta_i)\}$ is i **mixed best response set**, and $Br(\sigma) := \times_{i\in I} Br_i(\sigma)$ is the **(mutual) best response set**.*

In these terms, σ^* is a Nash eq'm iff $\sigma^* \in Br(\sigma^*)$. To check that σ^* is an eq'm, for each i, check that σ_i^* is a solution to the problem

$$\max_{\nu_i\in\Delta_i} U_i(\nu_i, \sigma^*_{-i}). \tag{4}$$

> The essential behavioral assumption is Nash equilibria is that each i acts as though they know what $-i$ is/are doing.

This is a VERY strong assumption. You should feel very uneasy about it. If we are to make this assumption, we should have a reason to believe that the maximization problem in (4) is descriptive, either because it's clear what $-i$ will do, think of Rational Pigs for example, or because we are talking about a situation in which we believe that everyone involved has figured out what the others are doing. It is too easy to (act as if one has) come to believe this.[4] Sometimes it seems impossible that people would have that knowledge. When we come to infinitely repeated games, it will/should seem really really impossible for some of our analyses.

So, Nash eq'm is a strong assumption. On the other hand, it is a minimal assumption. If σ^* is not an eq'm, then someone can make themselves strictly better off by playing something else. It is hard to see how we could settle on anything but a Nash eq'm. Those words, "settle on," will come back (to haunt us) later.

10. Two Prisoners

These two classic games have dominant strategies for at least one player.

[4] There is a story about two economists walking down the road, the younger one spots a $100 bill on the street, and excitedly bends down to pick it up, only to be stopped by the senior economist, looking at the younger one with sadness. "If it really was a hundred dollar bill, someone would have already picked it up." The elder had too thoroughly absorbed the assumption that everything we see is an eq'm.

Prisoners' Dilemma

	Squeal	Silent
Squeal	$(-B+r, -B+r)$	$(-b+r, -B)$
Silent	$(-B, -b+r)$	$(-b, -b)$

Rational Pigs

	Push	Wait
Push	$(-c+e, b-e-c)$	$(-c, b)$
Wait	$(\alpha b, (1-\alpha)b-c)$	$(0, 0)$

We've seen the story for the Rational Pigs. In the Dilemma, two criminals have been caught, but it is after they have destroyed the evidence of serious wrongdoing. Without further evidence, the prosecuting attorney can charge them both for an offense carrying a term of $b > 0$ years. However, if the prosecuting attorney gets either prisoner to give evidence on the other (Squeal), they will get a term of $B > b$ years. The prosecuting attorney makes a deal with the judge to reduce any term given to a prisoner who squeals by an amount r, $b \geq r > 0$, $B - b > r$ (equivalent to $-b > -B + r$). With $B = 15$, $b = r = 1$, this gives

	Squeal	Silent
Squeal	$(-14, -14)$	$(0, -15)$
Silent	$(-15, 0)$	$(-1, -1)$

In the Prisoners' Dilemma, Squeal dominates Silent for both players. Another way to put this, the only possibly rational action for either player is Squeal. This does not depend on i knowing what $-i$ is doing — whatever $-i$ does, Squeal is best. We might as well solve the optimization problems independently of each other.

What makes it interesting is that when you put the two solutions together, you have a disaster from the point of view of the players. They are both spending 14 years in prison, and by cooperating with each other and being Silent, they could both spend only 1 year in prison. Their individualistic choice of an action imposes costs on someone else. The individual calculation and the joint calculation are very very different.

One useful way to view many economists is as apologists for the inequities of a moderately classist version of the political system called laissez faire capitalism. Perhaps this is the driving force behind the large literature trying to explain why we should expect cooperation in this situation. After all, if economists' models come to the conclusion that equilibria without outside intervention can be quite bad for all involved, they become an attack on the justifications for laissez faire capitalism. Another way to understand this literature is that we are, in many ways, a cooperative species, so a model predicting extremely harmful non-cooperation is very counter-intuitive.

11. Tragedy of the commons

There are I different countries that can put out fishing fleets to catch from pelagic schools of fish. Use the number $a_i \in \mathbb{R}_+ = A_i$ to represent the number of fishing boats in the fleet of country i, $i = 1, \ldots, I$. To finish specifying the game, the utilities, $u_i : A \to \mathbb{R}$ need to be specified.

The marginal cost of a boat is constant, and equal to c. For given $a \in A$, let $n = n(a) = \sum_{i \in I} a_i$ and $n_{-i}(a) = \sum_{j \neq i} a_j$. When the total number of boats is n, the **per boat** return is $v(n)$ where $v(0) > 0$, $v'(n) < 0$, $v'(0) < c$, and $v''(n) < 0$. For country i, the benefit to putting out a fleet depends on the size of their own fleet, a_i, and the size of the other countries' fleets,

$n_{-i}(s)$,

$$u_i(a_i, n_{-i}(a)) = a_i v(a_i + n_{-i}(a)) - ca_i = a_i v(n(a)) - ca_i.$$

For fixed $n_{-i} = n_{-i}(a)$, $u_i(\cdot, n_{-i})$ is concave because $\frac{\partial u_i(a_i, n_{-i})}{\partial a_i} = v(n) + a_i v'(n) - c$, implying

$$\frac{\partial^2 u_i(a_i, n_{-i})}{\partial a_i^2} = v'(n) + a_i v''(n) + v'(n) < 0.$$

The simultaneous satisfaction of the best response conditions gives the Nash equilibrium. Compare these to the FOC from any Pareto optimal point. In this game, the Nash equilibrium is inefficient compared to binding agreements to limit fleet size, and the inefficiency grows with I. Strategic considerations do not lead to socially optimal solutions.

In this game, the players' strategy sets could be taken to be interval subsets of $\mathbb{R}$, and the equilibrium, being interior to the intervals, was not at a vertex. Theorem 4.1 of [**1**] tells us that the set of utility functions for which all non-vertex equilibria fail to be Pareto optimal is open and finitely prevalent.

In this particular game, there is a fairly easy way to see what is involved in their argument. Let $(a_1^*, a_2^*) \gg (0, 0)$ be an equilibrium. The indifference curve of 1 in the $A_1 \times A_2$ square comes tangent, from below, to the line $a_2^* \times [0, M]$ at the eq'm point (a_1^*, a_2^*), and vice versa. Draw these two, note that when we go down and to the left, both are better off.

Crucial to this argument is the observation that neither player is indifferent, at the eq'm, to what the other is doing. This allows for us to have an area where we can push both players' actions and make both better off. What Anderson and Zame do is to show that for essentially all utility functions, at an interior pure strategy eq'm, no-one is indifferent to the actions of others.

12. Cournot Competition and Some Dynamics

Two firms selling a homogeneous product to a market described by a known demand function and using a known technology decide on their quantities, $s^i \in [0, M]$, $i = 1, 2$. There is an initial state $\theta_0 = (\theta_{i,0})_{i \in I} = (\theta_{1,0}, \theta_{2,0}) \in S^2$. When t is an odd period, player 1 changes $\theta_{1,t-1}$ to $\theta_{1,t} = Br_1(\theta_{2,t-1})$, when t is an even period, player 2 changes $\theta_{2,t-1}$ to $\theta_{2,t} = Br_2(\theta_{1,t-1})$. Or, if you want to combine the periods,

$$(\theta_{1,t-1}, \theta_{2,t-1}) \mapsto (Br_1(\theta_{2,t-1}), Br_2(Br_1(\theta_{2,t-1}))).$$

In either case, note that if we set $S^0 = \{h^0\}$ (some singleton set), we have specified a **dynamic system**, that is, a class of functions $f_t : \Theta \times S^{t-1} \to S$, $t \in \mathbb{N}$. When we combine periods, the f_t has a form that is independent of the period, $f_t \equiv f$, and we have a **stationary dynamic system**. Whatever dynamic system we study, for each θ_0, the result is the outcome point

$$\mathbb{O}(\theta_0) = (\theta_0, f_1(\theta_0), f_2(\theta_0, f_2(\theta_0)), \ldots),$$

a point in $\Theta \times S^\infty$. When $f_t \equiv f$ is independent of t and depends only on the previous period's outcome,

$$\mathbb{O}(\theta_0) = (\theta_0, f(\theta_0), f(f(\theta_0)), f(f(f(\theta_0))), \ldots).$$

Definition 2.5. *A point $\hat{s}$ is* **stable** *for the dynamic system $(f_t)_{t\in\mathbb{N}}$ if $\exists\theta_0$ such that*

$$\mathbb{O}(\theta_0) = (\theta_0, \hat{s}, \hat{s}, \ldots).$$

With the best response dynamics specified above, the stable points are exactly the Nash equilibria.

12.0.1. *Convergence, stability, and local stability.* Suppose we have a way to measure the distance between points in S, e.g. $d(u, v) = \sqrt{(u_1 - v_1)^2 + (u_2 - v_2)^2}$ when $S = [0, M]^2$. The d-ball around u with radius ϵ is the set $B(u, \epsilon) = \{v \in S : d(u, v) < \epsilon\}$.

Problem 2.4. *A* **metric on a set** *X is a function $d : X \times X \to \mathbb{R}_+$ with the following three properties:*

(1) $(\forall x, y \in X)[d(x, y) = d(y, x)]$,
(2) $(\forall x, y \in X)[d(x, y) = 0$ *iff* $x = y]$, *and*
(3) $(\forall x, y, z \in X)[d(x, y) + d(y, z) \geq d(x, z)]$.

Show that $d(u, v) = \sqrt{(u_1 - v_1)^2 + (u_2 - v_2)^2}$ is a metric on the set $S = [0, M]^2$. Also show that $\rho(u, v) = |u_1 - v_1| + |u_2 - v_2|$ and $r(u, v) = \max\{|u_1 - v_1|, |u_2 - v_2|\}$ are metrics on $S = [0, M]^2$. In each case, draw $B(u, \epsilon)$.

There are at least two useful visual images for convergence: points s_1, s_2, s_3, etc. appearing clustered more and more tightly around u; or, looking at the graph of the sequence (remember, a sequence is a function) with $\mathbb{N}$ on the horizontal axis and S on the vertical, as you go further and further to the right, the graph gets closer and closer to u. Convergence is a crucial tool for what we're doing this semester.

Definition 2.6. *A sequence $(s_n) \in S^\infty$* **converges to** *$u \in S$* **for the metric** *$d(\cdot, \cdot)$ if for all $\epsilon > 0$, $\exists N$ such that $\forall n \geq N$, $d(s_n, u) < \epsilon$. A sequence* **converges** *if it converges to some u.*

In other notations, $s \in S^\infty$ converges to $u \in S$ if

$$(\forall\epsilon > 0)(\exists K)(\forall k \geq K)[d(z_k(s), u) < \epsilon],$$
$$(\forall\epsilon > 0)(\exists K)(\forall k \geq K)[z_k(s) \in B(u, \epsilon)].$$

These can be written $\lim_k z_k(s) = u$, or $\lim_n s_n = u$, or $z_k(s) \to u$, or $s_n \to u$, or even $s \to u$.

Example 2.7. *Some convergent sequences, some divergent sequences, and some cyclical sequences that neither diverge nor converge.*

There is yet another way to look at convergence, based on cofinal sets. Given a sequence s and an $N \in \mathbb{N}$, define the cofinal set $C_N = \{s_n : n \geq N\}$, that is, the values of the sequence from the N'th onwards. $s_n \to u$ iff $(\forall\epsilon > 0)(\exists M)(\forall N \geq M)[C_N \subset B(u, \epsilon)]$. This can be said "$C_N \subset B(u, \epsilon)$ for all large N" or "$C_N \subset B(u, \epsilon)$ for large N." In other words, the English phrases "for all large N" and "for large N" have the specific meaning just given.

Another verbal definition is that a sequence converges to u if and only if it gets and stays arbitrarily close to u.

The following is a real analysis kind of problem. It is not supposed to be difficult for economics grad students who have had the mathematics for economists first year course. For others students, this may be entirely unfamiliar, and you should talk to me or the T.A.

Problem 2.5. *Show that $s \in S^\infty$ converges to u in the metric $d(\cdot,\cdot)$ of Homework 2.4 iff it converges to u in the metric $\rho(\cdot,\cdot)$ iff it converges to u in the metric $r(\cdot,\cdot)$.*

Convergence is what we hope for in dynamic systems, if we have it, we can concentrate on the limits rather than on the complicated dynamics. Convergence comes in two flavors, local and global.

Definition 2.8. *A point $\hat{s} \in S$ is* **asymptotically stable** *or* **locally stable** *for a dynamic system $(f_t)_{t\in\mathbb{N}}$ if it is stable and $\exists\epsilon > 0$ such that for all $\theta_0 \in B(\hat{\theta}, \epsilon)$, $\mathbb{O}(\theta_0) \to \hat{s}$.*

Example 2.9. *Draw graphs of non-linear best response functions for which there are stable points that are not locally stable.*

When the f_t's are fixed, differentiable functions, there are derivative conditions that guarantee asymptotic stability. These results are some of the basic limit theorems referred to above.

Definition 2.10. *A point $\hat{s} \in S$ is* **globally stable** *if it is stable and $\forall\theta_0$, $\mathbb{O}(\theta_0) \to \hat{s}$.*

NB: If there are many stable points, then there cannot be a globally stable point.

12.0.2. *Subsequences, cluster points, and ω-limit points.* Suppose that $\mathbb{N}'$ is an infinite subset of $\mathbb{N}$. $\mathbb{N}'$ can be written as

$$\mathbb{N}' = \{n_1, n_2, \ldots\}$$

where $n_k < n_{k+1}$ for all k. Using $\mathbb{N}'$ and sequence $(s_n)_{n\in\mathbb{N}}$, we can generate another sequence, $(s_{n_k})_{k\in\mathbb{N}}$. This new sequence is called a **subsequence** of $(s_n)_{n\in\mathbb{N}}$. The trivial subsequence has $n_k = k$, the even subsequence has $n_k = 2k$, the odd has $n_k = 2k - 1$, the prime subsequence has n_k equal to the k'th prime integer, etc.

Definition 2.11. *A* **subsequence of** *$s = (s_n)_{n\in\mathbb{N}}$ is the restriction of s to an infinite $\mathbb{N}' \subset \mathbb{N}$.*

By the one-to-one, onto mapping $k \leftrightarrow n_k$ between $\mathbb{N}$ and $\mathbb{N}'$, every subsequence is a sequence in its own right. Therefore we can take subsequences of subsequences, subsequences of subsequences of subsequences, and so on.

Sometimes a subsequence of (s_n) will be denoted $(s_{n'})$, think of $n' \in \mathbb{N}'$ to see why the notation makes sense.

Definition 2.12. *u is a* **cluster point** *or* **accumulation point** *of the sequence $(s_n)_{n\in\mathbb{N}}$ if there is a subsequence $(s_{n_k})_{k\in\mathbb{N}}$ converging to u.*

s_n converges to u iff for all $\epsilon > 0$, the cofinal sets $C_N \subset B(u,\epsilon)$ for all large N. s_n clusters or accumulates at u iff for all $\epsilon > 0$, the cofinal sets $C_N \cap B(u,\epsilon) \neq \emptyset$ for all large N. Intuitively, u is a cluster point if the sequence visits arbitrarily close to u infinitely many times, and u is a limit point if the sequence does nothing else.

Example 2.13. *Some convergent sequences, some cyclical sequences that do not converge but cluster at some discrete points, a sequence that clusters "everywhere."*

Let $\text{accum}\,(s)$ be the set of accumulation points of an $s \in S^\infty$.

Definition 2.14. *The set of* ω**-limit points** *of the dynamic system* $(f_t)_{t\in\mathbb{N}}$ *is set*

$$\bigcup_{\theta\in\Theta} \text{accum}\,(\mathbb{O}(\theta)).$$

If a dynamic system cycles, it will have ω-limit points. Note that this is true even if the cycles take different amounts of time to complete.

Example 2.15. *A straight-line cobweb example of cycles, curve the lines outside of some region to get an attractor.*

The distance between a set S' and a point x is defined by $d(x, S') = \inf\{d(x, s') : s' \in S'\}$ (we will talk in detail about inf later, for now, if you haven't seen it, treat it as a min). For $S' \subset S$, $B(S', \epsilon) = \{x : d(x, S') < \epsilon\}$. If you had graduate micro from me, you've seen this kind of set.

When $\Theta = S$ and S is compact, a technical condition that we will spend a great deal of time with (later), we have

Definition 2.16. *A set* $S' \subset S$ *is* **invariant** *under the dynamical system* $(f_t)_{t\in\mathbb{N}}$ *if* $\theta \in S'$ *implies* $\forall k$, $z_k(\mathbb{O}(\theta)) \in S'$. *An invariant* S' *is an* **attractor** *if* $\exists\epsilon > 0$ *such that for all* $\theta \in B(S', \epsilon)$, $\text{accum}\,(\mathbb{O}(\theta)) \subset S'$.

Strange attractors are really cool, but haven't had much impact in the theory of learning in games, probably because they are so strange.

CHAPTER 3

Optimality? Who Told You the World is Optimal?

We're going to begin with a brief study of the strategic equivalence of games and its implication for Pareto optimality. We will also look at some of the ideas about changing preferences in games: negatively independent preferences by (Kockesen, Ok, and Sethi [**11**]), arguments about the evolution of preferences and what it is that people maximize, (Heifetz, Shannon, and Spiegel [**10**]), and arguments about generic non-optimality of interior pure strategy equilibria, (Anderson and Zame [**1**]).

1. A Little Bit of Decision Theory

We assume that people act so as to maximize their expected utility taking others' actions/choices as given. In other words, assuming that what they choose in their optimization does not affect what others choose. Here is a useful Lemma. It may seem trivial, but it turns out to have strong implications for our interpretations of equilibria in game theory.

LEMMA 3.1 (Rescaling). *Suppose that $u : A \times \Omega \to \mathbb{R}$ is bounded and measurable. $\forall Q_a, Q_b \in A$, $\forall P \in \Delta(\mathcal{F})$,*

$$\int_A \left[\int_\Omega u(x,\omega)\, dP(\omega) \right] dQ_a(x) \geq \int_A \left[\int_\Omega u(y,\omega)\, dP(\omega) \right] dQ_b(y)$$

iff

$$\int_A \left[\int_\Omega [\alpha \cdot u(x,\omega) + f(\omega)]\, dP(\omega) \right] dQ_a(x) \geq \int_A \left[\int_\Omega [\alpha \cdot u(y,\omega) + f(\omega)]\, dP(\omega) \right] dQ_b(y)$$

for all $\alpha > 0$ and P-integrable functions f.

Remember how you learned that Bernoulli utility functions were immune to multiplication by a positive number and the addition of a constant? Here the constant is being played by $F := \int_\Omega f(\omega)\, dP(\omega)$.

Proof: Suppose that $\alpha > 0$ and $F = \int f\, dP$. Define $V(x) = \int_\Omega u(x,\omega)\, dP(\omega)$. The Lemma is saying that $\int_A V(x)\, dQ_a(x) \geq \int_A V(y)\, dQ_b(y)$ iff $\alpha \left[\int_A V(x)\, dQ_a(x) \right] + F \geq \alpha \left[\int_A V(y)\, dQ_b(y) \right] + F$, which is immediate. ∎

2. Generalities About 2×2 Games

We are going to focus on games where there are no ties — for each $i \in I$ and a_{-i}, $u_i(a_i, a_{-i}) \neq u_i(b_i, a_{-i})$ for $a_i \neq b_i$. Within this class of 2×2 games, we've seen four types:

1. Games in which both players have a dominant strategy, e.g. Prisoners' Dilemma;
2. Games in which exactly one player has a dominant strategy, e.g. Rational Pigs;

(3) Games in which neither player has a dominant strategy and there are three equilibria, e.g. Stag Hunt, Battle of the Partners, Deadly Force, Chicken;
(4) Games in which neither player has a dominant strategy and there is only a mixed strategy equilibrium, e.g. Hide and Seek, Matching Pennies, Inspection.

3. Rescaling and the Strategic Equivalence of Games

Consider the 2×2 game

	Left	Right
Up	$(\boldsymbol{a}, e)$	$(\boldsymbol{b}, f)$
Down	$(\boldsymbol{c}, g)$	$(\boldsymbol{d}, h)$

where we've put 1's payoffs in bold for emphasis. Since we're assuming there are no ties for player 1, $\boldsymbol{a} \neq \boldsymbol{c}$ and $\boldsymbol{b} \neq \boldsymbol{d}$. Consider the function $f_1(a_2)$ given by $f_1(\text{Left}) = -\boldsymbol{c}$ and $f_1(\text{Right}) = -\boldsymbol{b}$. Lemma 3.1 tells us that adding f_1 to 1's payoffs cannot change either CEq or Eq. When we do this we get the game

	Left	Right
Up	$(a - c, \boldsymbol{e})$	$(0, \boldsymbol{f})$
Down	$(0, \boldsymbol{g})$	$(d - b, \boldsymbol{h})$

where we've now put 2's payoffs in bold for emphasis. Since we're assuming there are no ties for player 2, $\boldsymbol{e} \neq \boldsymbol{f}$ and $\boldsymbol{g} \neq \boldsymbol{h}$. Consider the function $f_2(a_1)$ given by $f_2(\text{Up}) = -\boldsymbol{f}$ and $f_1(\text{Down}) = -\boldsymbol{g}$. Lemma 3.1 tells us that adding f_2 to 2's payoffs cannot change either CEq or Eq. When we do this we get the game

	Left	Right
Up	(x, y)	$(0, 0)$
Down	$(0, 0)$	(r, s)

where $x = a - c$, $y = e - f$, $r = d - b$, $s = h - g$, and $x, y, r, s \neq 0$. We've just proved that all 2×2 games with no ties are equivalent to 2×2 games with $(0, 0)$'s in the off-diagonal positions.

Applying this procedure to the six of the 2×2 games we've seen yields

Prisoners' Dilemma

	Squeal	Silent
Squeal	$(1, 1)$	$(0, 0)$
Silent	$(0, 0)$	$(-1, -1)$

Rational Pigs

	Push	Wait
Push	$(-3.9, -1.1)$	$(0, 0)$
Wait	$(0, 0)$	$(1, -2)$

Matching Coins

	H	T
H	$(2, -2)$	$(0, 0)$
T	$(0, 0)$	$(2, -2)$

Stag Hunt

	Stag	Rabbit
Stag	$(9, 9)$	$(0, 0)$
Rabbit	$(0, 0)$	$(1, 1)$

Battle of the Partners

	Dance	Picnic
Dance	$(12, 8)$	$(0, 0)$
Picnic	$(0, 0)$	$(8, 12)$

Chicken

	Chicken	Thru
Chicken	$(-10, -10)$	$(0, 0)$
Thru	$(0, 0)$	$(-7, -7)$

Once games are in this form, what matters for strategic analysis are the signs of the utilities x, y, r, s, e.g. the sign patterns for the first two games are

Prisoners' Dilemma

	Squeal	Silent
Squeal	$(+, +)$	$(0, 0)$
Silent	$(0, 0)$	$(-, -)$

Rational Pigs

	Push	Wait
Push	$(-, -)$	$(0, 0)$
Wait	$(0, 0)$	$(+, -)$

There are 2^4 possible sign patterns, but there are not 2^4 strategically distinct 2×2 games. We're going to say that two games are strategically equivalent if they have the same sign pattern after any finite sequence of (1) applying Lemma 3.1 to arrive at 0's off the diagonal, (2) relabeling a player's actions, or (3) relabeling the players.

For example, in Rational Pigs, Little Pig was player 1 and Big Pig was player 2. If we relabeled them as 2 and 1 respectively, we would not have changed the strategic situation at all. We would have changed how we represent the game, but that should make no difference to the pigs. This would give a game with the sign pattern

	Push	Wait
Push	$(-, -)$	$(0, 0)$
Wait	$(0, 0)$	$(-, +)$

If we were to relabel the actions of one player in Chicken, we'd have the game

	a_2	b_2
a_1	$(0, 0)$	$(-10, -10)$
b_1	$(-7, -7)$	$(0, 0)$

which is equivalent, via Lemma 3.1, to a game with the sign pattern

	a_2	b_2
a_1	$(+, +)$	$(0, 0)$
b_1	$(0, 0)$	$(+, +)$

In just the same way, the following two sign patterns, from games like Matching Coins are equivalent, Matching Coins

	H	T
H	$(+,-)$	$(0,0)$
T	$(0,0)$	$(+,-)$

Coins Matching

	H	T
H	$(-,+)$	$(0,0)$
T	$(0,0)$	$(-,+)$

Problem 3.1. *Show that all 2×2 games without ties are equivalent to one of the four categories identified at the beginning of this section* (p. 21).

4. The gap between equilibrium and Pareto rankings

The defining charateristic of an equilibrium is the mutual best response property. Pareto optimality arguments are very peculiar from the mutual best response point of view.

4.0.3. *Stag Hunt reconsidered.* An implication of Lemma 3.1 is that the following two versions of the Stag Hunt are strategically equivalent.

Stag Hunt

	Stag	Rabbit
Stag	(S,S)	$(0,R)$
Rabbit	$(R,0)$	(R,R)

Hunting Stag

	Stag	Rabbit
Stag	$(S-R,S-R)$	$(0,0)$
Rabbit	$(0,0)$	(R,R)

Remember that $S > R > 0$, which makes the Pareto ranking of the pure strategy equilibria in the first version of the game easy and clear. However, the Pareto rankings of the two pure strategy equilibria agree across the two versions of the game only if $S > 2R$. If $R < S < 2R$, then the Pareto criterion would pick differently between the equilibria in the two strategically equivalent games.

4.0.4. *Prisoners' Dilemma reconsidered.* An implication of Lemma 3.1 is that the following two versions of the Prisoners' Dilemma are strategically equivalent.

Prisoners' Dilemma

	Squeal	Silent
Squeal	$(1,1)$	$(0,0)$
Silent	$(0,0)$	$(-1,-1)$

Dilemma of the Prisoners

	Squeal	Silent
Squeal	$(-14,-14)$	$(0,-15)$
Silent	$(-15,0)$	$(-1,-1)$

If we take the Pareto criterion seriously, we feel very differently about the equilibria of these two games. In the first one, the unique equilibrium is the Pareto optimal feasible point, in the second, the unique equilibrium is (very) Pareto dominated.

5. Conclusions about Equilibrium and Pareto rankings

From these examples, we should conclude that the Pareto criterion and equilibrium have little to do with each other. This does not mean that we should abandon the Pareto criterion — the two versions of the Prisoners' Dilemma are equivalent only if we allow player i's choice to add 15 years of freedom to player $j \neq i$. Such a change does not change the strategic considerations, but it drastically changes the social situation being analyzed.

In other words: the difference between the two versions of the Prisoners' Dilemma is that we have stopped making one person's action, Squeal, have so bad an effect on the other person's welfare. One might argue that we have made the game less interesting by doing this. In particular, if you are interested in (say) understanding how people become socialized to pick the cooperative action when non-cooperation is individually rational but socially disastrous, the new version of the Prisoners' Dilemma seems to be no help whatsoever. The new version has synchronized social welfare (in the Pareto sense) and individual optimization.

My argument about socialization would be phrased in terms of changes to the utility functions, though not necessarily the changes given in Lemma 3.1. Utility functions are meant to represent preferences, and preferences are essentially indistinguishable from revealed preferences, that is, from choice behavior. If one thinks that both being Silent is the right outcome, then you need to change the preferences so that the players prefer being Silent. Socialization is one very effective way to change preferences. Many people feel badly if their actions harm others, even others they do not personally know. I take it that they have preferences that include consequences to others. To study socialization to cooperative actions, one needs to study how preferences are changed. Once again, much of the mileage in game theory arises from the contrast between different games.[1]

6. Risk dominance and Pareto rankings

One possible reaction to the previous section is "Yeah, yeah, that's all fine so far as the mathematics of equilibrium is concerned, but when I write down a game with specified payoffs, I really mean that those payoffs represent preferences, they are not merely devices for specifying best responses." If you take this point of view (or many points of view like it), analyzing Pareto optimality again makes sense.[2] However, if you take this point of view, you are stuck when you come to games in which the players disagree about which equilibrium is better. One way to try to resolve this is using the idea of risk dominance.

In some coordination games, we (might have) favored one equilibrium outcome over another because it was better for everyone. In the following game (with the same best response pattern as the Stag Hunt), Pareto ranking does not work,

	L	R
T	$(5,6)$	$(3,2)$
B	$(0,2)$	$(6,4)$

One idea that does work to pick a unique equilibrium for this game is called **risk dominance**. The two pure strategy equilibria for this game are $e^1 = (T, L)$ and $e^2 = (B, R)$. The set of σ_2 for which T, 1's part of e^1, is a best response for player 1 is $S_1^{e1} = \{\sigma_2 : \sigma_2(L) \geq 3/8\}$. The set of σ_2 for which B, 1's part of e^2, is a best response for player 1 is $S_1^{e^2} = \{\sigma_2 : \sigma_2(L) \leq 3/8\}$. Geometrically, S_{e^1} is a larger set than S_{e^2}. One way to interpret this is to say that the set of beliefs that 1 might hold that make 1's part of e^1 a best response is larger that the set that make 1's part of e^2 a best response. In this sense, it is "more likely" that 1 plays his/her part of e^1 than his/her part of e^2. Similarly, $S_2^{e^1} = \{\sigma_1 : \sigma_1(T) \geq 1/3\}$ is geometrically larger than

[1]The NYT article covering blood flows in the brain.

[2]The previous section is just a bad dream to be ignored while you get on with the serious business of proving that all works out for the best in this best of all possible worlds.

the set $S_2^{e^2} = \{\sigma_1 : \sigma_1(B) \leq 1/3\}$, so that it is "more likely" that 2 plays his/her part of e^1 than his/her part of e^2. This serves as a definition of **risk dominance**, e^1 risk dominates e^2.

What we have just seen is that it is possible to invent a principle that takes over when Pareto ranking does not pick between equilibria. There are at least two more problems to overcome before we can reach an argument for systematically picking a single equilibrium, even in the set of 2×2 games that we have been looking at.

(1) The two players may disagree about which equilibrium risk dominates as in

	L	R
T	$(5,6)$	$(3,5)$
B	$(0,2)$	$(6,4)$

which is the same as the previous game, except that 2's payoff to (T,R) has been changed from 2 to 5. The sets $S_1^{e^1}$ and $S_1^{e^2}$ are unchanged, but $S_2^{e^1} = \{\sigma_1 : \sigma_1(T) \geq 2/3\}$ and $S_2^{e^2} = \{\sigma_1 : \sigma_1(B) \leq 2/3\}$. Now e^1 risk dominates e^2 for 1 but e^2 risk dominates e^1 for 2.

(2) Risk dominance may disagree with the Pareto ranking, so we actually need to decide whether we believe more strongly in risk dominance than in Pareto ranking. Return to the Stag Hunt,

	Stag	Rabbit
Stag	(S,S)	$(0,R)$
Rabbit	$(R,0)$	(R,R)

where $S > R > 0$. While $(S,S)^T \gg (R,R)^T$, making S look good, for each hunter the Rabbit strategy looks less 'risky' in the sense that they are less dependent on the actions of the other. Arguing directly in terms of the risk dominance criterion, the Stag equilibrium risk dominates if $S > 2R$, while Rabbit risk dominates of $2R > S > R$. However, Stag always Pareto dominates.

Even if Pareto rankings do not survive the utility transformations of Lemma 3.1, risk dominance rankings do.

PROBLEM 3.2. *Suppose that a 2×2 game without ties, $\Gamma = (A_i, u_i)$, has two pure strategy equilibria, e^1 and e^2 and that e^1 risk dominates e^2. Suppose that $\Gamma' = (A_i, v_i)$ where the v_i are derived from the u_i using any of the transformations allowed in Lemma* 3.1. *We know that e^1 and e^2 are equilibria of Γ'. Show that e^1 risk dominates e^2 in Γ'.*

7. Idle Threats

In Puccini's *Gianni Schicchi*, Buoso Donati has died and left his large estate to a monastery. Before the will is read by anyone else, the relatives call in a noted mimic, Gianni Schicchi, to play Buoso on his deathbed, re-write the will, and then convincingly die. The relatives explain, very carefully, to Gianni Schicchi, just how severe are the penalties for anyone caught tampering with a will (at the time, the penalities included having one's hand cut off). The plan is put into effect, but, on the deathbed, Gianni Schicchi, as Buoso Donati, rewrites the will leaving the entire estate to the noted mimic and great artist, Gianni Schicchi. The relatives can expose him, *and thereby expose themselves too*, or they can remain silent. With player 1 being Gianni Schicchi and player 2 being the relatives, and some utility numbers with $y \gg 200$ to make the point, with Gianni Shicchi being player 1, we have

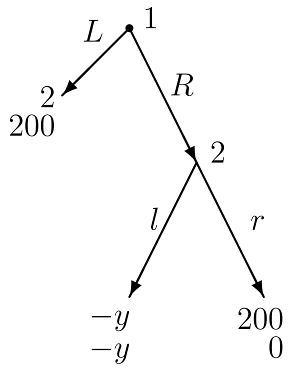

Gianni Shicchi

	l	r
L	$(2, 200)$	$(2, 200)$
R	$(-y, -y)$	$(200, 0)$

The box on the right gives the 2×2 **normal form** representation of the extensive form game. There are two equilibria. Compare the equilibrium sets when we do and when do not insist that people pick an action that is optimal for some beliefs at 0 probability events:

(1) If we insist, then 2 must play, or plan to play, r when they receive the signal that they should choose. This means that 1's best reply is R.

(2) If we do not insist, then 1 can play L in equilibrium. 1 playing L corresponds to 1 believing 2's threat. If we insist on optimality even at 0 probability events, we are insisting that 1 not believe in any threat that 2 would be unwilling to carry out.

8. Spying and Coordination

One of the interesting aspects of information is who has it and who chooses to have it. Information can be used to achieve coordination, and coordination is often crucial. One way to achieve coordination in two person strategic situations is to have one player "spy" on the other. This game is designed to get some of the simple implications of the possibility that one player may choose to expend the effort to spy on the other in a two person coordination game. The specific coordination game to be considered is

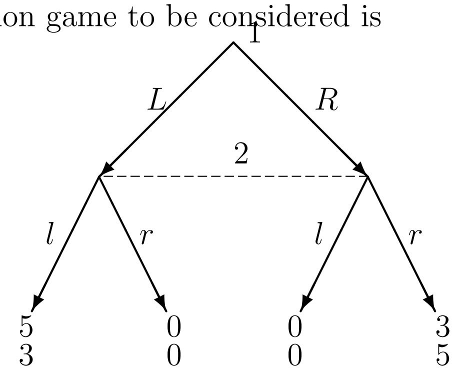

With the spying option, the game is altered so that player 2 can, before making their choice, pay a utility cost of $s > 0$ and learn what 1 has done. Specifically, the game is

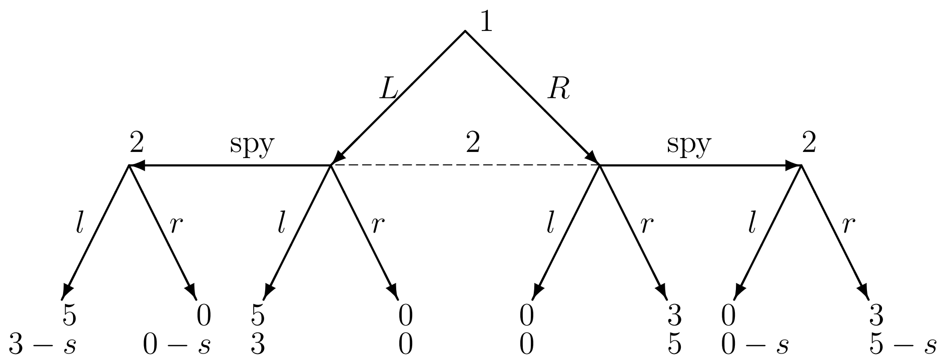

The normal form for this game is fairly large. A strategy in an extensive form game is best thought of as a **complete contingent plan**, it specifies what will happen at every eventuality. There is 1 place where the first player has to make a choice, and s/he has 2 choices. There are 3 places where 2 has to make a choice, and there are two choices at each, which means there are 2^3 complete contingent plans for player 2.

The coordination game has a mixed strategy equilibrium, the game with the possibility of spying does not, but spying does not happen in equilibrium. There is a lesson here: we never observe spying, so might be tempted not to include the possibility of spying in our description of the game; yet the possibility of spying changes the set of equilibria.

9. Problems

Problem 3.3. *Add the option of worker spying on the inspector, say tapping the inspector's phone line to find out where they are going on a particular day, to the inspection game with* $b = 10$, $c = 2$, $w = 12$, *and* $e = 3$, *and analyze the resulting game. Set the cost* s *to something positive, but small.*

Problem 3.4. *Add the option of inspector spying on the worker, say by checking on a secretly placed closed circuit tv, to the inspection game with* $b = 10$, $c = 2$, $w = 12$, *and* $e = 3$, *and analyze the resulting game. Set the cost* s *to something positive, but small.*

Problem 3.5 (A role for envy). [*This is loosely based on Koçkesen, Ok, and Sethi's "Evolution of Interdependent Preferences in Aggregative Games," Games and Economic Behavior, 31(2), 303-310, May 2000.*]

Firms $n \in \{1, \ldots, N\}$ *pick quantities* $q_n \in [0, 1]$. *When the vector* $\vec{q} = (q_1, \ldots, q_N)$ *has been picked, industry supply is* $q := \sum_n q_n$ *and industry price is* $p(q) = \max\{0, 1 - q\}$, *and the profits of firm* n *are given by*

$$\pi_n(\vec{q}) = q_n \cdot p(q).$$

For firms $m \in I_k = \{1, \ldots, k\}$, $1 \leq k < N$, *preferences are representable by any strictly increasing function of* π_m.

For any vector $\vec{q}$ *of quantities, define* $\overline{\pi} = \overline{\pi}(\vec{q}) = \frac{1}{N} \sum_{r=1}^{N} \pi_r(\vec{q})$ *as the average profits in the industry. For firms* $\ell \in J_k = \{k+1, \ldots, N\}$, *preferences are representable by the utility* *function*

$$u_\ell(\vec{q}) = \pi_\ell(\vec{q}) + \theta_\ell \left[F(\pi_\ell(\vec{q}), \overline{\pi}(\vec{q})) \right].$$

Assume that θ_ℓ is a small, strictly positive number. The function $F = F(x,y)$ is a twice continuously differentiable, concave envy function, capturing the idea that a firm $\ell \in J_k$ cares about their relative standing. Specifically, assume that $\partial F/\partial x > 0$, $\partial F/\partial y < 0$, $dF/d\pi_\ell = \partial F/\partial x + \frac{1}{N}\partial F/\partial y > 0$, and $d^2F/d\pi_\ell^2 < 0$.

(1) *For $0 < \pi_n < 1$, $n \in \{1,\ldots,N\}$, $F(\pi_\ell, \overline{\pi}) = \pi_\ell/\overline{\pi}$ is a twice continuously differentiable, concave envy function.*

For the following problems, do not assume that $F(\cdot,\cdot)$ has the special form just given.

(2) *An interior equilibrium has $q_n^* > 0$, $n \in \{1,\ldots,N\}$. Give the system of equations that characterize interior equilibria.*

(3) *At any interior equilibrium, for every $\ell \in J_k$ and any $m \in I_k$, $q_\ell^* > q_m^*$.*

(4) *At any interior equilibrium, for every $\ell \in J_k$ and any $m \in I_k$, $\pi_\ell(\vec{q}^*) > \pi_m(\vec{q}^*)$.*

(5) *The Cournot equilibrium is the equilibrium of the present game without envy (i.e. when $\theta_\ell = 0$ for $\ell \in J_k$). At an interior equilibrium with envy, industry profits are lower than at the Cournot equilibrium.*[3]

[3] *To summarize, "mutants" with the envying preferences do better than non-"mutants," but the population as a whole does worse.*

CHAPTER 4

Monotone Comparative Statics for Decision Theory

Readings: (Chambers and Echenique [**5**]), (Milgrom and Shannon [**13**]).

1. Overview

Throughout economics, we are interested in how changes in one variable affect another variable. We answer such questions **assuming** that what we observe is the result of optimizing behavior.[1] Given what we know about rational choice theory, optimizing behavior involves maximizing a utility function (in many of the interesting models we care about). In symbols, with $t \in T$ denoting a variable not determined by the individual, we let $x^*(t)$ denote the solution(s) to the problem $P(t)$,

$$\max_{x \in X} \ f(x,t), \tag{5}$$

and ask how $x^*(t)$ depends on t as t varies over T.

Since the problem $P(t)$ in (5) is meant as an approximation to rather than a quantitative representation of behavior, we are after "qualitative" results. These are results that are of the form "if t increases, then $x^*(t)$ will increase," and they should be "fairly immune to" details of the approximation.[2] If $x^*(\cdot)$ is differentiable, then we are after a statement of the form $dx^*/dt \geq 0$. If $x^*(\cdot)$ is not differentiable, then we are after a statement of the form that x^* is non-decreasing in t. We're going to go after such statements three ways.

Implicit function theorem: using derivative assumptions on f when X and T are 1-dimensional intervals.

Simple univariate Topkis's Theorem: using supermodularity assumptions on f when X and T are linearly ordered sets, e.g. 1-dimensional intervals.

Monotone comparative statics: using supermodularity when X is a lattice and T is partially ordered.

2. The Implicit Function Approach

Assume that X and T are interval subsets of $\mathbb{R}$, and that f is twice continuously differentiable on $X \times T$. f_x, f_t, f_{xx} and f_{xt} denote the corresponding partial derivatives of f. To

[1]Part of learning to "think like an economist" involves internalizing this assumption.

[2]Another part of learning to "think like an economist" involves developing an aesthetic sense of what "fairly immune" means. Aesthetics are complicated, subtle, and vary over time. An understanding of the currect aesthetic is best gained by immersion in and indoctrination by the culture of economists, as typically happens during graduate school.

have $f_x(x,t) = 0$ characterize $x^*(t)$, we must have $f_{xx} < 0$ (this is a standard result about concavity). From the implicit function theorem, we know that $f_{xx} \neq 0$ is what is needed for there to exist a function $x^*(t)$ such that

$$f_x(x^*(t), t) \equiv 0. \tag{6}$$

To find dx^*/dt, take the derivative on both sides with respect to t, and find

$$f_{xx}\tfrac{dx^*}{dt} + f_{xt} = 0, \tag{7}$$

so that $dx^*/dt = -f_{xt}/f_{xx}$. Since $f_{xx} < 0$, this means that dx^*/dt and f_{xt} have the same sign.

This ought to be intuitive: if $f_{xt} > 0$, then increases in t increase f_x; increases in f_x are increases in the marginal reward of x; and as the marginal reward to x goes up, we expect that the optimal level of x goes up. In a parallel fashion: if $f_{xt} < 0$, then increases in t decrease f_x; decreases in f_x are decreases in the marginal reward of x; and as the marginal reward to x goes down, we expect that the optimal level of x goes down.

Problem 4.1. *Let $X = T = \mathbb{R}_+$, $f(x,t) = x - \frac{1}{2}(x-t)^2$. Find $x^*(t)$ and verify directly that $dx^*/dt > 0$. Also find f_x, f_{xx}, and f_{xt}, and verify, using the sign test just given, that $dx^*/dt > 0$. If you can draw three dimensional figures (and this is a skill worth developing), draw f and verify from your picture that $f_{xt} > 0$ and that it is this fact that make $dx^*/dt > 0$. To practice with what goes wrong with derivative analysis when there are corner solutions, repeat this problem with $X = \mathbb{R}_+$, $T = \mathbb{R}$, and $g(x,t) = x - \frac{1}{2}(x+t)^2$.*

The following will help pinpoint the relation between sufficient conditions for monotone comparative statics result and being at an optimum.

Example 4.1. *The amount of a pollutant that can be emitted is regulated to be no more than $t \geq 0$. The cost function for a monopolist producing x is $c(x,t)$ with $c_t < 0$ and $c_{xt} < 0$. These derivative conditions means that increases is the allowed emission level lower costs and lower marginal costs, so that the firm will always choose t. For a given t, the monopolist's maximization problem is therefore*

$$\max_{x \geq 0} f(x,t) = xp(x) - c(x,t) \tag{8}$$

where $p(x)$ is the (inverse) demand function. Since $f_{xt} = -c_{xt}$, we know that increases in t lead the monopolist to produce more, ***provided*** *$f_{xx} < 0$.*

The catch in the previous analysis is that $f_{xx} = xp_{xx} + p_x - c_{xx}$, so that we need to know $p_{xx} < 0$, concavity of inverse demand, and $c_{xx} > 0$, convexity of the cost function, before we can reliably conclude that $f_{xx} < 0$. The global concavity of $f(\cdot, t)$ seems to have little to do with the intuition that the lowering of marginal costs makes x^* depend positively on t.

However, and this is the crucial point, the global concavity of $f(\cdot, t)$ is **not** what we need for the implicit function theorem, only the concavity of $f(\cdot, t)$ in the region of $x^*(t)$. This local concavity is an **implication** of the second order derivative conditions for $x^*(t)$ being a strict local maximum for $f(\cdot, t)$. What supermodularity does is to make it clear that the local maximum property is all that is being assumed, and to allow us to work with optima that are non-differentiable.

3. The simple supermodularity approach

The simplest case has X and T being linearly ordered sets. The most common example has X and T being intervals in $\mathbb{R}$ with the usual less-than-or-equal-to order. However, nothing rules out the sets X and T being discrete.

DEFINITION 4.2. *For linearly ordered X and T, a function $f : X \times T \to \mathbb{R}$ is* ***supermodular*** *if for all $x' \succ x$ and all $t' \succ t$,*

$$f(x', t') - f(x, t') \geq f(x', t) - f(x, t),$$

equivalently

$$f(x', t') - f(x', t) \geq f(x, t') - f(x, t).$$

It is ***strictly supermodular*** *if the inequalities are strict.*

At t, the benefit of increasing from x to x' is $f(x', t) - f(x, t)$, at t', it is $f(x', t') - f(x, t')$. This assumption asks that benefit of increasing x be increasing in t. A good verbal shorthand for this is that f **has increasing differences in** x **and** t. Three sufficient conditions in the differentiable case are: $\forall x$, $f_x(x, \cdot)$ is nondecreasing; $\forall t$, $f_t(\cdot, t)$ is nondecreasing; and $\forall x, t$, $f_{xt}(x, t) \geq 0$.

Another name for this condition is the **single crossing condition**, because for all $x' \succ x$, the function $h(t) := f(x', t) - f(x, t)$ is an increasing (or non-decreasing) function of t, hence, it crosses any value at most once, from below.

THEOREM 4.3. *If $f : X \times T \to \mathbb{R}$ is supermodular and $x^*(\tau)$ is the largest solution to $\max_{x \in X} f(x, \tau)$ for all τ, then $[t' \succ t] \Rightarrow [x^*(t') \succsim x^*(t)]$.*

If there are unique, unequal maximizers at t' and t, then $x^*(t') \succ x^*(t)$.

Proof: Suppose that $t' \succ t$ but that $x' = x^*(t') \prec x = x^*(t)$. Because $x^*(t)$ and $x^*(t')$ are maximizers, $f(x', t') \geq f(x, t')$ and $f(x, t) \geq f(x', t)$. Since x' is the largest of the maximizers at t' and $x \succ x'$, we know a bit more, that $f(x', t') > f(x, t')$. Adding the inequalities, we get $f(x', t') + f(x, t) > f(x, t') + f(x', t)$, or

$$f(x, t) - f(x', t) > f(x, t') - f(x', t').$$

But $t' \succ t$ and $x \succ x'$ and supermodularity imply that this inequality must go the other way. ∎

Going back to Example 4.1 (p. 32), we can substitute f in the relations of Definition 4.2 (p. 33). Then these relations for the supermodularity of f reduce to those for supermodularity of $-c$. Thus assuming $-c$ (and hence f) is supermodular, we can use Theorem 4.3 for f, which implies that $x^*(t)$ is increasing. None of the second derivative conditions except $c_{xt} < 0$ are necessary, and this can be replaced by the looser condition that $-c$ is supermodular.

Clever choices of T's and f's can make some analyses criminally easy.

EXAMPLE 4.4. *Suppose that the one-to-one demand curve for a good produced by a monopolist is $x(p)$ so that $CS(p) = \int_p^\infty x(r)\,dr$ is the consumer surplus when the price p is charged. Let $p(\cdot)$ be $x^{-1}(\cdot)$, the inverse demand function. From intermediate microeconomics, you should know that the function $x \mapsto CS(p(x))$ is nondecreasing.*

The monopolist's profit when they produce x is $\pi(x) = x \cdot p(x) - c(x)$ where $c(x)$ is the cost of producing x. The maximization problem for the monopolist is

$$\max_{x \geq 0} \; \pi(x) + 0 \cdot CS(p(x)).$$

Society's surplus maximization problem is

$$\max_{x \geq 0} \; \pi(x) + 1 \cdot CS(p(x)).$$

Set $f(x,t) = \pi(x) + tCS(p(x))$ where $X = \mathbb{R}_+$ and $T = \{0,1\}$. Because $CS(p(x))$ is nondecreasing, $f(x,t)$ is supermodular.[3] *Therefore $x^*(1) \geq x^*(0)$, the monopolist always (weakly) restricts output relative to the social optimum.*

Here is the externalities intuition: increases in x increase the welfare of people the monopolist does not care about, an effect external to the monopolist; the market gives the monopolist insufficient incentives to do the right thing.

To fully appreciate how much simpler the supermodular analysis is, we need to see how complicated the differentiable analysis would be.

Example 4.5. (↑*Example* 4.4 (p. 33)) *Suppose that for every $t \in [0,1]$, the problem*

$$\max_{x \geq 0} \; \pi(x) + t \cdot CS(p(x))$$

has a unique solution, $x^*(t)$, and that the mapping $t \mapsto x^*(t)$ is continuously differentiable. (This can be guaranteed if we make the right kinds of assumptions on $\pi(\cdot)$ and $CS(p(\cdot))$.) To find the sign of $dx^*(t)/dt$, we assume that the first order conditions,

$$\pi'(x^*(t)) + tdCS(p(x^*(t)))/dx \equiv 0$$

characterize the optimum. In general, this means that we need to assume that $x \mapsto \pi(x) + tCS(p(x))$ is a smooth, concave function. We then take the derivative of both sides with respect to t. This involves evaluating $d(\int_{p(x^*(t))}^{\infty} x(r)\,dr)/dt$. In general, $d(\int_{f(t)}^{\infty} x(r)\,dr)/dt = -f'(t)x(f(t))$, so that when we take derivatives on both sides, we have

$$\pi''(x^*(t))(dx^*/dt) + dCS(p(x^*(t)))/dx - p'(x^*)(dx^*/dt)x(p(x^*)) = 0.$$

Gathering terms, this yields

$$[\pi''(x^*) - p'(x^*)x(p(x^*))](dx^*/dt) + dCS(p(x^*(t)))/dx = 0.$$

Since we are assuming that we are at an optimum, we know that $\pi''(x^*) \leq 0$, by assumption, $p'(x^*) < 0$ and $x > 0$, so the term in the square brackets is negative. As argued above, $dCS(p(x^*(t)))/dx > 0$. Therefore, the only way that the last displayed equation can be satisfied is if $dx^*/dt > 0$. Finally, by the fundamental theorem of calculus (which says that the integral of a derivative is the function itself), $x^*(1) - x^*(0) = \int_0^1 \frac{dx^*(r)}{dt}\,dr$. The integral of a positive function is positive, so this yields $x^*(1) - x^*(0) > 0$.

[3] This is an invitation/instruction to check this last statement.

4. POSETs and Lattices

"POSET" stands for "partially ordered set." Many of the partially ordered sets that we care about have the property that there is a minimal element greater and a maximal element less than any two non-ordered points. Those are the lattices.

4.1. Definitions.

Definition 4.6 (POSETs). *$(X, \precsim)$ is a **partially ordered set (POSET)** if for all $x, y, z \in X$, $x \precsim x$ (reflexivity), $x \precsim y$ and $y \precsim x$ imply $x = y$ (antisymmetry), and, if $x \precsim y$ and $y \precsim z$, then $x \precsim z$ (transitivity).*

Definition 4.7 (Lattices). *A partially ordered $(X, \precsim)$ is a **lattice** if each pair of elements, $x, y \in X$, has a unique **greatest lower bound**, $\min\{x, y\}$, sometimes written $x \wedge y$, and a unique **least upper bound**, $\max\{x, y\}$, sometimes written $x \vee y$. These are defined by*

(1) $z = (x \wedge y)$ is the unique element of X such that, if $z \precsim x$, $z \precsim y$, and $z \precsim z'$ then either $z' \not\precsim x$ or $z' \not\precsim y$, and

(2) $z = (x \vee y)$ is the unique element of X such that, if $x \precsim z$, $y \precsim z$, and $z' \precsim z$ then either $x \not\precsim z'$ or $y \not\precsim z'$.

Sometimes the greatest lower bound of a pair of elements is known as the inf or **infimum** of the pair and the least upper bound of a pair of elements is known as the sup or **supremum** of the pair.

It is antisymmetry that distinguishes partial orderings from the preference orderings studied in the first semester of micro. The notation, however, is exactly the same, so be wary. The uniqueness of $x \wedge y$ and $x \vee y$ are important.

Sometimes every set has an infimum and a supremum, not just every two point set. Such lattices are modified by the adjective "complete."

Definition 4.8. *A lattice $(X, \precsim)$ is **complete** if every $S \subset X$ has an infimum and a supremum.*

Definition 4.9. *A partial order that is also a complete ordering is called a **total (or linear) ordering**, and $(X, \precsim)$ is called a **totally (or linearly) ordered set**. A **chain** in a partially ordered set is a subset, $X' \subset X$, such that $(X', \precsim)$ is totally ordered.*

Note that when the adjective "complete" modifies an order, it means something different than when it modifies a lattice.

The classical example of a linearly ordered set is $(\mathbb{R}, \leq)$. In a total ordering, any two elements x and y in A can be compared, in a partial ordering, there are noncomparable elements.

4.2. Some Classical Examples of Lattices.

Here are two classic examples of lattices.

Example 4.10. *$(X, \precsim) = (\mathbb{R}^n, \leq)$ with $x \leq y$ iff $x_i \leq y_i$, $i = 1, \ldots, n$ is a lattice with $x \wedge y = (\min(x_1, y_1), \ldots, \min(x_n, y_n))$ and $x \vee y = (\max(x_1, y_1), \ldots, \max(x_n, y_n))$. X is not a complete lattice, but $([0,1]^n, \leq)$ is, with $\inf S = (\inf\{x_1 : x \in S\}, \inf\{x_2 : x \in S\})$ and $\sup S = (\sup\{x_1 : x \in S\}, \sup\{x_2 : x \in S\})$.*

EXAMPLE 4.11. *Suppose that* $X \neq \emptyset$*, let* $\mathcal{P}(X)$ *denote the class of subsets of* X*, and for* $A, B \in \mathcal{P}(X)$*, define* $A \precsim B$ *if* $A \subset B$*.* $(\mathcal{P}(X), \precsim)$ *is a complete lattice, with* $\inf \mathcal{S} = \bigcap\{E : E \in \mathcal{S}\}$ *for* $\mathcal{S} \subset \mathcal{P}(X)$ *and* $\sup \mathcal{S} = \bigcup\{E : E \in \mathcal{S}\}$*.*

We can turn lattice orderings around and keep the lattice structure.

EXAMPLE 4.12. *Suppose that* $(X, \precsim)$ *is a lattice. Define* $x \precsim^- y$ *if* $y \precsim x$*. Show that* $(X, \precsim^-)$ *is a lattice. [Thus,* $(\mathbb{R}^n, \leq^-)$*, i.e.* $(\mathbb{R}^n, \geq)$*, is a lattice, as is* $(\mathcal{P}(X), \supset)$*.]*

EXAMPLE 4.13. $(\mathcal{P}(\mathbb{N}), \subseteq)$ *is a lattice with many noncomparable pairs of elements, but the class of sets* $\{\{1\}, \{1,2\}, \ldots, \{1,2,\ldots,n\}, \ldots\}$ *is a chain in* $\mathcal{P}(\mathbb{N})$*.*

PROBLEM 4.2. *Show that if* $A \subset B$ *and* B *is totally ordered, then* A *is totally ordered. As an application, show that, in* $(\mathbb{R}^2, \leq)$*, any subset of the graph of a non-decreasing function from* $\mathbb{R}$ *to* $\mathbb{R}$ *is a chain.*

There are partially ordered sets that are **not** lattices.

EXAMPLE 4.14. *Let* $T_1 \subset \mathbb{R}^2$ *be the triangle with vertices at* $(0,0)$*,* $(1,0)$ *and* $(0,1)$*.* $(T_1, \leq)$ *is partially ordered, but is not a lattice because* $(0,1) \vee (1,0)$ *is not in* T_1*. On the other hand,* $(T_2, \leq)$ *is a lattice when* T_2 *is the square with vertices at* $(1,1)$*,* $(1,0)$*,* $(0,1)$*, and* $(0,0)$*.*

5. Monotone comparative statics

We now generalize in two directions:

(1) we now allow for X and T to have more general properties, they are not just linearly ordered in what follows, and
(2) we now also allow for the set of available points to vary with t, not just the utility function.

5.1. Product Lattices. Suppose that $(X, \precsim_X)$ and $(T, \precsim_T)$ are lattices. Define the order $\precsim_{X \times T}$ on $X \times T$ by $(x', t') \succsim_{X \times T} (x, t)$ iff $x' \succsim_X x$ and $t' \succsim_T t$. (This is the unanimity order again.)

LEMMA 4.15. $(X \times T, \precsim_{X \times T})$ *is a lattice.*

Proof: $(x', t') \vee (x, t) = (\max\{x', x\}, \max\{t', t\}) \in X \times T$.
$(x', t') \wedge (x, t) = (\min\{x', x\}, \min\{t', t\}) \in X \times T$. ∎

5.2. Supermodular Functions.

DEFINITION 4.16. *For a lattice* $(L, \precsim)$*,* $f : L \to \mathbb{R}$ *is* ***supermodular*** *if for all* $\ell, \ell' \in L$*,*

$$f(\ell \wedge \ell') + f(\ell \vee \ell') \geq f(\ell) + f(\ell'),$$

equivalently,

$$f(\ell \vee \ell') - f(\ell') \geq f(\ell) - f(\ell \wedge \ell').$$

EXAMPLE 4.17. *Taking* $\ell' = (x', t)$ *and* $\ell = (x, t')$ *recovers Definition* 4.2*.*

PROBLEM 4.3. *Show that a monotonic convex transformation of a monotonic supermodular function is supermodular.*

Problem 4.4. *Let* $(L, \precsim) = (\mathbb{R}^n, \leq)$*. Show that* $f : L \to \mathbb{R}$ *is supermodular iff it has increasing differences in* x_i *and* x_j *for all* $i \neq j$*. Show that a twice continuously differentiable* $f : L \to \mathbb{R}$ *is supermodular iff* $\partial^2 f/\partial x_i \partial x_j \geq 0$ *for all* $i \neq j$*.*

5.3. Ordering Subsets of a Lattice.

Definition 4.18. *For* $A, B \subset L$*,* L *a lattice, the* ***strong set order*** *is defined by* $A \precsim_{Strong} B$ *iff* $\forall (a, b) \in A \times B$*,* $a \wedge b \in A$ *and* $a \vee b \in B$*.*

Recall that interval subsets of $\mathbb{R}$ are sets of the form $(-\infty, r)$, $(-\infty, r]$, (r, s), $(r, s]$, $[r, s)$, $[r, s]$, (r, ∞), or $[r, \infty)$.

Problem 4.5. *For intervals* $A, B \subset \mathbb{R}$*,* $A \precsim_{Strong} B$ *iff every point in* $A \setminus B$ *is less than every point in* $A \cap B$*, and every point in* $A \cap B$ *is less than every point in* $B \setminus A$*. Show that this is also true when* $(\mathbb{R}, \leq)$ *is replaced with any linearly ordered set.*

Problem 4.6. *Let* $T_1 \subset \mathbb{R}^2$ *be the triangle with vertices at* $(0, 0)$*,* $(1, 0)$ *and* $(0, 1)$*, and let* L *be the line segment connecting the points* $(\frac{1}{2}, \frac{1}{2})$ *and* $(1, 1)$*. Find the smallest* B *containing* L *such that* $T_1 \precsim_{Strong} B$*. How does the answer change if* L *is the line segment connecting the points* $(0, 0)$ *and* $(1, 1)$*?*

The strong set order is not, in general, reflexive. Subsets of $(\mathbb{R}, \leq)$ are linearly ordered, hence they are a lattice. Subsets of $(\mathbb{R}^2, \leq)$ are not necessarily lattices. For any non-lattice subset, A, of $\mathbb{R}^2$, $\neg[A \succsim_{Strong} A]$.

5.4. The Supermodularity Theorem. For $t \in T$, let $M(t)$ be the set of solutions to the problem $\max_\ell f(\ell, t)$, and let $M(t, S)$ be the set of solutions to the problem $\max_{\ell \in S} f(\ell, t)$, i.e. so that $M(t) = M(t, L)$.

Theorem 4.19. *If* $(L, \precsim_L)$ *is a lattice,* $(T, \precsim_T)$ *is a partially ordered set,* $f : L \times T \to \mathbb{R}$ *is supermodular in* ℓ *for all* t*, and has increasing differences in* ℓ *and* t*, then* $M(t, S)$ *is nondecreasing in* (t, S)*.*

Here is an easy observation and its implication. If $(X, \precsim_X)$ is linearly ordered, then **any** function $f : X \to \mathbb{R}$ is supermodular (with equalities in Definition 4.16). This means that if L is any subset of $\mathbb{R}$ in Theorem 4.19, then increasing differences is all that is needed for the monotonicity of the set of solutions.

Problem 4.7. *This asks you to prove Theorem* 4.19 *in a series of steps.*

(1) First, prove to yourself that the result is true in the case that S *is fixed at* L *and* $M(t)$ *is a singleton set.*
(2) Pick $(t', S') \succsim (t, S)$ *[so that we must show that* $M(t', S') \succsim_{Strong} M(t, S)$*].*
(3) Show that $\varnothing \succsim_{Strong} A \succsim_{Strong} \varnothing$ *for any set* A*.*
(4) Show that if $M(t', S') = \varnothing$ *or* $M(t, S) = \varnothing$*, the result is true, so that "all" that is left is to show is that the result is true when* $M(t', S')$ *and* $M(t, S)$ *are non-empty.*
(5) Pick $\ell' \in M(t', S') \subset S'$ *and* $\ell \in M(t, S) \subset S$*. By the definition of the strong set order, we need to show that* $\ell \vee \ell' \in M(t', S')$ *and* $\ell \wedge \ell' \in M(t, S)$*.*

Since $S' \succsim_{Strong} S$, $\ell' \wedge \ell \in S'$ and $\ell' \vee \ell \in S$. Because ℓ is optimal in S and $\ell' \wedge \ell \in S$, we know that $f(\ell, t) - f(\ell' \wedge \ell) \geq 0$. Show that combining the supermodularity of $f(\cdot, t)$ and this last inequality yields

$$f(\ell \vee \ell', t) - f(\ell', t) \geq f(\ell, t) - f(\ell \wedge \ell', t) \geq 0.$$

(6) *Show that increasing differences, $t' \succsim_T t$, and $\ell \vee \ell' \succsim_L \ell'$ and the previous inequality yield*

$$f(\ell \vee \ell', t') - f(\ell', t') \geq f(\ell \vee \ell', t) - f(\ell', t) \geq 0.$$

(7) *Since ℓ' is optimal in S' and $\ell \vee \ell' \in S'$, we have just discovered something about the optimality of $\ell \vee \ell'$ in S'. Show that what we have discovered completes this step of the proof.*

(8) *Finish the proof.*

6. Cones and Lattice Orderings

A number of useful lattices arise by defining $x \precsim y$ by $L(x) \leq L(y)$ for all L is a cone of linear functions.

6.1. Cone-Based Lattice Orderings on $\mathbb{R}^n$. Sometimes different lattice orderings in $\mathbb{R}^n$ are useful. The starting point is to notice that $\{y : x \leq y\}$ has a nice geometric shape, that of a convex cone opening up and to the right with its vertex at the point x. Essentially, we are going to shift the vertex to 0 by a sleight of hand, and see that all kinds of convex cones to define lattice orderings.

PROBLEM 4.8. *Let $A := \{a_m : m \in \{1, \ldots, M\}\}$ be a collection of vectors in $\mathbb{R}^n$. Define $x \leq_A y$ if $a_m \cdot x \leq a_m \cdot y$ for all $m \in \{1, \ldots, M\}$. With $A = \{e_i : i \in \{1, \ldots, n\}\}$ being the set of unit vectors, this recovers the ordering in the previous problem.*

(1) *Show that if A is a spanning, linearly independent collection of vectors, then $(\mathbb{R}^n, \leq_A)$ is a lattice.*

(2) *More generally, we define $C \subset \mathbb{R}^n$ to be a* ***pointy, closed, convex cone*** *if (i) it is closed, (ii) it is a convex cone, that is, for all $x, y \in C$ and all $\alpha, \beta \geq 0$, $\alpha x + \beta y \in C$, and (iii), 0 is the point of the cone, that is, if $x, y \in C$ and $0 = \alpha x + (1 - \alpha) y$ for some $\alpha \in [0, 1]$, then either $x = 0$ or $y = 0$.*
   (a) *Show that if A is a linearly independent collection of vectors, then $C = \{x \in \mathbb{R}^n : \forall a \in A,\ x \cdot a \leq 0\}$ is a pointy closed convex cone.*
   (b) *Show that if C is a pointy closed convex that spans $\mathbb{R}^n$ and we define $x \leq_C y$ if $(x - y) \in C$, then $(\mathbb{R}^n, \leq_C)$ is a lattice.*
   (c) *Let $x \neq 0$ be a point in $\mathbb{R}^n$ and $r > 0$ such that $\|x\| > r$. Show that there is a smallest closed cone with a point containing $\{y : \|y - x\| \leq r\}$, and that it spans $\mathbb{R}^n$. [For $n \geq 3$, this is a cone that is* ***not*** *defined by finitely many linear inequalities.]*

(3) *Give conditions on A such that the ordering $\leq_A$ is trivial, that is, $x \leq_A y$ iff $x = y$.*

(4) *For A a spanning, linearly independent set and a twice continuously differentiable $f : \mathbb{R}^n \to \mathbb{R}$, give a condition on the derivatives that is equivalent to f being supermodular on the lattice $(\mathbb{R}^n, \leq_A)$. [Doing this in $\mathbb{R}^2$ is enough.]*

Problem 4.9. *Let $Non \subset \mathbb{R}^n$ be the set of vectors $x' = (x_1, \ldots, x_n)$ with the property that $x_1 \leq x_2 \leq \cdots \leq x_n$, that is, the set of all vectors with non-decreasing components. Define $x \geq_{Non} y$ if $(x - y) \in Non$, equivalently, $x \leq_{Non} y$ if $(y - x) \in Non$.*

(1) *Show that Non is a convex cone in $\mathbb{R}^n$, that is, for all $x, y \in Non$ and all $\alpha, \beta \geq 0$, $\alpha x + \beta y \in Non$.*
(2) *Show that $(\mathbb{R}^n, \leq_{Non})$ is a* ***not*** *a lattice because Non is not pointy.*
(3) *For $x, y \in \mathbb{R}^2$, draw Non and show that $\min\{x, y\}$ and $\max\{x, y\}$ are not defined.*
(4) *In $\mathbb{R}^2$, $x \leq_{Non} y$ iff $x \leq_A y$ for some set of vectors A. Give the set A. Generalize to $\mathbb{R}^n$. [Here you should see that $\leq_{Non}$ is not a lattice ordering because the set A does not span.]*

6.2. Cone-Based Lattice Orderings on Sets of Functions. A function defined on a set of functions is often called a **functional**. The idea of ordering functions using cones of linear functionals extends what we did in $\mathbb{R}^n$.

Example 4.20. *For each $f \in C[0, 1]$ and $t \in [0, 1]$, define $L_t(f) = f(t)$. For each t, the function $f \mapsto L_t(f)$ is continuous and linear, as is the function $f \mapsto \alpha \cdot L_t(f)$ for $\alpha \in \mathbb{R}$. Let $\mathcal{C} = \{\sum_{k \leq K} \alpha_k \cdot L_{t_k} : K \in \mathbb{N},\ \alpha_k \geq 0,\ t_k \in [0, 1]\}$. It is easy to check that $\mathcal{C}$ is a cone of continuous linear functionals. We define $f \leq g$ be $L(f) \leq L(g)$ for all $L \in \mathcal{C}$. This is a fancy way of understanding the pointwise ordering of functions, $f \leq g$ iff $f(t) \leq g(t)$ for all $t \in [0, 1]$.*

Here is another cone-based way to order functions.

Problem 4.10. *Let $\mathcal{N}$ be the set of non-decreasing functions from $[0, 1]$ to $\mathbb{R}$ with $u(0) = 0$.*

(1) *Show that $\mathcal{N}$ is a pointy convex cone and give its point.*
(2) *Define $f \geq_{\mathcal{N}} g$ if $(f - g) \in \mathcal{N}$, equivalently, $f \leq_{\mathcal{N}} g$ if $(g - f) \in \mathcal{N}$.*
(3) *Show that $(C[0, 1], \leq_{\mathcal{N}})$ is a lattice and explicitly given $f \wedge g$ and $f \vee g$.*

6.3. Cone-Based Lattice Orderings on Probability Distributions. $\Delta(\mathbb{R})$ denotes the set of probability distributions on $\mathbb{R}$. A basic result is that these are uniquely identified by their cdfs. The following should be (mostly) review.

$C(\mathbb{R})$ is the set of continuous functions on $\mathbb{R}$, $C_b(\mathbb{R})$ is the set of continuous bounded functions, and for probabilities P_n and P, $P_n \to_w P$ iff $\int f\, dP_n \to \int f\, dP$ for all $f \in C_b(\mathbb{R})$. In particular, if u is a bounded continuous expected utility function, the expected utility function is continuous.

$C_{con} \subset C(\mathbb{R})$ the set of concave functions, which contains no bounded elements except the constant functions.

ND the set of non-decreasing functions from $\mathbb{R}$ to $\mathbb{R}$, whether or not they are continuous, ND_b the set of bounded non-decreasing functions, and $\mathcal{N} \subset ND_b$ is the set of functions with $\lim_{t \downarrow -\infty} f(t) = 0$.

Example 4.21. *All of these sets of functions, $C(\mathbb{R})$, $C_b(\mathbb{R})$, C_{con}, ND, ND_b and $\mathcal{N}$, are convex cones. However, only $\mathcal{N}$ is pointy.*

Definition 4.22. *For $P, Q \in \Delta$,*

(1) *P* ***first order stochastically dominates*** *Q, $P \succsim_{FOSD} Q$, if for all $u \in \mathcal{N}$, $\int u\,dP \geq \int u\,dQ$,*
(2) *P* ***is riskier than*** *Q if if for all $u \in C_{con}$, $\int u\,dP \leq \int u\,dQ$ (note the change in the direction of the inequality here), and*
(3) *P* ***second order stochastically dominates*** *Q, $P \succsim_{SOSD} Q$, if for all $u \in (C_{con} \cap ND)$, $\int u\,dP \geq \int u\,dQ$.*

We will see the multi-dimensional version of "riskier" when we study (one of the many results called) Blackwell's Theorem.

Some useful observations:

(1) $C_{con} \cap ND_b$ contains only the constant functions.
(2) $P \succsim_{FOSD} Q$ iff for all $r \in \mathbb{R}$, $P(-\infty, r] \leq Q(-\infty, r]$, i.e. iff $F_P \leq F_Q$ where F_P, F_Q are the cdfs associated with P, Q and "$\leq$" is the pointwise order.
(3) $P \succsim_{FOSD} Q$ iff if for all $u \in ND$, $\int u\,dP \geq \int u\,dQ$, but now you need to allow for infinite integrals.
(4) If P is the distribution of a random variable X and Q the distribution of Y, then $P \succsim_{FOSD} Q$ iff there are random variables X', Y' with distributions P, Q, such that $X' \geq Y'$ almost everywhere. [This uses the probability integral transform.]
(5) Sometimes people restrict second order stochastic dominance to pairs of distributions with the same mean.

7. Quasi-supermodularity

Sometimes people make the mistake of identifying supermodular utility functions as the ones for which there are complementarities. This is wrong.

Problem 4.11. *For $(x_1, x_2) \in \mathbb{R}^2_{++}$, define $u(x_1, x_2) = x_1 \cdot x_2$ and $v(x_1, x_2) = \log(u(x_1, x_2)$.*

(1) *Show that $\partial^2 u/\partial x_1 \partial x_2 > 0$.*
(2) *Show that $\partial^2 v/\partial x_1 \partial x_2 = 0$.*
(3) *Find a monotonic transformation, f, of v such that $\partial^2 f(v)/\partial x_1 \partial x_2 < 0$ at the point $(x_1^\circ, x_2^\circ) \in \mathbb{R}^2_{++}$.*

The problem is that supermodularity is not immune to monotonic transformations. That is, supermodularity, like expected utility theory, is a cardinal rather than an ordinal theory. Milgrom and Shannon fixed that problem.

Definition 4.23. *A function $u : X \to \mathbb{R}$ is* ***quasisupermodular*** *on the lattice X if, $\forall x, y \in X$,*

$$[u(x) \geq u(x \wedge y)] \Rightarrow [u(x \vee y) \geq u(y)], \text{ and}$$
$$[u(x) > u(x \wedge y)] \Rightarrow [u(x \vee y) > u(y)].$$

By way of contrast, from Definition 4.16 (p. 36), $f : X \to \mathbb{R}$ is supermodular if, $\forall x, y \in X$, $f(x \vee y) + f(x \wedge y) \geq f(x) + f(y)$, which directly implies that it is quasi-supermodular. From micro, you should remember that a monotonically increasing transformation of a concave utility function is quasi-concave.

LEMMA 4.24. *A monotonic increasing transformation of a supermodular function is quasi-supermodular.*

PROBLEM 4.12. *Prove Lemma* 4.24.

Recall that a binary relation, $\precsim$, on a set X has a representation $u : X \to \mathbb{R}$ iff $[x \precsim y] \Leftrightarrow [u(x) \leq u(y)]$. For choice theory on **finite** lattices with monotonic preferences, quasi-supermodularity and supermodularity of preferences are indistinguishable.

THEOREM 4.25 (Chambers and Echenique). *A binary relation on a finite lattice X has a weakly increasing and quasi-supermodular representation iff it has a weakly increasing and supermodular representation.*

Proof: Since superm implies q-superm, we need only show that a weakly increasing q-superm representation can be monotonically transformed to be superm. Let u be q-superm, set $u(X) = \{u_1 < u_2 < \cdots < u_N\}$, and define $g(u_n) = 2^{n-1}$, $n = 1, \ldots, N$. Show that the function $v(x) = g(u(x))$ is superm. ∎

PROBLEM 4.13. *There are three results that you should know about the relation between monotonicity and quasi-supermodularity.*

(1) *Strong monotonicity implies quasi-supermodularity, hence supermodularity: Show that if a binary relation $\precsim$ on a finite lattice X has a strictly increasing representation, then that representation is quasi-supermodular.* [*By the Chambers and Echenique result above, this implies that the binary relation has a supermodular representation.*]
(2) *Weak monotonicity does not imply quasi-supermodularity: Let $X = \{(0,0),(0,1),(1,0),(1,1)\}$ with $(x,y) \precsim (x',y')$ iff $(x,y) \leq (x',y')$. Show that the utility function $u(x,y) = 0$ if $x = y = 0$ and $u(x,y) = 1$ otherwise is weakly monotonic, but not monotonic transformation of it is q-supermodular.*
(3) *Supermodularity does not imply monotonicity: Let $X = \{(0,0),(0,1),(1,0),(1,1)\}$ with $(x,y) \precsim (x',y')$ iff $(x,y) \leq (x',y')$. Show that the utility function $u(0,0) = 0$, $u(0,1) = -1$, $u(1,0) = 2$ and $u(2,2) = 1.5$ is strictly supermodular but not monotonic.*

8. More Practice with Orderings and Comparative Statics

PROBLEM 4.14 (Tragey of the Commons (again) ↑§11). *There are I different countries that can put out fishing fleets to catch from pelagic schools of fish. Use the number $a_i \in \mathbb{R}_+ = A_i$ to represent the number of fishing boats in the fleet of country i, $i = 1, \ldots, I$. To finish specifying the game, the utilities, $u_i : A \to \mathbb{R}$ need to be specified.*

The marginal cost of a boat is constant, and equal to c. For given $a \in A$, let $n = n(a) = \sum_{i \in I} a_i$ and $n_{-i}(a) = \sum_{j \neq i} a_j$. When the total number of boats is n, the **per boat** *return is $v(n)$ where $v(0) > c > 0$, $v'(n) < 0$, and $v''(n) < 0$. For country i, the benefit to putting out a fleet depends on the size of their own fleet, a_i, and the size of the other countries' fleets, $n_{-i}(s)$,*

$$u_i(a_i, n_{-i}(a)) = a_i v(a_i + n_{-i}(a)) - ca_i = a_i v(n(a)) - ca_i.$$

For fixed $n_{-i} = n_{-i}(a)$, $u_i(\cdot, n_{-i})$ *is concave because* $\frac{\partial u_i(a_i,n_{-i})}{\partial a_i} = v(n) + a_i v'(n) - c$, *implying*

$$\frac{\partial^2 u_i(a_i, n_{-i})}{\partial a_i^2} = v'(n) + a_i v''(n) + v'(n) < 0.$$

(1) *Give the system equations that must be satisfied at any interior equilibrium. Show that they imply that the equilibrium number,* n_I^*, *of boats when there are* I *countries in the game satisfies* $v(n) + \alpha n v'(n) = c$ *where* $\alpha = \frac{1}{I}$.
(2) *Show that the equilibrium with* $I \geq 2$ *is Pareto inefficient.*
(3) *More generally, show that as* $\alpha \downarrow$, *the total benefit of a fleet of size* n_α *decreases, where* n_α *satisfies* $v(n) + \alpha n v'(n) = c$.

There are many other lattice orderings on $C[0,1]$. The standard partial ordering for $f, g \in C[0,1]$ is $f \leq g$ iff $f(x) \leq g(x)$ for all $x \in [0,1]$. This is the ordering used in the lattice version of the Stone-Weierstrass Theorem, and we saw it using cones of continuous linear functionals above. Here we introduce another partial ordering, useful for a class of mechanism design problems.

Definition 4.26. $f \succsim g$ *iff* $f \leq g$ *and for all* $0 \leq r < s \leq 1$, $f(s) - f(r) > g(s) - g(r)$.

Intuitively, the order is $f \succsim g$ iff "f is below and steeper than g." With concave $u(\cdot)$ switching to a reward function that is below and steeper makes the person "hungrier," because their marginal utility is higher, and more hard-working, because their reward to more effort is steeper. We will only apply this order to non-decreasing functions.

Problem 4.15. *If* $f(x) = x$ *and* $g(x) = x^2$, *find* $f \vee g$ *and* $f \wedge g$.

Problem 4.16. *Suppose that for each* $e \in E \subset \mathbb{R}$, F_e *is the cdf of a distribution on* $[0,1]$, *that is,* $F_e(r) = 0$ *for all* $r < 0$, *and* $F_e(1) = 1$. *Further suppose that for each* $e < e'$, $e, e' \in E \subset \mathbb{R}$, $F_{e'}$ *first order stochastically dominates* F_e.

For any non-decreasing reward/wage function $w : [0,1] \to \mathbb{R}_+$, *consider the problem* $\max_{e \in E} \int u(w(\pi))\, dF_e(\pi) - v(e)$ *where* $u(\cdot)$ *is an increasing, concave function, and* $v(\cdot)$ *is a non-decreasing function. Let* $e^*(w)$ *denote the set of optimal* e*'s for a given* w.

Show that $w_1 \succsim w_2$ *implies that* $e^*(w_1)$ *dominates* $e^*(w_2)$ *in the set order.*

Throughout all of this, remember that supermodularity assumptions are sufficient for monotone results, but they are NOT necessary.

Problem 4.17. X_a *is the agent's random income depending on their action* $a \geq 0$. *The distribution of* X_a *is* $R_a = cQ_a + (1-c)\mu$. *Here,* μ *does not depend on* a, *and if* $a > a'$, *then* Q_a *first order stochastically dominates* $Q_{a'}$. *The parameter* c *is the amount of "control" that the agent has,* $c = 0$ *and they have no control,* $c = 1$, *and they have the most possible. The question is how* a^* *depends on* c. *Intuitively, increasing* c *ought to increase the optimal action, however …*

The agent maximizes $f(a,c) = Eu(X_a - a)$, $u(\cdot)$ *strictly increasing and strictly concave. Increases in* a *pull down* $X_a - a$ *by increasing the direct cost* a, *but increase* $X_a - a$ *by stochastically increasing* X_a. *These two effects mean that* $f(\cdot,\cdot)$ *is not supermodular. But think about* *taking derivatives —*

$$f_a = -\int u'(x-a)d\mu(x) + cd/da[\textit{messy term with } Q_a \textit{ and } \mu].$$

(1) *Give a utility function* $u(\cdot)$, *a mapping* $a \mapsto Q_a$, *and a* μ *such that* $f(a,c)$ *fails to be supermodular.*
(2) *Check that if* $f_a = 0$, *the derivative of the messy term at the optimum is* > 0.
(3) *Show that* $a^*(c)$ *is increasing.*

CHAPTER 5

Monotone Comparative Statics and Related Dynamics for Game Theory

Readings: (Milgrom and Roberts [**12**]), (Conlisk [**6**]), (Hart and MasColell [**9**]). Also, look at the sections on dominant strategies, rationalizable strategies, and correlated equilibria in the textbook, [**8**].

1. Some examples of supermodular games

Van Huyck, Battalio and Beil [**19**] studied variants the following type of game: $A_i = \{1, 2, \ldots, A\}$, $S_n(a) = S_n(a_1, \ldots, a_I) = n$'th order statistic, $u_i(a_i, a_{-i}) = S_n(a) - c(a)$ where $c(\cdot)$ is non-decreasing. Note that

(1) $\forall i \in I,\ (A_i, \geq)$ is a complete lattice,
(2) $\forall i \in I,\ u_i$ is bounded above, and every kind of continuous that you can think of because the A_i's are finite,
(3) $\forall i \in I$ and each a_{-i}, $u_i(\cdot, a_{-i})$ is supermodular,
(4) $\forall i \in I$, $u_i(\cdot, \cdot)$ has increasing differences in a_i and a_{-i}.

This game has multiple, Pareto ranked, pure strategy equilibria.

Fable of the Bees #1: Each $i \in I$ chooses an action $a_i \in \mathbb{R}_+$ to solve

$$\max_{a_i \in [0,M]} u_i(a_i, \overline{a}) - c_i a_i$$

where $c_i > 0$, u_i is monotonic in both arguments, $\overline{a} = \frac{1}{I}\sum_{i \in I} a_i$. We assume that $\partial u_i / \partial a_i > 0$, $\frac{\partial^2 u_i(\cdot,\cdot)}{\partial a_i \partial \overline{a}} > 0$. Define $\alpha(\overline{a}) = \frac{1}{I}\sum_{i \in I} a_i^*(\overline{a})$. Any $\overline{a}$ such that $\alpha(\overline{a}) = \overline{a}$ is an equilibrium aggregate level of activity. Note that

(1) $\forall i \in I,\ (A_i, \geq)$ is a complete lattice,
(2) $\forall i \in I,\ u_i$ is bounded above, and every kind of continuous that you can think of because they are smooth,
(3) $\forall i \in I$ and each a_{-i}, $u_i(\cdot, a_{-i})$ is supermodular,
(4) $\forall i \in I$, $u_i(\cdot, \cdot)$ has increasing differences in a_i and a_{-i}.

This game may have many equilibria, and they are Pareto ranked.

Diamond search model (a quasi-theory of money): $A_i = [0, M]$ interpreted as effort spent trying to find someone to trade with, probability of success is proportional to own effort and effort of others according to $\tau a_i \cdot \sum_{j \neq i} a_j$ and utilities are $u_i(a) = \tau a_i \cdot \sum_{j \neq i} a_j - C(a_i)$. Note that

(1) $\forall i \in I,\ (A_i, \geq)$ is a complete lattice,

(2) $\forall i \in I$, u_i is bounded above, and every kind of continuous that you can think of because they are smooth,
(3) $\forall i \in I$ and each a_{-i}, $u_i(\cdot, a_{-i})$ is supermodular,
(4) $\forall i \in I$, $u_i(\cdot, \cdot)$ has increasing differences in a_i and a_{-i}, and
(5) $\forall i \in I$ and all a_{-i}, $u(\cdot, a_{-i})$ has increasing differences in a_i and τ.

If $C(\cdot)$ is increasing, $a = 0$ is always an equilibrium, if $C(\cdot)$ is smooth and convex, all of the solutions to $\tau(I-1)e = C'(e)$ are equilibria, they are Pareto ranked, and the upper equilibrium increases in with τ.

2. A bit more about lattices

Complete lattices, sublattices versus lattice subsets, upper semi-continuous functions, order upper semi-continuous functions. The following is due to Tarski.

THEOREM 5.1 (Tarski). *Every increasing map, R, from a complete lattice, $(X, \precsim)$, to itself has a largest fixed point, $x^*_{max} = \sup(\{x \in X : x \precsim R(x)\})$, and a least fixed point, $x^*_{min} = \inf(\{x \in X : R(x) \precsim x\})$.*

PROOF. Suppose that $(X, \precsim)$ is a complete lattice, and that R is an increasing map from X to itself. We will first show that $x^*_{max} = \sup(\{x \in X : x \precsim R(x)\})$ is a fixed point.

Define $U = \{x \in X : x \precsim R(x)\}$. Because X is complete, $\inf(X)$ exists. $U \neq \emptyset$ because $\inf(X) \in U$. Because $U \neq \emptyset$ and $(X, \precsim)$ is complete, $x^*_{max} = \sup U$ exists. To show that x^*_{max} is a fixed point, it is sufficient to show that $x^*_{max} \precsim R(x^*_{max})$ and $R(x^*_{max}) \precsim x^*_{max}$ (by antisymmetry).

By the definition of supremum, for every $u \in U$, $u \precsim x^*_{max}$. Because R is increasing, for all $u \in U$, $u \precsim R(u) \precsim R(x^*_{max})$ so that $R(x^*_{max})$ is an upper bound for U. Since x^*_{max} is the least upper bound, $x^*_{max} \precsim R(x^*_{max})$.

Since $x^*_{max} \precsim R(x^*_{max})$ and R is increasing, $R(x^*_{max}) \precsim R(R(x^*_{max}))$ so that $R(x^*_{max}) \in U$. Since x^*_{max} is the least upper bound for U, $R(x^*_{max}) \precsim x^*_{max}$.

We now show that x^*_{max} is the largest fixed point of R. Let x^* be any fixed point, i.e. satisfy $R(x^*) = x^*$. This implies that $x^* \precsim R(x^*)$, which in turn implies that $x^* \in U$. Therefore, $x^* \precsim x^*_{max}$.

The proof that x^*_{min} is the least fixed point proceeds along very similar lines, and is left as an exercise. □

PROBLEM 5.1. *Complete the proof of Tarski's fixed point theorem by showing that x^*_{min} is the least fixed point of T.*

3. A bit more about solution concepts for games

Nash equilibrium requires the satisfaction of some stringent conditions (though we've already seen cases where we want to strengthen them), correlated equilibrium asks for the satisfaction of less stringent conditions. There are also the **rationalizable** sets of strategies, and the serially undominated strategies.

3.1. Rationalizability. In the setting where one has beliefs β_s about ω, and maximizes $\int u(a,\omega)\,d\beta_s(\omega)$, an action $a \in A$ is **potentially Rational (**pR**)** if there exists some β_s such that $a \in a^*(\beta_s)$. An action a **dominates action** b if $\forall\omega\, u(a,\omega) > u(b,\omega)$. The following example shows that an action b can be dominated by a random choice.

EXAMPLE 5.2. *$\Omega = \{L, R\}$, $A = \{a, b, c\}$, and $u(a,\omega)$ is given in the table*

$A\downarrow, \Omega \rightarrow$	L	R
a	5	9
b	6	6
c	9	5

Whether or not a is better than c or vice versa depends on beliefs about ω, but $\frac{1}{2}\delta_a + \frac{1}{2}\delta_c$ dominates b. Indeed, for all $\alpha \in (\frac{1}{4}, \frac{3}{4})$, $\alpha\delta_a + (1-\alpha)\delta_c$ dominates b.

Let $R_i^0 = \text{pR}_i \subset A_i$ denote the set of potentially rational actions for i when they believe that others' actions have some distribution. Define $R^0 := \times_{i\in I} R_i^0$ so that $\Delta(R^0)$ is the largest possible set of outcomes that are at all consistent with rationality. (In Rational Pigs, this is the set $\delta_{\text{Wait}} \times \Delta(A_2)$.) As we argued above, it is too large a set. Now we'll start to whittle it down.

Define R_i^1 to be the set of maximizers for i when i's beliefs β_i have the property that $\beta_i(\times_{j\neq i} R_j^0) = 1$. Since R_i^1 is the set of maximizers against a smaller set of possible beliefs, $R_i^1 \subset R_i^0$. Define $R^1 = \times_{i\in I} R_i^1$, so that $\Delta(R^1)$ is a candidate for the set of outcomes consistent with rationality. (In Rational Pigs, you should figure out what this set is.)

Given R_i^n has been define, inductively, define R_i^{n+1} to be the set of maximizers for i when i's beliefs β_i have the property that $\beta_i(\times_{j\neq i} R_j^n) = 1$. Since R_i^n is the set of maximizers against a smaller set of possible beliefs, $R_i^{n+1} \subset R_i^n$. Define $R^{n+1} = \times_{i\in I} R_i^{n+1}$, so that $\Delta(R^n)$ is a candidate for the set of outcomes consistent with rationality.

LEMMA 5.3. *For finite games, $\exists N \forall n \geq N\; R^n = R^N$.*

We call $R^\infty := \bigcap_{n\in\mathbb{N}} R^n$ the set of **rationalizable strategies**. $\Delta(R^\infty)$ is then the set of **signal rationalizable outcomes**.[1]

There is (at least) one odd thing to note about $\Delta(R^\infty)$ — suppose the game has more than one player, player i can be optimizing given their beliefs about what player $j \neq i$ is doing, so long as the beliefs put mass 1 on R_j^∞. There is no assumption that this is actually what j is doing. In Rational Pigs, this was not an issue because R_j^∞ had only one point, and there is only one probability on a one point space. The next pair of games illustrate the problem.

3.2. Variants on iterated deletion of dominated sets. A strategy $\sigma_i \in \Delta_i$ **dominates (or strongly dominates)** $t_i \in A_i$ **relative to** $T \subset \Delta$ if

$$(\forall\sigma^\circ \in T)[u_i(\sigma^\circ\backslash\sigma_i) > u_i(\sigma^\circ\backslash t_i)].$$

[1] I say "signal rationalizable" advisedly. **Rationalizable outcomes** involve play of rationalizable strategies, just as above, but the randomization by the players is assumed to be stochastically independent.

If $T = \Delta$, this is the previous definition of dominance. Let $D_i(T)$ denote the set of $t_i \in A_i$ that are dominated relative to T. Smaller T's make the condition easier to satisfy.

In a similar fashion, a strategy $\sigma_i \in \Delta_i$ **weakly dominates** $t_i \in A_i$ **relative to** $T \subset \Delta$ if

$$(\forall \sigma^\circ \in T)[u_i(\sigma^\circ \backslash \sigma_i) \geq u_i(\sigma^\circ \backslash t_i)], \text{ and}$$
$$(\exists \sigma' \in T)[u_i(\sigma' \backslash \sigma_i) > u_i(\sigma' \backslash t_i)].$$

Let $WD_i(T)$ denote the set of $t_i \in A_i$ that are weakly dominated relative to T.

Lemma 5.4. *If Γ is finite, then for all $T \subset \Delta$, $A_i \setminus D_i(T) \neq \emptyset$ and $A_i \setminus WD_i(T) \neq \emptyset$.*

This is not true when Γ is infinite.

Problem 5.2. Two variants of 'pick the largest integer'.

(1) $\Gamma = (A_i, u_i)_{i \in I}$ *where* $I = \{1, 2\}$, $A_i = \mathbb{N}$, $u_i(n_i, n_j) = 1$ *if* $n_i > n_j$, *and* $u_i(n_i, n_j) = 0$ *otherwise. Every strategy is weakly dominated, and the game has no equilibrium.*

(2) $\Gamma = (A_i, v_i)_{i \in I}$ *where* $I = \{1, 2\}$, $A_i = \mathbb{N}$, *and* $v_i(n_i, n_j) = \Phi(n_i - n_j)$, $\Phi(\cdot)$ *being the cdf of a non-degenerate Gaussian distribution, every strategy is strongly dominated (hence the game has no equilibrium).*

Iteration sets $S_i^1 = A_i$, defines $\Delta^n = \times_{i \in I} \Delta(S_i^n)$, and if S^n has been defined, set $S_i^{n+1} = S_i^n \setminus D_i(\Delta^n)$. If Γ is finite, then Lemma 7.2 implies

$$(\exists N)(\forall n, n' \geq N)[S_i^n = S_i^{n'} \neq \emptyset].$$

There are many variations on this iterative-deletion-of-dominated-strategies theme. In all of them, $A_i^1 = \Delta_i$.

(1) **serially undominated** $S_i^{n+1} = S_i^n \setminus D_i(\Delta^n)$. If this reduces the strategy sets to singletons, then the game is **dominance solvable** (a term due to Herve Moulin).
(2) **serially weakly undominated** $S_i^{n+1} = S_i^n \setminus WD_i(\Delta^n)$ where $WD_i(T)$ is the set of strategies weakly dominated with respect to T.
(3) Set $S_i^2 = S_i^1 \setminus WD_i(\Delta^1)$, and for $n \geq 2$, set $S_i^{n+1} = S_i^n \setminus D_i(\Delta^n)$. [**7**], [**4**] show that the most that can be justified by appealing to common knowledge of the structure of the game and common knowledge of expected utility maximization is this kind of iterated deletion procedure.

3.3. The basic result. Every element of a rationalizable strategy set is serially undominated, and if $a \in A$ receives positive mass in any correlated equilibrium, then a is serially undominated.

4. Supermodular games

Definition 5.5. $\Gamma = ((A_i \succsim_i), u_i)_{i \in I})$ *is a* ***supermodular game***

(1) $\forall i \in I$, $(A_i, \succsim_i)$ *is a complete lattice,*
(2) $\forall i \in I$, u_i *is bounded above,* $u_i(\cdot, a_{-i})$ *is order upper semi-continuous for all* a_{-i}, $u_i(a_i, \cdot)$ *is order continuous for all* a_i,
(3) $\forall i \in I$ *and each* a_{-i}, $u_i(\cdot, a_{-i})$ *is supermodular,*
(4) $\forall i \in I$, $u_i(\cdot, \cdot)$ *has increasing differences in* a_i *and* a_{-i}.

There are three results on monotone properties of equilibria:

THEOREM 5.6 (Milgrom and Roberts (Theorem 5)). *If* Γ *is a supermodular game, then for each* $i \in I$*, there exist largest and smallest serially undominated strategies,* $\underline{a}_i$ *and* $\overline{a}_i$*. Further,* $\underline{a}$ *and* $\overline{a}$ *are pure strategy Nash equilibria.*

THEOREM 5.7 (Milgrom and Roberts (Theorem 6)). *If* $\Gamma = ((A_i \succsim_i), u_i(\cdot, \cdot, \tau))_{i \in I})$ *is a supermodular game for each* $\tau \in T$*,* T *a POSET, and each* $u_i(\cdot, a_{-i}, \cdot)$ *has increasing differences in* τ*, then the equilibria* $\underline{a}(\tau)$ *and* $\overline{a}(\tau)$ *are non-decreasing in* τ*.*

THEOREM 5.8 (Milgrom and Roberts (Theorem 7)). *If* Γ *is a supermodular game and* $b \succsim c$ *are two equilibria of* Γ *and* $\overline{a}$ *and* $\overline{a}$ *are the lowest and highest equilibria, then*

(1) if $\forall i \in I$*,* $u_i(\underline{a}, \cdot)$ *is increasing in* a_{-i}*, then* $\forall ii$*,* $u_i(b) \geq u_i(c)$*,*
(2) if $\forall i \in I$*,* $u_i(\underline{a}, \cdot)$ *is decreasing in* a_{-i}*, then* $\forall ii$*,* $u_i(b) \leq u_i(c)$*, and*
(3) if $\forall i \in I$*, either* $u_i(\underline{a}, \cdot)$ *is increasing in* a_{-i} *or* $u_i(\underline{a}, \cdot)$ *is decreasing in* a_{-i}*, then all of the players satisfying the increasing condition prefer* b *to* c*, and the remainder prefer* c *to* b*.*

CHAPTER 6

Positive and Negative Results for Game Dynamics

Readings: (Milgrom and Roberts [**12**]), (Conlisk [**6**]), (Hart and MasColell [**9**]).

Problem 6.1. *(Give Barnum his due.*[1]*) At a cost of* t*, people can develop a line of bullshit that makes someone with no self-discipline*[2] *give them \$1,000. At a cost of* a*,* $0 < a < t < 1,000$*, people can enforce a self-discipline that makes them immune to a trickster's line of bullshit. Call this the "avoider" strategy. At a cost of* 0*, a person is left vulnerable to tricksters, but loses nothing to avoiders. A pair of people meet and play one of the three strategies. The* 3×3 *symmetric game representing this is*

	Trickster	*Avoider*	*Sucker*
Trickster	$(-t, -t)$	$(-t, -a)$	$(1-t, -1)$
Avoider	$(-a, -t)$	$(-a, -a)$	$(-a, 0)$
Sucker	$(-1, 1-t)$	$(0, -a)$	$(0, 0)$

1. *Show that there is only one equilibrium and find it.*
2. *Let* Δ° *be the interior of the set of probability distributions on* $\{T, A, S\}$*. We understand* $x(t) \in \Delta^\circ$ *as the proportions of a population that are playing the strategies* T*,* A*, and* S*. Suppose that we have a continuous time dynamic for the proportion of a population playing this game with the property* $(*)$ *that* $\dot{x}_a > \dot{x}_b$ *whenever* $u(a, x) > u(b, x)$*. The interpretation is that those actions with higher payoffs tend to replicate faster than those with lower payoffs.*

   Show that such a system can have only one rest point and give it. Give conditions on a function f *with* $\dot{x} = f(x)$ *satisfying* $(*)$ *that guarantee that the system is globally stable.*

One version of the conclusion to be drawn from the previous problem is that the distribution of "rationality" can be an equilibrium result, that is, can be the result of thinking through costs and benefits. This ought to have some implications for cheap talk. In a bit more detail, once one has invested in developing a line of bullshit, the marginal cost of using it again might be thought of as being very small, that is, the talk might be cheap. However, it seems that in equilibrium, **no-one** should believe it, because, if they do, they lose \$1,000. However, if it is costly to learn to see through bullshit, this logic goes away. Modifying

[1]This problem is from to John Conlisk's "Costly Predation and the Distribution of Competence," *American Economic Review*, June 2001, 91(3), pp. 475-484.

[2]*Call such a person a "sucker" in deference to P. T. Barnum's dictum, "Every minute a sucker is born, and two to take him."*

cheap talk games to make this work is shockingly more complicated than it seems it ought to be, see Vince Crawford's "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions," *American Economic Review*, 133-149, March 2003.

CHAPTER 7

Equilibrium Refinement

Readings: extensive form games, Kuhn's Theorem, perfect and proper equilibria and stable sets of equilibria from the text, [**8**].

1. Iterative deletion procedures in extensive form games

We begin with an example, asking the ever-burning question, **"Will that be beer, or quiche for breakfast?"**

This game is due to Cho and Kreps (1987), who tell a version of the following story: There is a fellow who, on 9 out of every 10 days on average, rolls out of bed like Popeye on spinach. When he does this we call him "strong." When strong, this fellow likes nothing better than Beer for breakfast. On the other days he rolls out of bed like a graduate student recovering from a comprehensive exam. When he does this we call him "weak." When weak, this fellow likes nothing better than Quiche for breakfast. In the town where this schizoid personality lives, there is also a swaggering bully. Swaggering bullies are cowards who, as a result of deep childhood traumas, need to impress their peers by picking on others. Being a coward, he would rather pick on someone weak. He makes his decision about whether to pick, p, or not, n, after having observed what the schizoid had for breakfast. With payoffs, the game tree is

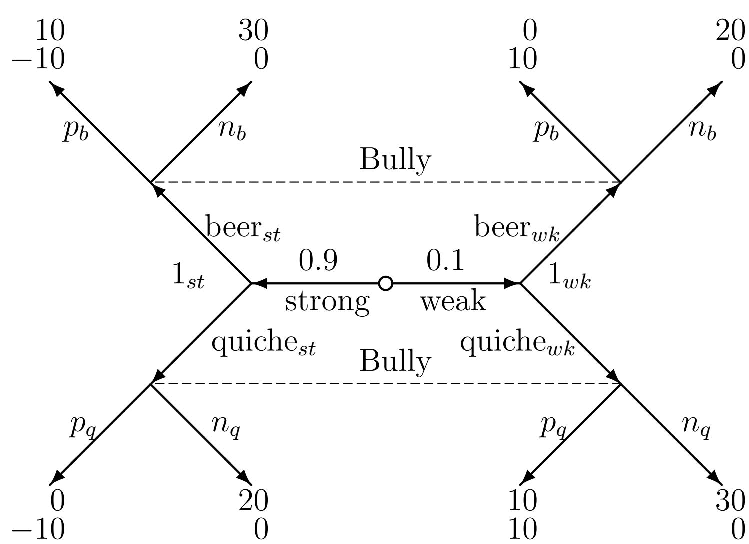

This is not only a psycho-drama, it's also a story about entry-deterrence. To get to that story, re-interpret the schizoid as an incumbent who may or may not have found a cost-saving technology change. If they have, it makes them a stronger competitor. This private information is not known to player 2, the potential entrant. The incumbent can start an aggressive advertising campaign, something they'd like better if they were strong, that is, if they have found the cost-saving change to their technology. The potential entrant would rather compete against a weak incumbent than a strong incumbent, but can condition their entry decision only on whether or not they see the aggressive advertising campaign.

1.1. Normal form analysis. We now give the 4×4 normal form for this game and find the two set of equilibria. Taking expectations over Nature's move, the 4×4 normal form for the Beer-Quiche game is

(st,wk)\(b,q)	(p, p)	(p, n)	(n, p)	(n, n)
(b, b)	$(9, -8)$	$(9, -8)$	$(29, 0)$	$(29, 0)$
(b, q)	$(10, -8)$	$(12, -9)$	$(28, 1)$	$(30, 0)$
(q, b)	$(0, -8)$	$(18, 1)$	$(2, -9)$	$(20, 0)$
(q, q)	$(1, -8)$	$(21, 0)$	$(1, -8)$	$(21, 0)$

The equilibrium set for this game can be partitioned into E_1 and E_2 where

$$E_1 = \{((q, q), (0, \beta, 0, 1 - \beta)) : 21 \geq 12\beta + 30(1 - \beta), \text{ i.e. } \beta \geq \frac{1}{2}\},$$

and

$$E_2 = \{((b, b), (0, 0, \beta, 1 - \beta)) : 29 \geq 28\beta + 30(1 - \beta) \text{ i.e. } \beta \geq \frac{1}{2}\}.$$

Import to note: E_1 and E_2 are connected sets of equilibria, and the outcome function, and hence payoffs, is constant on them. We will see this pattern in all the games that we look at. For some really silly non-generic games, we may not see this.

EXAMPLE 7.1. $I = \{1, 2, 3\}$, 1 *chooses which of the following two matrix games are played between* 2 *and* 3*, so* $A_1 = \{$*Left Box, Right Box*$\}$*,* $A_2 = \{$*Up, Down*$\}$*, and* $A_3 = \{$*Left, Right*$\}$*, and the payoffs are*

	Left	Right
Up	$(0, 0, 1)$	$(0, 0, 1)$
Down	$(0, 0, 2)$	$(0, 0, 2)$

	Left	Right
Up	$(0, 0, 3)$	$(0, 0, 3)$
Down	$(0, 0, 4)$	$(0, 0, 4)$

Notice that for all $\sigma \in \Delta$*, all* $i \in I$*, and all* $a_i \neq b_i \in A_i$*,* $U_i(\sigma \backslash a_i) = U_i(\sigma \backslash b_i)$*. Thus,* $Eq = \Delta$*, which is a nice closed connected set. However, the outcome function is* ***not*** *constant on this set, nor are the utilities, which are anyplace in the line segment* $[(0, 0, 1), (0, 0, 4)]$.

Returning to beer-quiche, there are dominance relations in this game, a mixed strategy dominates a pure strategy for 1, and after iterated elimination of dominated normal form strategies, only E_2 survives.

1.2. Behavioral strategies and an agent normal form analysis. Consider a mixed strategy $(\alpha, \beta, \gamma, \delta)$ for player 1. In the agent normal form, we take extraordinarily seriously the idea that every is the sum total of their experiences, and that different experiences make different people. This turns 1 into two people, having independent randomization at each information set. Give the Kuhn reduction to behavioral strategies, give Kuhn's Theorem.

There are no dominated strategies in the agent normal form of the game. However, there is something else, something that we will spend a great deal of time with.

1.3. Variants on iterated deletion of dominated sets. A strategy $\sigma_i \in \Delta_i$ **dominates (or strongly dominates)** $t_i \in A_i$ **relative to** $T \subset \Delta$ if

$$(\forall \sigma^\circ \in T)[u_i(\sigma^\circ \backslash \sigma_i) > u_i(\sigma^\circ \backslash t_i)].$$

If $T = \Delta$, this is the previous definition of dominance. Let $D_i(T)$ denote the set of $t_i \in A_i$ that are dominated relative to T. Smaller T's make the condition easier to satisfy.

In a similar fashion, a strategy $\sigma_i \in \Delta_i$ **weakly dominates** $t_i \in A_i$ **relative to** $T \subset \Delta$ if

$$(\forall \sigma^\circ \in T)[u_i(\sigma^\circ \backslash \sigma_i) \geq u_i(\sigma^\circ \backslash t_i)], \text{ and}$$
$$(\exists \sigma' \in T)[u_i(\sigma' \backslash \sigma_i) > u_i(\sigma' \backslash t_i)].$$

Let $WD_i(T)$ denote the set of $t_i \in A_i$ that are weakly dominated relative to T.

LEMMA 7.2. *If Γ is finite, then for all $T \subset \Delta$, $A_i \setminus D_i(T) \neq \emptyset$ and $A_i \setminus WD_i(T) \neq \emptyset$.*

This is not true when Γ is infinite.

PROBLEM 7.1. Two variants of 'pick the largest integer'.

(1) $\Gamma = (A_i, u_i)_{i \in I}$ *where* $I = \{1, 2\}$, $A_i = \mathbb{N}$, $u_i(n_i, n_j) = 1$ *if* $n_i > n_j$, *and* $u_i(n_i, n_j) = 0$ *otherwise. Every strategy is weakly dominated, and the game has no equilibrium.*

(2) $\Gamma = (A_i, v_i)_{i \in I}$ *where* $I = \{1, 2\}$, $A_i = \mathbb{N}$, *and* $v_i(n_i, n_j) = \Phi(n_i - n_j)$, $\Phi(\cdot)$ *being the cdf of a non-degenerate Gaussian distribution, every strategy is strongly dominated (hence the game has no equilibrium).*

Iteration sets $S_i^1 = A_i$, defines $\Delta^n = \times_{i \in I} \Delta(S_i^n)$, and if S^n has been defined, set $S_i^{n+1} = S_i^n \setminus D_i(\Delta^n)$. If Γ is finite, then Lemma 7.2 implies

$$(\exists N)(\forall n, n' \geq N)[S_i^n = S_i^{n'} \neq \emptyset].$$

There are many variations on this iterative-deletion-of-dominated-strategies theme. In all of them, $A_i^1 = \Delta_i$.

(1) **serially undominated** $S_i^{n+1} = S_i^n \setminus D_i(\Delta^n)$. If this reduces the strategy sets to singletons, then the game is **dominance solvable** (a term due to Herve Moulin).
(2) **serially weakly undominated** $S_i^{n+1} = S_i^n \setminus WD_i(\Delta^n)$ where $WD_i(T)$ is the set of strategies weakly dominated with respect to T.
(3) Set $S_i^2 = S_i^1 \setminus WD_i(\Delta^1)$, and for $n \geq 2$, set $S_i^{n+1} = S_i^n \setminus D_i(\Delta^n)$. [**7**], [**4**] show that the most that can be justified by appealing to common knowledge of the structure of the game and common knowledge of expected utility maximization is this kind of iterated deletion procedure.

1.4. Self-referential tests. These iterated procedures become really powerful when we make them self-referential. Let us ask if a set of equilibria, $E \subset Eq(\Gamma)$, is "sensible" or "internally consistent" are "stable" by asking if it passes an E-test. This kind of self-referential test is (sometimes) called an **equilibrium dominance test**. Verbally, this makes (some kind of) sense because, if everyone knows that only equilibria in a set E are possible, then everyone knows that no-one will play any strategy that is either weakly dominated or that is strongly dominated *relative to E itself*. That is, E should survive an E-test.

There is a problem with this idea, one that can be solved by restricting attention to a class $\mathcal{E}$ of subsets of $Eq(\Gamma)$. The class $\mathcal{E}$ is the class of closed and connected[1] subsets of $Eq(\Gamma)$.

Formally, fix a set $E \subset Eq(\Gamma)$, set $S_i^1 = A_i$, $E^1 = E$, given A_i^n for each $i \in I$, set $\Delta^n = \times_{i \in I} \Delta(S_i^n)$, and iteratively define S_i^{n+1} by

$$S_i^{n+1} = S_i^n \setminus \{WD_i(\Delta^n) \cup D_i(E^n)\}.$$

$E \in \mathcal{E}$ **passes the iterated equilibrium dominance test** if at each stage in the iterative process, there exists a non-empty $E^{n+1} \in \mathcal{E}$, $E^{n+1} \subset E^n$, such that for all $\sigma \in E^{n+1}$ and for all $i \in I$, $\sigma_i(\{WD_i(\Delta^n) \cup D_i(E^n)\}) = 0$. This means that something in E^n must be playable in the game with strategy sets S^{n+1}.

We will examine this workings of this logic first in a "horse" game, then return to beer-quiche, which belongs to a class of games known as signaling games.

1.5. A horse game. These games are called horse games because the game tree looks like a stick figure horse, not because they were inspired by stories about the Wild West.

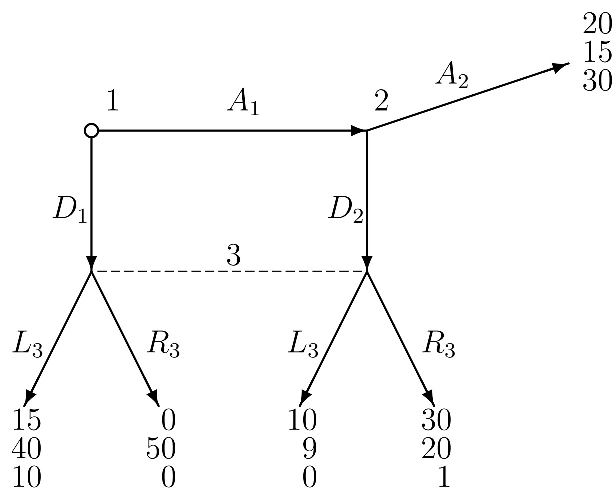

There are three sets of equilibria for this game, Listing 1's and 2's probabilities of playing D_1 and D_2 first, and listing 3's probability of playing L_3 first, the equilibrium set can be

[1] If you've had a reasonable amount of real analysis or topology, you will know what the terms "closed" and "connected" mean. We will talk about them in more detail later. Intuitively, you can draw a connected set (in our context) without taking your pencil off of the paper.

partitioned into $Eq(\Gamma) = E_A \cup E_B \cup E_C$,

$$E_A = \{((0,1),(0,1),(\gamma,1-\gamma)) : \gamma \geq 5/11\}$$

where the condition on γ comes from $15 \geq 9\gamma + 20(1-\gamma)$,

$$E_B = \{((1,0),(\beta,1-\beta),(1,0)) : \beta \geq \frac{1}{2}\}$$

where the condition on β comes from $15 \geq 10\beta + 20(1-\beta)$, and

$$E_C = \{((0,1),(1,0),(0,1))\}.$$

Note that $\mathbb{O}(\cdot)$ is constant on the sets E_A, E_B, and E_C. In particular, this means that for any $\sigma, \sigma' \in E_k$, $u(\sigma) = u(\sigma')$. I assert without proof that the E_k are closed connected sets.[2]

There are no weakly dominated strategies for this game:

(1) $u_1(s\backslash D_1) = (15,15,0,0)$ while $u_1(s\backslash A_1) = (10,20,30,20)$ so no weakly dominated strategies for 1,
(2) $u_2(s\backslash D_2) = (40,9,50,20)$ while $u_2(s\backslash A_2) = (40,15,50,15)$ so no weakly dominated strategies for 2,
(3) $u_3(s\backslash L_3) = (10,0,10,30)$ while $u_3(s\backslash R_3) = (0,1,0,3)$ so no weakly dominated strategies for 3.

Each E_k survives iterated deletion of weakly dominated strategies. However, E_A and E_B do not survive self-referential tests, while E_C does.

(1) E_A — the strategy D_1 is dominated for 1 **relative to** E_A. Removing D_1 makes L_3 weakly dominated for 3, but every $\sigma \in E_A$ puts mass on the deleted strategy, violating the iterative condition for self-referential tests. (We could go further, removing L_3 make A_2 dominated for 2, and every $\sigma \in E_A$ puts mass on A_2.)
(2) E_B — the strategy R_3 is dominated for 3 relative to E_B, removing R_3 make D_2 weakly dominated for 2, meaning that every $\sigma \in E_B$ puts mass on the deleted strategy, violating the iterative condition for self-referential tests.

The set E_C contains only one point, and it is easy to check that 1 point survives iterated deletion of strategies that are either weakly dominated or weakly dominated relative to E_C.

PROBLEM 7.2. *For the following horse game, partition $Eq(\Gamma)$ into closed and connected sets on which the outcome function, $\mathbb{O}(\cdot)$, is constant and find which of the elements of the partition survive the iterative condition for self-referential tests.*

[2]Intuitively, the sets are closed because they are defined by weak inequalities, and they are connected because, if you were to draw them, you could move between any pair of points in any of the E_k without lifting your pencil.

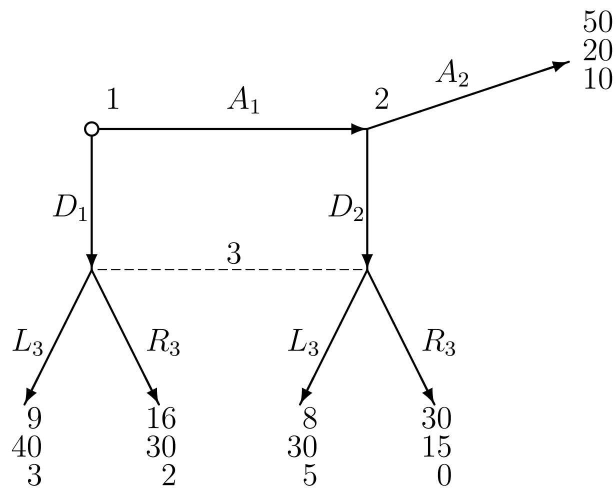

1.6. Back to beer and quiche. Return to beer-quiche, and use the self-referential tests, interpreting them in terms of reasonable beliefs.

2. Correspondences and fixed point theorems

We begin with reminders about what we need for our correspondences to have fixed points. For finite games, the best response correspondence is a closed, non-empty valued, convex valued, and upper-hemicontinuous, which, with a compact range, is equivalent to have a closed graph. Compare this with upper semi-continuous functions.

PROBLEM 7.3. *Show that the graph of $\sigma \mapsto Br(\sigma)$ is a closed subset of $\Delta \times \Delta$ by expressing it as the intersection of closed sets.*

Combining, the correspondence $\sigma \mapsto Br(\sigma)$ maps the non-empty, compact, convex Δ to itself, is non-empty valued, convex valued, and has a closed graph. These are the conditions for Kakutani's fixed point theorem.

3. Kakutani's fixed point theorem and equilibrium existence results

A **correspondence Ψ from X to Y** is a function from X to subsets of Y. The graph of Ψ is the set

$$gr_\Psi = \{(x, y) : y \in \Psi(x)\}.$$

Just as functions can be identified with their graphs, correspondences can be identified with their graphs.

A point x^* in X is a **fixed point of the correspondence Ψ from X to X** if $x^* \in \Psi(x^*)$. If we identify a function f from X to X with the correspondences $F(x) = \{x\}$, then x^* is a fixed point of f if and only if it is a fixed point of F. From before, we know that $\sigma^* \in Eq(\Gamma)$ if and only if $\sigma^* \in Br(\sigma^*)$, that is, if and only if σ^* is a fixed point of the best response correspondence.

Definition 7.3. *A correspondence Ψ from a non-empty, compact, convex subset, K, of $\mathbb{R}^n$ to $\mathbb{R}^n$ is a* **game theoretic correspondence (gtc)** *if*

(1) *for all $x \in K$, $\Psi(x) \neq \emptyset$,*
(2) *for all $x \in K$, $\Psi(x)$ is convex,*
(3) *gr_Ψ is a closed subset of $K \times K$.*

Theorem 7.4 (Kakutani). *Every gtc has a fixed point.*

Since $Br(\cdot)$ is a gtc, an immediate implication is

Theorem 7.5 (Nash). *Every finite game has an equilibrium.*

Another is

Theorem 7.6 (Brouwer). *If K is a compact convex subset of $\mathbb{R}^N$ and $f : K \to K$ is continuous, then there exists an x^* such that $f(x^*) = x^*$.*

None of the conditions in Kakutani's theorem can be omitted.

Problem 7.4. *Give a correspondence Ψ that has no fixed point and satisfies all of the assumptions of a gtc except:*

(1) *K is compact,*
(2) *K is convex,*
(3) *for all $x \in K$, $\Psi(x) \neq \emptyset$,*
(4) *for all $x \in K$, $\Psi(x)$ is convex,*
(5) *gr_Ψ is a closed subset of $K \times K$.*

Let $D \subset X \times X$ be the diagonal, i.e. $D = \{(x, x) : x \in X\}$. D is closed. If Ψ is a correspondence from X to X, then the fixed points of Ψ are the set of x^* such that $(x^*, x^*) \in (gr_\Psi \cap D)$. Because the intersection of closed sets is closed, the set of fixed points of a gtc is closed.

4. Perturbation based theories of equilibrium refinement

In §**??**, we removed equilibria by a variety of iterative procedures. In this section, we will perturb games in a variety of ways then take limits as the perturbations go to 0. In the various iterated deletion procedures, we noticed that there were some implications for beliefs that players might have had at different information sets. Perturbations also have implications for beliefs.

4.1. Overview of perturbations. An ϵ-perturbation of a game Γ is a game Γ_ϵ that we are willing to say is close to Γ. The ϵ may be a vector. The perturbation approach to refining $Eq(\Gamma)$ is to consider only those $\sigma \in Eq(\Gamma)$ that are of the form $\sigma = \lim_{\epsilon \to 0} \sigma^\epsilon$ where $\sigma_\epsilon \in Eq(\Gamma_\epsilon)$. This perturbation approach has three basic variants:

(1) (perfection) accept σ if there exists any Γ_ϵ, $\epsilon \to 0$, $\sigma^\epsilon \in Eq(\Gamma_\epsilon)$, $\sigma^\epsilon \to \sigma$,
(2) (properness) accept σ if there exists any <u>reasonable</u> Γ_ϵ, $\epsilon \to 0$, $\sigma^\epsilon \in Eq(\Gamma_\epsilon)$, $\sigma^\epsilon \to \sigma$,
(3) (stability) accept σ only if for all Γ_ϵ, $\epsilon \to 0$, there exists $\sigma^\epsilon \in Eq(\Gamma_\epsilon)$, $\sigma^\epsilon \to \sigma$.

Imagine a ball resting on a wavy surface. Any rest point is an equilibrium. The rest point is perfect if there exists some kind of tiny tap that we can give the ball and have it stay in the same area. The rest point is proper if it stays in the same area after reasonable kinds of taps (e.g. we might agree that taps from straight above are not reasonable). The rest point is stable if no matter what kind of tiny tap it receives, the ball stays in the same area after being tapped.

There are a variety of different ways to say that $d(\Gamma, \Gamma_\epsilon) < \epsilon$. One is to perturb the set of strategies, e.g. for $\epsilon_i \in \mathbb{R}^{A_i}_{++}$, let

$$\Delta^{\epsilon_i} = \{\sigma_i \in \Delta_i : (\forall a_i \in A_i)[\sigma_i(s_i) \geq \epsilon_i(s_i)]\}.$$

This notion of perturbation leads directly to Selten's perfection, with the appropriate notion of "reasonable" ϵ_i's, to Myerson's properness, and to Kohlberg and Merten's stability.

Another form of perturbation adds independent full support "noise," η, to each component of the utilities $u \in \mathbb{R}^{Z \cdot I}$, takes $P(\|\eta\| < \epsilon) > 1 - \epsilon$, tells i the realization of the $\eta_{i,z}$ but not the realization of the $\eta_{j,z}$, $j \neq i$, and has the players play the game Γ. This gives van Damme's $\mathcal{P}$-firm equilibria (perfection). Restricting to 'reasonable' η as the ones satisfying $E\|\eta\| < \epsilon$ gives Stinchcombe and Zauner's $\mathcal{P}$-good equilibria (properness). Asking for robustness to all reasonable η gives Stinchcombe and Zauner's $\mathcal{P}$-stable equilibria (stability).

Yet another form of perturbation identifies a game Γ with its best response correspondence, $Br(\cdot)$, and uses perturbations of the correspondence. This leads to Hillas stability, has not (to my knowledge) been used for a perfection or a properness concept. We will return to this in detail soon enough.

4.2. Perfection by Selten. Let $\epsilon = (\epsilon_i)_{i \in I}$ where $\epsilon_i \in \mathbb{R}^{A_i}_{++}$. Define Γ_ϵ to be the game $(\Delta^{\epsilon_i}, u_i)_{i \in I}$ where

$$\Delta^{\epsilon_i} = \{\sigma_i \in \Delta_i : (\forall a_i \in A_i)[\sigma_i(s_i) \geq \epsilon_i(s_i)]\}$$

and u_i is extended to Δ_i in the usual fashion.

Definition 7.7. *σ is* **perfect** *if there exists $\epsilon^n \to 0$, $\sigma^{\epsilon^n} \in Eq(\Gamma_{\epsilon^n})$ such that $\sigma^{\epsilon^n} \to \sigma$.*

Let $Per(\Gamma)$ denote the set of perfect equilibria. Kakutani's fixed point theorem implies that each $Eq(\Gamma_\epsilon) \neq \emptyset$, the compactness of Δ implies that any sequence σ^{ϵ^n} has a convergent subsequence, hence perfect equilibria exist.

The simplest example of the use of perfection is in the following 2×2 game where $Per(\Gamma) \subsetneq Eq(\Gamma)$,

	Left	Right
Up	$(1,1)$	$(0,0)$
Down	$(0,0)$	$(0,0)$

Fix a vector of perturbations (ϵ_1, ϵ_2). If 1 is playing a strategy σ_1 with $\sigma_1(\text{Up}) \geq \epsilon_1(\text{Up}) > 0$ and $\sigma_1(\text{Down}) \geq \epsilon_1(\text{Down}) > 0$, then 2 payoffs satisfy

$$u_2(\sigma_1, \text{Left}) \geq \epsilon_1(\text{Up}) > 0, \text{ and } u_2(\sigma_1, \text{Right}) = 0.$$

This means that in any perturbed game, Left strictly dominates Right for 2. Therefore, in any equilibrium of the perturbed game, 2 puts as much mass as possible on Left, that is, 2 plays the strategy $\sigma_2(\text{Left}) = 1 - \epsilon_2(\text{Right})$, $\sigma_2(\text{Right}) = \epsilon_2(\text{Right})$. By symmetry, 1 puts as much mass as possible on Up. Taking limits as $(\epsilon_1, \epsilon_2) \to 0$, the unique perfect equilibrium is (Up,Left).

A note here, since the $\epsilon_i \gg 0$, weak dominance turns into strict dominance in the perturbed games. Therefore, in perfect equilibria, weakly dominated strategies are not played.

Perfect equilibria do not survive deletion of strictly dominated strategies, an observation due to Roger Myerson.

	L	R	A_2
T	$(1,1)$	$(0,0)$	$(-1,-2)$
B	$(0,0)$	$(0,0)$	$(0,-2)$
A_2	$(-2,-1)$	$(-2,0)$	$(-2,-2)$

Delete the two strictly dominated strategies, A_1 and A_2 and you are back at the previous 2×2 game. In the present game, (B, R) is a perfect equilibrium. To see why, suppose that 2's perturbations satisfy $\epsilon_2(A_2) > \epsilon_2(L)$ and that 2 is playing the perturbed strategy $\sigma_2 = (\epsilon_2(L), 1 - (\epsilon_2(L) + \epsilon_2(A_2)), \epsilon_2(A_2))$. In this case, the payoff to T is strictly less than 0 while the payoff to B is 0. Therefore, if 1's perturbations satisfy the parallel pattern, there is only one equilibrium, (B, R) played with as much probability as possible. Taking limits as the perturbations go to 0, (B, R) is perfect. After deletion of the strongly dominated strategies, (B, R) is not perfect. Oooops.

Problem 7.5. *Show that all of the equilibria in the normal form of the Beer-Quiche game are perfect. Show the same in the agent normal form of the game.*

There is a bit more theory to be had.

Theorem 7.8 (Selten). *$Per(\Gamma)$ is a non-empty, closed subset of $Eq(\Gamma)$, and $\sigma^* \in Per(\Gamma)$ implies that for all $i \in I$, $\sigma_i^*(D_i) = 0$.*

Proof: (Non-emptiness) We have seen that $Per(\Gamma) \neq \emptyset$.

(Every perfect equilibrium is a Nash equilibrium) Suppose $\sigma \in Per(\Gamma)$, but assume, for the purposes of contradiction, that $\sigma \notin Eq(\Gamma)$. Since σ is perfect, there exists $\sigma^{\epsilon^n} \in Eq(\Gamma_{\epsilon^n})$, $\sigma^{\epsilon^n} \to \sigma$ as $\epsilon^n \to 0$. Since σ is not an equilibrium, there $(\exists i \in I)(\exists a_i \in A_i)[u_i(\sigma backslash s_i) > u_i(\sigma)]$. This implies that $Br_i^P(\sigma)$ contains points strictly better than σ for i. We will show that

$$(\ddagger) \quad (\exists N)(\forall n \geq N)[\sigma_i^{\epsilon^n}(Br_i^P(\sigma)) = 1 - \textstyle\sum_{t_i \notin Br_i^P(\sigma)} \epsilon_i^n(t_i)].$$

This implies that $\sigma_i^{\epsilon^n}(Br_i^P(\sigma)) \to 1$, which implies that $\sigma_i(Br_i^P(\sigma)) = 1$ which implies that σ is an equilibrium, a contradiction that completes the proof.

For any $s_i' \in Br_i^P(\sigma)$, continuity of the u_i implies that there exists an N such that for all $n \geq N$, $u_i(\sigma^{\epsilon^n} \backslash s_i') > u_i(\sigma^{\epsilon^n} \backslash t_i)$ for any $t_i \notin Br_i^P(\sigma)$. Since $\sigma^{\epsilon^n} \in Eq(\Gamma_{\epsilon^n})$, this delivers $(\ddagger)$.

(Perfect equilibria put no mass on weakly dominated strategies) Suppose $t_i \in A_i$ is weakly dominated. Then for all $\epsilon \gg 0$ and all $\sigma \in \times_{i \in I} \Delta^{\epsilon_i}$, $u_i(\sigma \backslash t_i) < u_i(Br_i(\sigma))$. Since $\sigma^{\epsilon^n} \in Eq(\Gamma_{\epsilon^n})$, this means that $\sigma_i^{\epsilon^n}(t_i) = \epsilon_i^n(t_i) \to 0$. Therefore $\sigma_i(t_i) = 0$.

(Closedness) For $\delta \in \mathbb{R}_{++}$, let E_δ denote the closure of the set of all $Eq(\Gamma_\epsilon)$ where $0 \ll \epsilon \ll \delta \cdot \tilde{1}$. (Note that $\delta < \delta'$ implies that $E_\delta \subset E'_\delta$.) It is easy to check that $Per(\Gamma) = \cap_\delta E_\delta$, that is, $Per(\Gamma)$ is the the intersection of closed sets, hence closed. ∎

The extensive form version of perfect is called "trembling hand perfection." To get it, just rephrase perfection using behavioral strategies. Let B_i denote i's set of behavioral strategies, that is, $B_i = \times_{H \in U_i} \Delta(A(H))$. For each $i \in I$, let $\epsilon_i = (\epsilon_H)_{H \in U_i}$ where $\epsilon_H \in \mathbb{R}^{A(H)}_{++}$. Define

$$B_i^{\epsilon_i} = \{b_i \in B_i : (forallH \in U_i)(\forall a_i \in A(H))[b_i(a_i) \geq \epsilon_H(a_i)]\}.$$

Note that if Γ is a normal form game, then $B_i^{\epsilon_i}$ is just Δ^{ϵ_i}.

For an extensive form game Γ, define Γ_ϵ to be the game $(B_i^{\epsilon_i}, u_i)_{i \in I}$ where u_i is composed with the outcome function $\mathbb{O}(\cdot)$.

Definition 7.9. *b is* **trembling hand perfect** *if there exists $\epsilon \to 0$, $b^\epsilon \in Eq(\Gamma_\epsilon)$ such that $b^\epsilon \to b$.*

Often times, trembling hand perfect equilibria are simply called perfect equilibria. Theorem 7.8 applies, except, and this is an important exception, no trembling hand perfect equilibrium puts positive mass on any strategy that is weakly dominated in the agent normal form of the game.

4.3. Properness by Myerson. A strategy $\sigma \gg 0$ is **ϵ-perfect**, $\epsilon \in \mathbb{R}_{++}$, if for all $i \in I$, $\sigma_i(Br_i^P(\sigma)) \geq 1 - \epsilon$. It is easy to show that σ is perfect if and only if it is the limit of ϵ^n-perfect equilibria, $\epsilon^n \to 0$. An ϵ perfect equilibrium σ is **ϵ-proper** if

$$(\forall i \in I)(\forall a_i, t_i \in A_i)\left[[u_i(\sigma\backslash a_i) < u_i(\sigma\backslash t_i)] \Rightarrow [\sigma_i(a_i) \leq \epsilon\sigma_i(t_i)]\right].$$

σ is **proper** if it is the limit of ϵ^n-proper equilibria as $\epsilon^n \to 0$.

Properness limits the set of perturbations that are allowed and then takes the set of limits. It has not limited the set of perturbations too far to preclude existence. Let $Pro(\Gamma)$ denote the set of proper equilibria.

Theorem 7.10 (Myerson). *$Pro(\Gamma)$ is a non-empty, closed subset of $Per(\Gamma)$.*

The proof is, to a first order of approximation, infinitely clever.

Proof: Pick $\epsilon \in (0,1)$. Set $\epsilon_i(a_i) = \epsilon^{\#A_i}/\#A_i$ for each $s_i \in A_i$. The perturbed strategy set $\Delta_i(\epsilon_i)$ is compact, convex and non-empty, as is the set $\times_{i \in I}\Delta_i(\epsilon_i)$. Define the correspondence $\Psi_i(\cdot)$ from $\times_{i \in I}\Delta_i(\epsilon_i)$ to $\times_{i \in I}\Delta_i(\epsilon_i)$ by

$$\Psi_i(\sigma) = \{\sigma_i \in \Delta_i(\epsilon_i) : [u_i(\sigma\backslash a_i) < u_i(\sigma\backslash t_i)] \Rightarrow [\sigma_i(a_i) \leq \sigma_i(t_i)]\}.$$

To see that $\Psi_i(\sigma) \neq \emptyset$, we introduce the ranking of a pure strategy $a_i \in A_i$ against σ, defining it by

$$r_i(\sigma, a_i) = \#\{t_i \in A_i : u_i(\sigma\backslash t_i) > u_i(\sigma\backslash a_i)\}.$$

Thus, is $a_i \in Br_i^P(\sigma)$, then $r_i(\sigma, a_i) = 0$. Consider the mixed strategy

$$\sigma_i(a_i) = \frac{\epsilon^{r_i(\sigma,a_i)}}{\sum_{t_i \in A_i} \epsilon^{r_i(\sigma,t_i)}}.$$

This is clearly a mixed strategy, and it belongs to $\Delta_i(\epsilon_i)$, so $\Psi_i(\sigma) \neq \emptyset$. It is also fairly easy to show that Ψ is a g.t.c., therefore it has a fixed point. The fixed point is an ϵ-proper equilibrium.

Take a sequence $\epsilon^n \downarrow 0$ and a corresponding sequence σ^n of ϵ^n-proper equilibria. Taking a subsequence if necessary, we can find a σ such that $\sigma^{n'} \to \sigma$ because Δ is compact.

To show that $Pro(\Gamma)$ is closed, let T_ϵ be the closure of the set of δ-proper equilibria, $\delta \in (0, \epsilon)$. It is nearly immediate that $Pro(\Gamma) = \cap_{\epsilon>0} T_\epsilon$, i.e. that the set of proper equilibria is the intersection of a collection of closed sets, hence it is itself closed. ∎

Problem 7.6 (Optional). *Show that the Ψ used in the previous proof is a gtc.*

Let us reconsider the previous 3×3 example,

	L	R	A_2
T	$(1,1)$	$(0,0)$	$(-1,-2)$
B	$(0,0)$	$(0,0)$	$(0,-2)$
A_2	$(-2,-1)$	$(-2,0)$	$(-2,-2)$

Note that if $\sigma \gg 0$, then for both $i \in I$ for this game, and all $t_i \neq A_i$, $u_i(\sigma \backslash A_i) < u_i(\sigma \backslash t_i)$. For an ϵ-proper equilibrium, we need to have $\sigma_i(A_i) \leq \epsilon \cdot \sigma_i(t_i)$ for both $t_i \neq A_i$. The only way to make B better than T in our calculations for (B, R) being perfect was to have $\epsilon_2(L) > \epsilon_2(A_2)$. This cannot happen in an ϵ-proper equilibrium. Therefore, there is only one proper equilibrium, (T, L). This survives deletion of strictly dominated strategies. There are 3 person games for which the set of proper equilibria changes after deletion of a strictly dominated strategy.[3]

4.4. Sequential equilibria. These are a subset of the trembling hand perfect equilibria. From that point of view, they are not exciting. What is really useful is that sequential equilibria are *defined* in terms of beliefs. This change unleashed a huge torrent of creative energy in the quest for a good refinement of $Eq(\Gamma)$. Let $b \gg 0$ be a strictly positive behavioral strategy. Bayes' Law gives beliefs, $\beta(\cdot|b)$ at each information set H.

Definition: A strategy-belief pair, (b, β) is a **sequential equilibrium** if there exists $b^n \gg 0$, $b^n \to b$ such that $\beta(\cdot|b^n) \to \beta$, and b is a best response to β at each H.

$Seq(\Gamma)$ denotes the set of sequential equilibria.

Definition: A strategy-belief pair, (b, β) is a **Bayesian Nash equilibrium** if $\beta(\cdot|b)$ is defined by Bayes Law at all H reached with positive probability when b is played, and b is a best response to β at each H.

Give some examples, emphasizing the beliefs at H that arise from Bayes Law and $b^n \gg 0$.

Every trembling hand perfect equilibrium is sequential, but not the reverse. Consider the following game

[3] Oh well, we need more work if we are to get S's that satisfy the self-referential tests above.

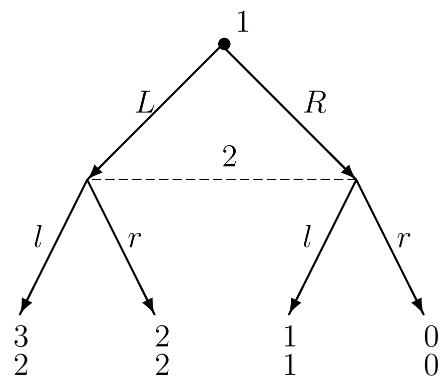

Problem 7.7. *Show that $Per(\Gamma) \subsetneq Seq(\Gamma) = Eq(\Gamma)$ in the above game.*

4.5. Strict perfection and stability by Kohlberg and Mertens. A strategy σ is **strictly perfect** if for all sequences Γ_{ϵ^n}, $\epsilon^n \to 0$, there exists a sequence $\sigma^n \in Eq(\Gamma_{\epsilon^n})$ such that $\sigma^n \to \sigma$.

Strictly perfect equilibria do not always exist, e.g.

	L	M	R
T	$(1,1)$	$(1,0)$	$(0,0)$
B	$(1,1)$	$(0,0)$	$(1,0)$

In this game, $Pro(\Gamma) = Per(\Gamma) = Eq(\Gamma)$, so none of the refinements work here. Perhaps they shouldn't, but in any case, it is easy to show that there is no strictly perfect equilibria, just take $\epsilon_2^n(M) > \epsilon_2^n(R) > 0$, $\epsilon_2^n(M) \to 0$. The only limit of equilibria of these perturbed games is (T, L). Now take $\epsilon_2^n(R) > \epsilon_2^n(M) > 0$, $\epsilon_2^n(R) \to 0$. The only limit of equilibria of these perturbed games is (B, L). Oooops, no strictly perfect equilibrium.

If x is a point and F is a set, define $d(x, F) = \inf\{d(x, f) : f \in F\}$.

Definition: A closed set $E \subset Eq(\Gamma)$ has property (SP) if it satisfies

(SP) for all sequences Γ_{ϵ^n}, $\epsilon^n \to 0$, there exists a sequence $\sigma^n \in Eq(\Gamma_{\epsilon^n})$ such that $d(\sigma^n, E) \to 0$.

Problem 7.8. *Show that if $E = \{\sigma\}$ and E satisfies* (SP)*, then σ is strictly perfect.*

Problem 7.9. *Show that for all Γ, $Eq(\Gamma)$ satisfies* (SP)*. (If this is too daunting, just show it for two non-trivial games.)*

In light of the last problem, what is needed is a smallest closed set E with property (SP).

Definition: A closed $E \subset Eq(\Gamma)$ is K-M stable if it satisfies Condition SP and no closed, non-empty, proper subset of E satisfies Condition SP.

In the 2×3 game above, we see that the unique K-M stable set is $\{(U, L), (R, L)\}$, i.e. it is not connected.

Time for the quick detour about connectedness. A set $E \subset \mathbb{R}^N$ is **path-connected** if for all $x, y \in E$, there is a continuous function $f : [0,1] \to E$ such that $f(0) = x$ and $f(1) = y$. Intuitively, this means that you can draw a continuous path from x to y that stays inside of E. It is clear that all convex sets are path-connected. It is also clear that if E and E' are path-connected and $E \cap E' \neq \emptyset$, then $E \cup E'$ is path-connected, though $E \cap E'$ may not be. There is a difference between connected and path-connected, and every path-connected set is connected. I am not going to tell you about the distinction because for our purposes path-connected is enough. However, this is a warning that when you see proofs of connectedness in the literature, you may not recognize the concept.

Look at Faruk Gul's example.

4.6. Stability by Hillas. Define the Hausdorff distance between two compact sets K_1 and K_2 by

$$d_H(K_1, K_2) = \inf\{\epsilon > 0 : K_1 \subset K_2^\epsilon, \text{ and } K_2 \subset K_1^\epsilon\},$$

where $K^\epsilon = \cup_{x \in K} B(x, \epsilon)$ is the open ϵ-ball around K.

Define the Hillas distance between two gtc's mapping K to K by

$$d(\Psi, \Psi') = \max_{x \in X} d_H(\Psi(x), \Psi'(x)).$$

Definition: A closed set $E \subset Eq(\Gamma)$ has property (S) if it satisfies

(S) for all sequences of gtc's Ψ^n, $d(\Psi^n, Br) \to 0$, there exists a sequence σ^n of fixed points of Ψ^n such that $d(\sigma^n, E) \to 0$.

Definition: A closed set $E \subset Eq(\Gamma)$ is (Hillas) stable if it satisfies (S) and no closed, non-empty, proper subset of E satisfies (S).

This is often said as "E is (Hillas) stable if it is minimal with respect to property (S)."

In the 2×3 game that we used to demonstrate that strictly perfect equilibria do not exist, it is (perhaps) possible to visualize the perturbations that Hillas uses. It is a good exercise to show that $((1/3, 2/3), (1, 0, 0))$ is the limit of fixed points of Ψ's close to Br.

THEOREM 7.11 (Hillas). *There exist a non-empty stable set E of equilibria. Further, E is connected, satisfies the self-referential tests given above, is a subset of $Per(\Gamma)$, and $Pro(\Gamma) \cap E \neq \emptyset$.*

In practice, the way that this is used is to (1) find the connected sets of equilibria and (2) test whether or not they fail the various criteria. For example, if a connected set E contains strategies that put positive mass on weakly dominated strategies, then $E \not\subset Per(\Gamma)$, and so E is not (Hillas) stable. If none of the proper equilibria belong to E, then E is not (Hillas) stable. If those tests don't work, and they often don't in signaling games, then you can turn to the self-referential tests. If you know that all but one of the connected sets of equilibria fails the criteria, then the existence part of the above theorem tells you that the remaining one must be (Hillas) stable. Directly checking stability can be difficult.

5. Signaling game exercises in refinement

Here are a variety of signaling games to practice with. The presentation of the games is a bit different than the extensive form games we gave above, part of your job is to draw extensive forms. Recall that a pooling equilibrium in a signaling game is an equilibrium in which all the different types send the same message, a separating equilibrium is one in which each types sends a different message (and can thereby be separated from each other), a hybrid equilibrium has aspects of both behaviors.

The presentation method is taken directly from Banks and Sobel's (1987) treatment of signaling games. Signaling games have two players, a Sender S and a Receiver R. The Sender has private information, summarized by his type, t, an element of a finite set T. There is a strictly positive probability distribution ρ on T; $\rho(t)$, which is common knowledge, is the ex ante probability that S's type is t. After S learns his type, he sends a message, m, to R; m is an element of a finite set M. In response to m, R selects an action, a, from a finite set $A(m)$; S and R have von Neumann-Morgenstern utility functions $u(t, m, a)$ and $v(t, m, a)$ respectively. Behavioral strategies are $q(m|t)$, the probability that S sends the message m given that his type is t, and $r(a|m)$, the probability that R uses the pure strategy a when message m is received. R's set of strategies after seeing m is the $\#A(m) - 1$ dimensional simplex Δ_m, and utilities are extended to $r \in \Delta_m$ in the usual fashion. For each distribution λ over T, the receiver's best response to seeing m with prior λ is

$$Br(\lambda, m) = \text{arg max}_{r(m)\in\Delta_m} \sum_{t\in T} v(t, m, r(m))\lambda(t). \tag{9}$$

Examples are represented with a bi-matrix $B(m)$ for each $m \in M$. There is one column in $B(m)$ for each strategy in $A(m)$ and one row for each type. The (t, a)'th entry in $B(m)$ is $(u(t, m, a), v(t, m, a))$. With t_1 being the strong type, t_2 the weak, m_1 being beer, m_2 being quiche, a_1 being *p*ick a fight, and a_2 being *n*ot, the Beer-Quiche game is

$B(m_1)$	a_1	a_2
t_1	$10, -10$	$30, 0$
t_2	$0, 10$	$20, 0$

$B(m_2)$	a_1	a_2
t_1	$0, -10$	$20, 0$
t_2	$10, 10$	$30, 0$

You should carefully match up the parts of this game and the extensive form of B-Q given above.

Here is a simple example to start on:

$B(m_1)$	a_1
t_1	$2, 2$
t_2	$2, 2$

$B(m_2)$	a_1	a_2
t_1	$3, 3$	$0, 0$
t_2	$0, 0$	$3, 3$

PROBLEM 7.10. *Draw the extensive form for the game just specified. Find the* 3 *connected sets of equilibria. Show that all equilibria for this game are both perfect and proper. Show that the* 3 *connected sets of equilibria are both Hillas and K-M stable.*

The following game is Cho's (1987, Example 2.1): the types are A, B, and C, $\rho(A) = \rho(C) = 3/8$, $\rho(B) = 1/4$, the messages are L and R, and the actions are as given.

$B(L)$	U	D
A	$2,1$	$-1,0$
B	$2,1$	$0,0$
C	$2,0$	$0,1$

$B(R)$	U	M	D
A	$0,2$	$0,0$	$0,2$
B	$0,2$	$4,3$	$1,-1$
C	$0,-3$	$1,-2$	$4,0$

Problem 7.11. *Draw the extensive form for the game just specified and analyze the equilibrium set.*

The following is a sequential settlement game of a type analyzed by Sobel (1989): There are two types of defendants, S: type t_2 defendants are negligent, type t_1 defendants are not, $\rho(t_1) = 1/2$. S offers a low settlement, m_1, or a high settlement, m_2. R, the plaintiff, either accepts, a_1, or rejects a_2. If R accepts, S pays R an amount that depends on the offer but not S's type. If R rejects the offer, S must pay court costs and a transfer depending only on whether or not S is negligent. With payoffs, the game is

$B(m_1)$	a_1	a_2
t_1	$-3,3$	$-6,0$
t_2	$-3,3$	$-11,5$

$B(m_2)$	a_1	a_2
t_1	$-5,5$	$-6,0$
t_2	$-5,5$	$-11,5$

Problem 7.12. *Draw the extensive form for the game just specified. Analyze the equilibria of the above game, picking out the perfect, the proper, the sequential, and the Hillas stable sets.*

One more game! This one has $\rho(t_1) = 0.4$.

$B(m_1)$	a_1
t_1	$0,0$
t_2	$0,0$

$B(m_2)$	a_1	a_2	a_3	a_4
t_1	$-1,3$	$-1,2$	$1,0$	$-1,-2$
t_2	$-1,-2$	$1,0$	$1,2$	$-2,3$

Problem 7.13. *Draw the extensive form for the game just specified. Find the pooling and the separating equilibria, if any, check the perfection and properness of any equilibria you find, and find the Hillas stable sets.*

Problem 7.14. *A little more practice with sequential equilibria.*

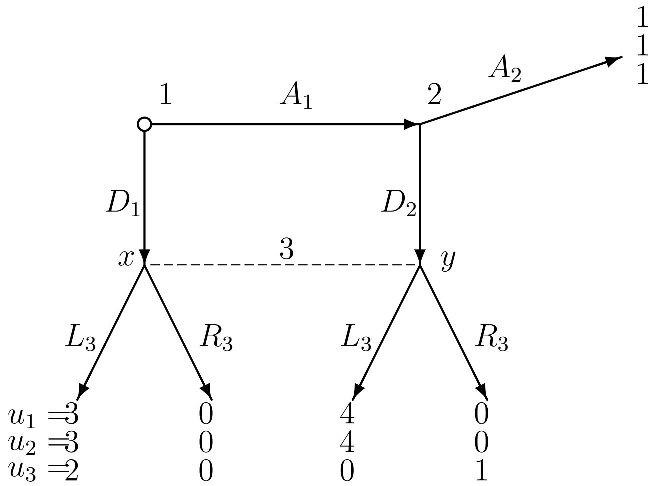

This game has two kinds of Nash equilibria, one kind involves player 1 *playing down, that is, playing* D_1*, the other involves player* 1 *playing across, that is, playing* A_1*.*

One of the equilibria in which 1 *plays* D_1 *is* (D_1, A_2, L_3)*, i.e. the behavioral strategy*

$$\pi = (\pi_1, \pi_2, \pi_3) = ((1, 0), (0, 1), (1, 0))$$

where the following table contains the key to reading these 1*'s and* 0*'s*

Player	*Strategy*
1	$\pi_1 = (\pi_1(D_1), \pi_1(A_1))$
2	$\pi_2 = (\pi_2(D_2), \pi_2(A_2))$
3	$\pi_3 = (\pi_3(L_3), \pi_3(R_3))$

(1) *Find the set of* π_2 *and* π_3 *for which* $((1, 0), \pi_2, \pi_3)$ *is a Nash equilibrium.*

(2) *Show that none of the* π_2 *you found in the previous problem is part of a sequential equilibrium.*

(3) *Find the set of Nash equilibria in which* 1 *plays* $\pi_1 = (0, 1)$*, that is, in which* 1 *plays* A_1*.*

(4) *Find the set of beliefs, that is,* (μ_x, μ_y)*,* $\mu_y = 1 - \mu_x$*, for* 3 *for which* L_3 *is a strict best response, for which* R_3 *is a strict best response, and for which* 3 *is indifferent between* L_3 *and* R_3*.*

(5) *Show that all of the Nash equilibria of Exercise 3 are sequential.*

Problem 7.15. *Yet more practice with sequential equilibria. Consider a market with one incumbent firm and a potential entrant. The potential entrant has low costs with probability* 0.2*, and has high costs with probability* 0.8*. The actual costs (low or high) of the entrant are private information to the entrant, who decides whether to stay "out" or "enter." The outside option for the low cost entrant has an expected utility of* 25*, while the outside option* *for the high cost entrant has an expected utility of* 0*. If the potential entrant stays out, then the incumbent has an expected utility of* 25*. If the potential entrant enters, then the incumbent decides whether to "fight" or "acquiesce." If the incumbent fights, then the payoffs for the entrant and incumbent respectively are* $(20, -25)$ *when the entrant is low cost, and* $(-10, 40)$ *when the entrant is high cost. If the incumbent acquiesces, then the payoffs are* $(40, 10)$ *when the entrant is low cost, and* $(20, 36)$ *when the entrant is high cost.*

One version of the game tree for this game is

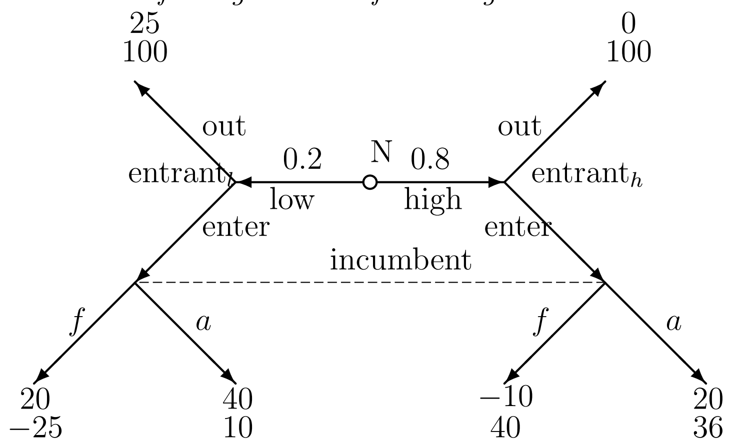

(1) *A* **separating equilibrium** *is one in which all the different types of Senders take different actions, thereby separating themselves from each other. A* **pooling equilibrium** *is one in which all the different types of Senders take the same action. Show that the only Nash equilibria of this game are pooling equilibria, and that all of them are sequential.*

(2) *Show that one of the sequential equilibria of the previous game is still a Nash equilibrium, but is not sequential in the following, changed version of the above game.*

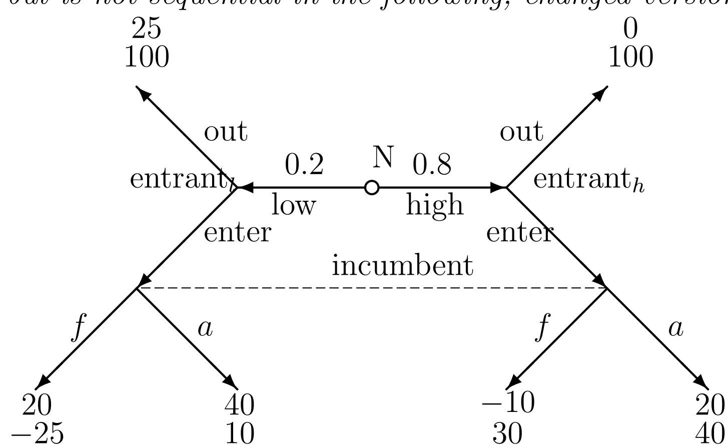

CHAPTER 8

Wandering in the Valley of the Mechanisms

I is the set of people involved. A is the set of possible outcomes. Each $i \in I$ has preferences over A.

DEFINITION 8.1. *A* ***mechanism*** *or* ***game form*** *is a pair* (M, g) *where* $\emptyset \neq M := \times_{i\in I} M_i$ *and* $g : M \to A$*. With preferences* $u := \{u_i \in \mathbb{R}^A : i \in I\} \in \mathcal{U}$*, a* ***solution*** $S((M, g), u)$ *is a subset of* M*. The* ***outcome correspondence*** *is* $\mathbb{O}_S((M, g), u) = g(S((M, g), u)) = \{a \in A : \exists m \in S((M, g), u), g(m) = a\}$*. A choice correspondence* $F : \mathcal{U} \to A$ *is* ***implemented by*** (M, g) *if* $\forall u \in \mathcal{U}$*,* $O_S((M, g), u) = F(u)$*.*

The technical name for a definition like that is *generalized abstract non-sense.* It only makes sense when we say something like, and bear in mind, this is only one possible example, "u has a distribution P, each $i \in I$ knows their own u_i and the distribution P, $S((M, g), u)$ is the set of Nash equilibria of the game $\Gamma = (M_i, v_i)_{i\in I}$ where $v_i(m) := u_i(g(m))$, and $F(u)$ is the set of Pareto optimal allocations when the utilities are u." Even then, it doesn't make much sense until we've used it in some examples.

1. Hurwicz 1972 through Maskin 1977

In honor of recent news from Stockholm, we'll start here.

$I = \{1, 2\}$, we have an Edgeworth box economy with endowments $e^1 = (0, 1)$, $e^2 = (1, 0)$. $a^i = (a^i_1, a^i_2)$ is i's allocation of goods 1 and 2, $A = \{(a^1, a^2) \in \mathbb{R}^4 : \sum_i a^i = \sum_i e^i\}$. There are two states of the world, α, β, both happening with some probabilities. In state 1, u_α is given by both people having standard Cobb-Douglas preferences. In state β, 2's preferences are the same, while 1's are $u_1(x_1, x_2) = x_2 - \frac{1}{1+x_1}$.

Let $F(u)$ be the Walrasian correspondence, i.e. $F(u_s) = \{$Walrasian eq'm allocations for the Edgeworth box economy $\mathcal{E} = (\mathbb{R}^2_+, u^i_s, e^i)_{i\in I}$. Check that $F(u_\alpha) = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$ and $F(u_\beta) = ((\frac{1}{2}, \frac{7}{9}), (\frac{1}{2}, \frac{2}{9}))$.

Suppose that M_i is i's possible set of utilities, so that M_1 has two points and M_2 has one point, that the mechanism takes reported vectors of utilities and gives the Walrasian eq'm, and that $S((M, g), u)$ is the set of Nash equilibria of this game. Player 1 will always misreport that they have the preferences u^1_b. This set up, where people report their own private information, is called a **direct mechanism**.

Now allow M_i to be any subset of the set of all utility functions and $g : M \to A$ to be arbitrary. There is still no Nash equilibrium that implements the Walrasian correspondence.[1]

[1]Jackson and Manelli (JET citation) go into approximate implementation in large economies.

Now suppose that 2 knows the state of the world and can report it. We can now design a mechanism that, essentially, checks on 1's veracity. Draw pictures. What is at work here is what Maskin (1977, 1999) called **monotonicity**. This is (yet) another place where the same word is used to mean very different things.

Definition 8.2. *The correspondence $F : \mathcal{U} \to A$ is **monotonic** if for any $u, u^\circ \in \mathcal{U}$, and any $a \in F(u)$ with $a \notin F(u^\circ)$, exists $i \in I$ and $b \in A$ such that $u_i(a) \geq u_i(b)$ and $u_i^\circ(b) > u_i^\circ(a)$.*

For allocations in $\mathbb{R}^2_+$, this involves either a crossing at a or a tangency at a with nested upper contour sets.

Theorem 8.3 (Maskin). *If F is Nash implementable, then it must be monotonic.*

In these terms, the Walrasian correspondence fails monotonicity, hence cannot be Nash implementable. And that was what was at work in Hurwicz's example. The idea of the proof is pretty simple: if an allocation is to be implemented at one profile of utilities but not at another, then in moving from the first to the second profile of utilities, the allocation must have fallen is someone's eyes in order to break the Nash-ness of the original allocation. Having preferences with indifference curves that cross only once (in 2 dimensional cases), called a single-crossing or a Spence-Mirrlees condition, tends to be an important and sensible kind of assumption.

2. Monopolist selling to different types

We now turn to the other main use of mechanism design, one in which someone has a utility function over the correspondences F and maximizes their preferences over the implementable correspondences.

Roughly 20% of the consumers of peanut butter buy 80% of the peanut butter that is sold. Roughly the same is true for hard liquor and wine. These are goods that are sold in larger and smaller quantities, and it is usually the case that the per unit price is cheaper for the larger quantities.[2] One explanation for this phenomenom is that the per unit packaging cost is smaller for the larger containers, after all, volume goes as the cube of dimension while packaging goes as the square. However, that is too prosaic an explanation for our present tastes and purposes, so we are going to look for another explanation.

Suppose that there are two types of consumers of peanut butter, good x, and that they are described by the two utility functions,

$$u_l(x, w) = r_l v(x) + w, \quad u_h(x, w) = r_h v(x) + w,$$

where $0 < r_l < r_h$, $v(\cdot)$ is a differentiable, strictly concave, strictly increasing function satisfying $v(0) = 0$ (a harmless normalization), $\lim_{x\to\infty} v'(x) = 0$ (the consumer eventually grows tired of peanut butter), $\lim_{x\downarrow 0} r_l v'(x) > c$ where c is the marginal cost of producing peanut butter (the type l consumer's marginal utility of peanut butter is larger than the marginal cost of peanut butter if they have nearly none of it), and w represents wealth to spend on

[2]This is much less true at my local grocery store than it was at the local grocery stores of my youth.

other items in the consumer's life. In the x-w plane, the slopes of the indifference curves are

$$\frac{dw}{dx} = -\frac{\frac{\partial u_l}{\partial x}}{\frac{\partial u_l}{\partial w}} = -\frac{r_l v'(x)}{1} = -r_l v'(x), \quad \frac{dw}{dx} = -\frac{\frac{\partial u_h}{\partial x}}{\frac{\partial u_h}{\partial w}} = -\frac{r_h v'(x)}{1} = -r_h v'(x),$$

so that the indifference curve of the h type through any point (x, w) is strictly steeper than the indifference curve of the l type through the same point. This is the famous **single-crossing property** — the indifference curves of the different types cross at most a single time.

Before advancing to the fancy explanation for two sizes, let us suppose that you were a monopolist selling to one of these types, that you knew their type, t, and that your marginal cost of production is c. You are going to choose a size, q, to sell and a \$ price p to charge the consumer. The profit maximization problem is

$$\max_{p,q} p - cq \text{ subject to } u_t(q, w-p) \geq u(0, w^\circ), \quad t = l, h$$

where w° is the initial wealth of the consumer. Notice that revenues are p, not p times q here, p is a dollar price, not a dollar per unit price. The constraint comes from the observation that consumer can always not buy peanut butter (just as they could always opt out of paying for the public good above). The Lagrangean for this problem is

$$\mathcal{L}(p, q; \lambda) = p - cq - \lambda[w^\circ - (r_t v(q) + w^\circ - p)] = p - cq + \lambda[r_t v(q) - p] \quad t = l, h.$$

Since $\lim_{x \downarrow 0} r_l v'(x) > c$, the solution is interior, and it is clear that the constraint is binding, therefore the Kuhn-Tucker conditions reduce to

$$\frac{\partial \mathcal{L}}{\partial p} = 1 - \lambda = 0, \quad \frac{\partial \mathcal{L}}{\partial q} = -c + \lambda r_t v'(q) = 0, \quad \frac{\partial \mathcal{L}}{\partial \lambda} = r_t v(q) - p = 0.$$

The first two equations say that marginal utility is equal to marginal cost. This means that it is Pareto efficient for the monopolist to be able to tell the consumers apart. This is no surprise, this is the famous case of the perfectly discriminating monopolist. The third equation says that the solution lies on the indifference curve of type t through their initial holdings, w°. It is easy to show that at solutions (p_l, q_l) and (p_h, q_h) to this pair of problems, $q_l < q_h$, and that there are $v(\cdot)$ such that $p_l/q_l > p_h/q_h$, that is, it is always the case that the l types are given smaller quantities than the h types, and it is sometimes the case that they pay a higher per unit cost than the h types.

PROBLEM 8.1. *Show it is always the case that the l types are given smaller quantities than the h types, and it is sometimes the case that they pay a higher per unit cost than the h types.*

We are going to assume that $p_l/q_l > p_h/q_h$ (because I like this case better, but it really doesn't matter), and examine what happens when the monopolist cannot distinguish the types of consumers from each other. The starting point is

PROBLEM 8.2. *Show that $u_h(p_l, q_l) > u_h(p_h, q_h)$.*

In words, the h type consumers would prefer to consume the l type's q_l at \$ price p_l. Now, the monopolist extracts all the surplus from the consumers when s/he can distinguish between them. However, if the h types can lie about their type and get the (p_l, q_l) combination, they will. The monopolist might try all kinds of complicated mechanisms to get the consumers to reveal their types so as to extract more from them, but by the revelation principle, any and all such mechanisms must be reduceable to maximizing profits subject to IR and IC constraints. Letting N_l and N_h represent the numbers of l and h types respectively, the monopolist's problem is

$$\max_{(p_l,q_l),(p_h,q_h)} N_l(p_l - cq_l) + N_h(p_h - cq_h) \text{ subject to}$$
$$u_l(p_l, q_l) \geq w^\circ,\ u_h(p_h, q_h) \geq w^\circ,$$
$$u_l(p_l, q_l) \geq u_l(p_h, q_h),\ u_h(p_h, q_h) \geq u_h(p_l, q_l).$$

The first line of constraints are the IR constraints, the second line are the IC constraints.

PROBLEM 8.3. *Letting $\alpha = N_l/(N_l + N_h)$, find how the solutions to the above problem vary as a function of $\alpha \in [0, 1]$.*

3. Hiring People to Run Your Business

Suppose somebody, call them a principal, wants to hire someone else, call them an agent, to do a particular job. A possible complication is that the principal cannot observe, at least not at any reasonable cost in time or money, exactly how the agent does the job. Another possible complication is that the principal doesn't know everything about the job that the agent knows. However, the principal does observe the final outcome. The principal wants to design a reward schedule for the agent, a mechanism if you will, that depends on the observable final outcome, and that induces the agent to behave as the principal desires.

For example, the owner of a factory (the principal) may hire a manager (the agent) to run the factory while the principal runs off to the South Pacific. It is too costly for the principal to watch what the agent does from Tahiti. However, each quarter the principal's accountant calls and tells the principal the realized profits. Now, profits are presumably related to the agent's actions, but they also have a random component. This is a peculiar market, the principal is buying the services sold by someone else, but the exact services are not specified in the contract. Here the mechanism is a wage schedule depending on the observable profit.

For another example, the principal sees a doctor or a lawyer (agent) when something is wrong. Often, the extent and treatment for the problem are not known to the principal and only probabilistically known to the agent. The principal eventually observes the outcome, that is, when/if the problem goes away. This too is a peculiar market, the principal is asking the agent what is needed and then buying the needed services from the agent. In the doctor example, the wage schedule should depend on the health outcome.

For yet another example, people differ in their real willingness and ability to pay for public goods. Public goods should be offered if the sum, over the affected people, of the willingness to pay is greater than or equal to the cost of the public good. However, if actual payment depends on reported willingness to pay, people have a motivation to "free ride," that is, to say that the are less willing to pay than they truly are in order to receive the public good at less cost to themselves. Here, one invents a rule or mechanism for turning reported willingnesses to pay into a decision about the public good. Given the mechanism/rule, the people best respond. In particular, they don't play weakly dominated strategies. Here, the principal is an abstraction sometimes thought of as society, there are many agents, and society, in the guise of the economist writing down the model, wants to get a particular kind of outcome. Some specific models are in order.

3.1. Hiring a manager. The principal offers a contract to the agent. The contract specifies what reward or wage the agent will get as a function of the observable outcome, say profit — S_1 is the set of functions $\pi \mapsto w(\pi)$. The agent examines the contract, $w(\cdot)$, and either decides to decline the job offer, or accepts the offer and chooses a managerial behavior, $b \in B$ — S_2 is the set of functions from S_1 to $\{\text{decline}\} \cup B$. If a point in B is chosen, then a random profit, Π_b, is realized. The expected utilities when the job offer is not declined are $E\,u_1(\Pi_b - r(\Pi_b))$ and $E\,u_2(r(\Pi_b), b)$. The principal's utility depends only on the difference between the realized profit and the wage they pay, the agents utility depends only on the wage they receive and their choice of managerial behaver. If the job offer is declined, the expected utilities are $\underline{u}_1$ and $\underline{u}_2$.

Just as in Atomic Handgrenades, the manager could bluff, saying that s/he won't take the job unless offered a very large salary schedule, one that guarantees more (in expected utility) than $\underline{u}_2$. However, when offered something more reasonable, by which I mean some lower wage schedule offering an expected utility greater than $\underline{u}_2$, declining is a weakly dominated strategy. Thus, in looking at an equilibrium for this game, we can proceed in two stages: first, figure out 2's best responses to different $w(\cdot)$'s; second, figure out the highest $E\,u_1$ as a function of $w(\cdot)$ coupled with best responses by 2.

We are going to start with the simplest case: $B = \{b_l, b_h\}$ where l and h stand for low and high respectively, the possible profit levels are $\{\pi_l, \pi_h\}$, $\pi_l < \pi_h$, the probability of π_h conditional on action b_l being taken is $P(\pi_h|b_l) = \alpha$, the probability of π_h conditional on action b_h being taken is $P(\pi_h|b_h) = \beta$, and $0 < \alpha < \beta < 1$. Also, suppose that u_2 is of the form $u(w) - v(b)$ where u is concave in wages, where $v(b_l) < v(b_h)$, and suppose that u_1 is linear in profits minus wage payments.

Problem 8.4. *For some values of the parameters of this game, the equilibrium involves the manager choosing b_l. Charaterize the equilibrium $w(\cdot)$ for such equilibria. For other values of the parameters of this game, the equilibrium involves the manager choosing b_h. Show that the corresponding equilibrium $w(\cdot)$ involves performance rewards, that is, $w(\pi_l) < w(\pi_h)$.* [You can use graphs to answer this question.]

Problem 8.5. *Change the set of profit levels to $\{\pi_1 < \pi_2 < \cdots < \pi_M\}$, have $P(\pi_m|b_l) = \alpha_m$, $P(\pi_m|b_h) = \beta_m$ where $\sum_{m=1}^M \alpha_m \pi_m < \sum_{m=1}^M \beta_m \pi_m$, but otherwise leave the game unchanged.*

1. *For some values of the parameters of this game, the equilibrium involves the manager choosing b_l. Charaterize the equilibrium $w(\cdot)$ for such equilibria.*
2. *For other values of the parameters of this game, the equilibrium involves the manager choosing b_h. For such an equilibrium:*
   - (a) *Give the Kuhn-Tucker optimality conditions that must be satisfied by the equilibrium $w(\pi_m)$, $m = 1, \ldots, m$.*
   - (b) *Show that the equilibrium $w(\cdot)$ may not involve performance rewards, that is, $w(\pi_m) > w(\pi_{m+1})$ is possible.*
   - (c) *Show that if α_m/β_m decreases in m, the equilibrium $w(\cdot)$ involves performance rewards.*

(3) Add an extra stage to this game — after observing π_m, the manager can report any $\pi_{m-k} \leq \pi_m$ (say by overpaying a supplier), and wage depends on the reported π. Show that in any equilibrium involving the manager picking b_h, the equilibrium $w(\cdot)$ involves weak performance rewards, that is, $w(\pi_{m-1}) \leq w(\pi_m)$.

4. Funding a public good

A group of neighborhood families, $i \in \{1, \ldots, I\}$, must decide whether or not to pool their money to hire a morning crossing guard for a busy intersection that the children walking to school must cross. The cost of a guard is C, which must be financed by family contributions, t_i. The set of feasible alternatives can be modeled as

$$X = \{x = (y, m_1 - t_1, \ldots, m_I - t_I) : y \in \{0,1\},\ t_i \geq 0,\ \sum_i t_i \geq yC\},$$

where m_i is family i's initial endowment of money. Family i's preferences are given by

$$u_i(x, \theta_i) = \theta_i y + (m_i - t_i),$$

where $\theta_i \in \Theta_i = \{\theta_L, \theta_H\}$ where $0 < \theta_L < \theta_H$. The θ_i are independent, and have probability $\frac{1}{2}$ of being either θ_H or θ_L. The set of possible vectors of preferences is Θ with typical element $\theta = (\theta_1, \ldots, \theta_I) \in \Theta = \times_{i \in I} \Theta_i$. We are going to assume that each family knows their own θ_i but not the θ_j of any $j \neq i$. (Draw a tree here).

A social choice rule is a function f mapping Θ to X. Family i's maximal willingness to pay, p_i, is the number that solves

$$\theta_i + (m_i - p_i) = m_i$$

because paying p_i and having the public good (crossing guard) leaves them just as well off as they are without the public good. In particular, $p_i = \theta_i$. A social choice rule is ex post efficient if for all θ such that $\sum_i \theta_i = \sum_i p_i > C$, $f(\theta)$ has $y = 1$, and if $\sum_i \theta_i = \sum_i p_i < C$, $f(\theta)$ has $y = 0$.

The general form of a political decision process involves each family i picking an m_i in some set M_i ("M" for Message). Since the M_i at this point are rather abstract, this seems to cover all of the relevant territory — M_i might include canvassing the neighbors, putting up placards, calling city hall, sacrificing chickens to the traffic gods, whatever. However, we think that contributions to the crossing guard fund are voluntary, so the M_i should include an m_i^0 corresponding to "i will make no contribution," and if m_i^0 is played, then family i makes no contribution. A strategy for family i a point in $A_i = M_i^{\{\theta_L, \theta_H\}}$. A vector s of strategies gives a mapping from Θ to $\times_{i \in I} M_i$. A mechanism is a mapping, $\mathcal{M}$, from $\times_{i \in I} M_i$ to allocations. Putting these together, each strategy gives a mapping from Θ to allocations. Since allocations have associated utilities, $s \mapsto E\, u_i(\mathcal{M}(s))$, we have a game. Since a given s leads to a distribution over the terminal nodes and it is $E\, u$ that is maximized, the equilibria for these games are sometimes called **Bayesian Nash** equilibria.

Since no-one can be forced to pay, the mechanism must have the property that for any m with $m_i = m_i^0$, then i's part of the allocation in $\mathcal{M}(m)$ involves i not paying. One interesting kind of question is whether or not there exist any mechanisms with the property that the equilibrium of the corresponding game is ex post efficient. At first glance, this seems impossibly difficult to answer, after all, there are no restrictions on the M_i nor the $\mathcal{M}$ besides the existence of a "no pay" option. However, there is a result, called the **Revelation Principle**, that allows us to study this question with some hope of success.

Some terminology: a social choice rule $f(\cdot)$ is **implemented by a mechanism** $\mathcal{M}$ if there is an equilibrium s^* such that $\forall \theta \in \Theta$, $\mathcal{M}(s^*(\theta)) = f(\theta)$. We are trying to answer the question "Which social choice rules are implementable?"

Intuition: when a family of type θ_i plays in a pure strategy equilibrium, s^*, they play $m_i = s_i^*(\theta_i)$. We can, if we wish, interpret any given $s_i^*(\theta_i)$ as if the family had said θ_i. Equivalently, we can imagine a mechanism that asks the family for their θ_i and guarantees that if told θ_i, $s_i^*(\theta_i)$ will be played. Since the mechanism delivers the same allocation, they family is just as happy revealing their true willingness to pay.

It's important not to get too excited here — everyone truthfully revealing their types is possible, but only if it is part of an equilibrium. It is still the case that if revealing their type hurts them, they won't do it. In particular, everyone must have the option of not paying.

Formally, replace each M_i with the simpler two point space, $M_i' = \{\theta_L, \theta_H\}$, and define a simpler mechanism, $\mathcal{M}'(\theta) = \mathcal{M}(s^*(\theta))$. Let $f(\cdot)$ be the mapping from Θ to allocations induced by s^* and $\mathcal{M}$, equivalently, by truth-telling and $\mathcal{M}'$, and let us examine the implications for $f(\cdot)$ of these strategies being an equilibrium.

First, it must be that case that $\forall i \in I$, $E\, u_i(f) \geq w_i$. If this is not true, then s^* is not an equilibrium because the family has the option of playing the strategy $s_i' \equiv m_i^0$, which delivers an expected utility greater than or equal to w_i. These conditions, one for each $i \in I$, are called **individual rationality (IR)** constraints. Second, after seeing any θ_i, the expected utility to reporting θ_i to $\mathcal{M}'$ must be at least as large as reporting any $\theta_i' \neq \theta_i$. If this is not true, then s^* is not an equilibrium, after all, one of the options with the original mechanism is to play $s_i'(\theta_i) = s^*(\theta_i')$, but it is an equilibrium to play $s_i^*(\theta_i)$. In terms of $f(\cdot)$, it must be the case that for all $i \in I$, $E\, u_i(\mathcal{M}'(\theta \backslash \theta_i)) \geq E\, u_i(\mathcal{M}'(\theta \backslash \theta_i'))$ for all θ_i'. These conditions are called **incentive compatibility (IC)** constraints.

The IR and IC constraints must be satisfied if f is implementable. If f satisfies the constraints, then it is implementable by the truth telling strategies in the simpler game described above. Therefore

THEOREM 8.4. *The social choice rule $f(\cdot)$ is implementable if and only if it is feasible and satisfies the* IR *and* IC *constraints.*

Now the real work starts, what's implementable?

PROBLEM 8.6. *Suppose that $\sum_i \theta_H > C > \sum_i \theta_L$ and $f(\cdot)$ is the unanimity rule, "hire the guard and spread the cost evenly between the families only when all report θ_H." Show that $f(\cdot)$ is implementable but not generally efficient.*

PROBLEM 8.7. *Suppose that f is the rule, "hire the guard and spread the cost evenly among those reporting θ_H only when the sum of the reported θ_H's is greater than C." Is this implementable? Efficient?*

4.1. Detour Through Vickrey Auctions. Second price sealed bid auction of a single object: I is the set of bidders, $A_i = \mathbb{R}_+$ are the possible bids, $u_i = v_i - p$ if i has the good,

$u_i = 0 - p$ if they do not have it. Assume that $v_i \geq 0$ with $(v_i)_{i \in I} \in \mathbb{R}^I_+$ distributed Q, and that i knows v_i before deciding what to bid.

Vickrey auction: The good is awarded to the person with the highest bid who pays the second highest price. Bidding v_i is a dominant strategy.

This is not true if: the game is for more than one object; or another auction will happen later; or if the seller will use your bids to gain information about you that will later allow them to extract more surplus from you; or . . . .

The ascending bid, Dutch bulb auctions, have the same strategic structure.

4.2. The Vickrey-Clark-Groves Mechanism for Public Goods. Returning to the public good provision question above, let us suppose that every individual pays:

(1) some function $t_i(\theta_{-i})$ of everybody elses' reported willingness to pay if their report does not make a difference between whether or not the public good is provided;
(2) the difference between the total cost of the good and the sum of everyone else's reported valuations if the person's report is **pivotal**, that is, if $\sum_{j \neq i} \theta_j < C$ and $\theta_i + \sum_{j \neq i} \theta_j \geq C$; and
(3) nothing if the sum of willingnesses to pay is less than or equal to C.

Truthfully reporting your value is a dominant strategy. Depending on the $t_i(\cdot)$ functions, this mechanism will generate budget surpluses or deficits. If people start acting with getting some share of that surplus or paying some share of that deficit, you change the incentives. Here's another way to write down the rule just described.

PROBLEM 8.8. *Give conditions on the $t_i(\cdot)$ so that the following pivotal voter rule satisfies the* IR *and* IC *constraints.*

$$f_i(\theta_1, \ldots, \theta_I) = \begin{cases} (1, w_i - t_i(\theta_{-i})) & \text{if } \sum_{j \neq i} \theta_j \geq C, \\ (1, w_i - (C - \sum_{j \neq i} \theta_j)) & \text{if } C - \theta_i \leq \sum_{j \neq i} \theta_j < C, \\ (0, w_i) & \text{otherwise.} \end{cases}$$

Let $y^*(\theta_1, \ldots, \theta_I) = 1$ if $\sum_k \theta_k \geq C$, 0 otherwise. Let $\tilde{\theta}_{-i}$ denote the random realization of the type of families other than i. Define a class of transfer functions, mechanisms if you will, depending on reported θ's, by

$$t_i(\theta) = -E \sum_{j \neq i} \tilde{\theta}_j \cdot y^*(\theta_i, \tilde{\theta}_{-i}) + h_i(\theta_{-i}).$$

So far as equilibrium calculations are concerned, the $h_i(\theta_{-i})$ cannot affect anything, this from Lemma 3.1. When each i maximizes the first term, they are maximizing a term that we can think of as containing the sum of their externalities on other players, assuming that the other players truthfully reveal their own θ_j. When all players do this, the mechanism satisfies the IC constraints, by fooling with the $h_i(\cdot)$, it is also possible to take care of feasibility and budget balancing. This is especially easy if the number of people is large and you only ask for budget balancing on average.

5. Selling a Car When the Buyer and Seller Have Different Values

The value of a used car to a Buyer is v_B, the value to a Seller v_S, $v = (v_B, v_S) \sim Q$. Values are private information, and it is very intuitive to those who have ever haggled over a price that they should stay that way.

Efficient trade and no coercion require that whenever $v_B > v_S$, that the ownership of the car be transfered from the Seller to the Buyer for some price (aka transfer of value) $t \in [v_S, v_B]$. One question we are after is "Can efficiency be achieved by any type of game?" that is, "Is some efficient allocation implementable?" To answer this, we ask for a mapping from the vector of values to a (trade,transfer) pair with the property that each agent, after learning their own private information but before learning other's information, has expected utility from the allocation function being implemented higher (or at least, no lower) reporting their true type than in lying about it.

A related question, which we will consider first is whether it is *ex post* implementable. This asks that, for each agent, after ALL the private information has been revealed, does the agent like the allocation being implemented more than any one I could have ended up by lying about my private information? This is a MUCH stronger condition, hence much harder to satisfy. Why would we ask for such a thing? Well, it has the advantage that implementing in this fashion gives one a mechanism that does not depend on the mechanism designer having, for example, a complete a description of the environment, e.g. the joint distribution of the Buyer's and Seller's values.

5.1. Ex post implementability and efficiency. The short version of what we are about to see is that the two criteria, ex post implementability and efficiency, are mutually exclusive in the interesting cases.

Suppose that Q is a joint distribution over the four Buyer-Seller valuations $a = (6, 3)$, $b = (12, 3)$, $c = (12, 9)$, and $d = (6, 9)$. Trade should happen at a, b, c, but not at d. Let $t(s)$ be the transfer from the Buyer to the Seller at state s, $s = a, b, c, d$. We know that, because of the no coercion condition, $t(d) = 0$. What can we figure out about the others?

Inequalities for ex post implementability:

(1) First, inequalities from the Buyer, assuming truthfulness by the Seller:
   (a) If the Seller truthfully says their value is 3, then
      (i) the low value Buyer, 6, must prefer $6-t(a)$ to 0, the no coercion inequality, and must prefer $6 - t(a)$ to $6 - t(b)$.
      (ii) the high value Buyer, 12, must prefer $12 - t(b)$ to 0, the no coercion inequality, and must prefer $12 - t(b)$ to $12 - t(a)$.
      (iii) Combining, $t(a) \leq 6$, and $t(a) = t(b)$.
   (b) If the Seller truthfully says their value is 9, then there are some more inequalities that reduce to $t(c) \in [6, 12]$.

(2) Now, inequalities from the Seller, assuming truthfulness by the Buyer:
   (a) If the Buyer truthfully says their value is 12, then
      (i) the low value Seller, 3, must prefer $t(b)$ to 3, the no coercion inequality, and must prefer $t(b)$ to $t(c)$.

(ii) the high value Seller, 9, must prefer $t(c)$ to 9, the no coercion inequality, and must prefer $t(c)$ to $t(b)$.
(iii) Combining, $t(c) \geq 9$, and $t(b) = t(c)$.

Combining all of these, $t(a) = t(b) = t(c)$, i.e. a posted price, and $t(a) \leq 6$ while $t(c) \geq 9$. Oooops.

PROBLEM 8.9. *Verify, using the Chung and Ely ex post taxation principle in the example just given, that no efficient allocation is ex post implementable.*

The posted price intuition for ex post implementability is pretty clear. It gives a great deal of inefficiency for the interesting Q's. Elaborate on this.

5.2. Implementability and efficiency. Above, we made no assumptions about Q in finding inefficiency. That was the advantage of ex post implementability. Now let us suppose that the joint distribution of $v_B = 6, 12$ and $v_S = 3, 9$ is given by

	3	9
6	x	$\frac{1}{2} - x$
12	$\frac{1}{2} - x$	x

PROBLEM 8.10. *As a function of x, find when efficient trade is implementable.*

6. Selling a Car When the Seller Knows the Quality but the Buyer Does Not

Now suppose that the Seller knows the "true" value, v, of a car to a potential Buyer, and the car is only worth $r \cdot v$ to the Seller, $r \in (0, 1)$. This means that efficiency requires that trade happen all the time. Quasi-linear utility again, what is implementable?

PROBLEM 8.11. *Under what (restrictive) conditions is efficient trade ex post implementable?* [*Hint: suppose that the transfer price depends on v. What must be true about the Seller's reports? How does this interact with efficiency?*]

Looking at implementability gives Akerlof's market for lemons story. Here what is at work is that one side of the market knows more than the other. Parametric example, $v = 6, 9$ with prob's $(\alpha, 1 - \alpha)$, $r = 0.9$, $r = 0.1$.

If quality will be revealed, say by whether or not the car breaks down, a Seller with a high value car who can credibly commit to making any repairs can usefully differentiate her/himself from the Seller of a low value car. If the gain to doing so is high enough, we expect that they will do it, and efficiency can be restored.

This has started to get us back into strategic considerations of information transmission, the idea that we can order the people on one side of a transaction according to their information in a fashion that is correlated with the cost of taking some action.

The essential problem with the idea is that signaling activities can be costly.[3]

[3]Schudson's aphorism here.

7. Going to College as a Signal

Some variants on Spence's labor market model.

Nature chooses ability, $a \in \{2, 5\frac{1}{2}\}$ with prob $\frac{1}{2}$ each, a observed by worker not firm, worker chooses education level $s \in \{0, 1\}$, 2 (or more) firms Betrand compete for worker, bidding wages up to productivity of worker, which is equal to a, $u_W = w - 8s/a$ if worker a accepts wage contract w and acquires education level s, $u_f = w - a$ for firm with contract w accepted by worked with ability a.

Pooling eq'm is not stable. Change to allow $s \in [0, \infty)$.

Now suppose that employers compete by offering a wage contract $w(s)$ before the workers choose their education level $s \in [0, 1]$. No pooling eq'm, screening with $w(s) = 2$ for $s < 0.875$, 5.5 else.

8. When Signaling Kills People

We are here looking at going to war as a signal.

With probability 2/3, country 1's secret military research program makes their armies deadlier (i.e. giving higher expected utility in case of war through higher probability of winning and lower losses), and with probability 1/3 the research project is a dud (i.e. making no change in the army's capacities). Knowing whether or not the research program has succeeded, country 1 decides whether or not to declare war on or to remain at peoce with country 2. Country 2 must decide how to respond to war, either $f$ighting or $c$eding territory, all of this without knowing the outcome of 1's research program. With payoffs, one version of the game tree is:

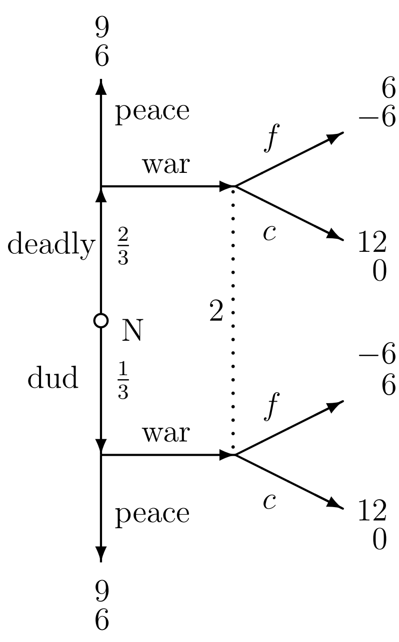

	f	c
(p, p)	$(9, 6)$	$(9, 6)$
(p, w)	$(4, 6)$	$(10, 4)$
(w, p)	$(7, -2)$	$(11, 2)$
(w, w)	$(2, -2)$	$(12, 0)$

Let us modify the previous game by adding an earlier move for country 2. Before country 1 starts its research, country 2 can, at a cost $s > 0$, insert sleepers (spies who will not act for years) into country 1. Country 1 does not know whether or not sleepers have been inserted, and if sleepers are inserted, country 2 will know whether or not 1's military research has made them deadlier. One version of the game tree is:

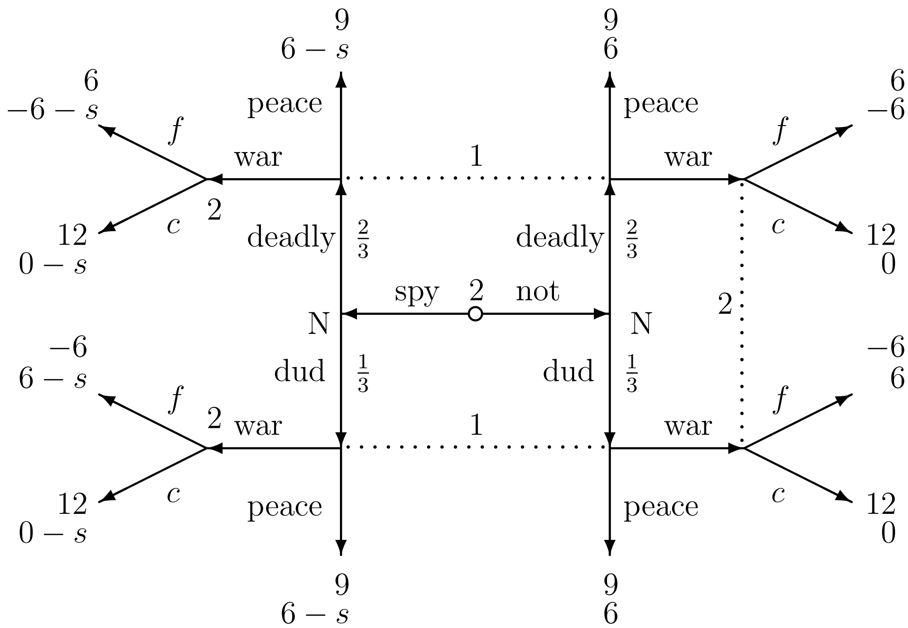

We will now evaluate, in the context of these two versions of this game, the statement, "When spying is cheap, it lowers the probability of war."

CHAPTER 9

Information and its Quality

All random variables are defined on a probability space $(\Omega, \mathcal{F}, P)$. Throughout, $X, X', Y : \Omega \to \mathbb{R}$ are random variables having finite range, and $P(Y = 1) = P(Y = 2) = \frac{1}{2}$.

Time line: first, $\omega \in \Omega$ is drawn according to P; then the value of a signal $X(\omega)$ or $X'(\omega)$ is observed; then a choice $a \in A$ is made; and finally the utility $u(a, Y(\omega))$ is received.

Notation and definitions:

1. $\beta_x^X := P(Y = 1 \big| X = x)$ is the **posterior probability**, aka **belief** that $Y = 1$ given that $X = x$;
2. $V_{u,A}(X) := \max_{f:S \to A} E\, u(f(X), Y)$;
3. the **convex hull** of $C \subset \mathbb{R}^k$ is written $\text{con}\,(C)$; A^S is the set of functions $f : S \to A$;
4. for a (finite) set C, $\Delta(C) = \{p \in \mathbb{R}_+^C : \sum_{c \in C} p_c = 1\}$ is the set of probability distributions over C;
5. the range of a function $X : \Omega \to \mathbb{R}$ is $\mathcal{R}(X) = \{r \in \mathbb{R} : (\exists \omega \in \Omega)[X(\omega) = r]\}$;
6. for p, q distribution on $\mathbb{R}^k$, $k \geq 2$, p is **riskier** that q if for all continuous, concave $g : \mathbb{R}^k \to \mathbb{R}$, $\int g(x)\, dp(x) \leq \int g(x)\, dq(x)$;
7. for distributions p, q on $\mathbb{R}^1$, q is **riskier** that p if for all continuous, concave, non-decreasing $g : \mathbb{R}^1 \to \mathbb{R}$, $\int g(x)\, dp(x) \geq \int g(x)\, dq(x)$.
8. for distributions p, q on $\mathbb{R}^1$, q is a **mean preserving spread of** p if $\int x\, dq(x) = \int x\, dp(x)$ and there is an interval $|a, b| \subset \mathbb{R}$, such that
   a. for all $E \subset (a, b]$, $q(E) \leq p(E)$,
   b. for all $E \subset (b, \infty)$, $q(E) \geq p(E)$, and
   c. for all $E \subset (-\infty, a]$, $q(E) \geq p(E)$.
9. For $r \in \mathbb{R}$, δ_r is point mass on $\mathbb{R}$, that is, $\delta_r(E) = 1_E(r)$.
10. for distributions p, q on $\mathbb{R}^1$, p **first order stochastically dominates** q if for all non-decreasing $g : \mathbb{R} \to \mathbb{R}$, $\int g(x)\, dp(x) \geq \int g(x)\, dq(x)$.

1. Basic Problems About Information Structures and Riskiness

Blackwell's name for what we're calling information structures is "experiments" (Blackwell [**3**], Blackwell [**2**]).

Problem 9.1. *Consider the problem (u, A) and the random variables X, X' given by*

$a = 3$	0	40
$a = 2$	30	30
$a = 1$	50	0
	$Y = 1$	$Y = 2$

$X = 2$	0.3	0.8
$X = 1$	0.7	0.2
	$Y = 1$	$Y = 2$

$X' = 2$	0.1	0.7
$X' = 1$	0.9	0.3
	$Y = 1$	$Y = 2$

where, e.g. $50 = u(a = 1, Y = 1)$ *and* $0.3 = P(X = 2\big|Y = 1)$.

a. *Let* $(\beta, 1-\beta) \in \Delta(\mathcal{R}(Y))$ *be a distribution over the range of* Y. *Give the set of* β *for which*
 i. $\arg\max_{a\in A} \int u(a, y)\, d\beta(y) = \{1\}$,
 ii. $\arg\max_{a\in A} \int u(a, y)\, d\beta(y) = \{2\}$, *and*
 iii. $\arg\max_{a\in A} \int u(a, y)\, d\beta(y) = \{3\}$.

b. *Give the* β^X_x *and the* $\beta^{X'}_{x'}$. *Using the previous problem, give the solutions* f^*_X *and* $f^*_{X'}$ *to the problems* $\max_{f\in A^S} E\, u(f(X), Y)$ *and* $\max_{f\in A^S} E\, u(f(X'), Y)$. *From these calculate* $V_{(u,A)}(X)$ *and* $V_{(u,A)}(X')$. [*You should find that* $X' \succ_{(u,A)} X$.]

c. *Let* $M = (\beta^X_x, P(X = x))_{x=1,2} \in \Delta(\Delta(\mathcal{R}(Y))$ *and* $M' = (\beta^{X'}_{x'}, P(X' = x'))_{x'=1,2} \in \Delta(\Delta(\mathcal{R}(Y))$. *Show directly that* M *is not risker than* M' *and that* M' *is not riskier than* M.

d. *Graph, in* $\mathbb{R}^2$, *the sets of achievable,* Y*-dependent utility vectors for the random variables* X *and* X'. *That is, graph*

$$F_{(u,A)}(X) = \{(E\,(u(f(X), Y)\big|Y = 1), E\,(u(f(X), Y)\big|Y = 2)) \in \mathbb{R}^2 : f \in A^S\}$$

and

$$F_{(u,A)}(X') = \{(E\,(u(f(X'), Y)\big|Y = 1), E\,(u(f(X'), Y)\big|Y = 2)) \in \mathbb{R}^2 : f \in A^S\}.$$

e. *If we allow random strategies, that is, pick* f *according to some* $q \in \Delta(A^S)$, *then the sets of achievable* Y*-dependent utility vectors become* $con\,(F_{(u,A)}(X))$ *and* $con\,(F_{(u,A)}(X'))$. *Show that the same is true if we allow "behavioral strategies," that is,* $f \in \Delta(A)^S$.

f. *Show that* $con\,(F_{(u,A)}(X)) \not\subset con\,(F_{(u,A)}(X'))$ *and* $con\,(F_{(u,A)}(X')) \not\subset con\,(F_{(u,A)}(X))$.

g. *Give a different problem,* (u°, A°) *for which* $X \succ_{(u^\circ,A^\circ)} X'$.

PROBLEM 9.2. *Let* X, X' *be two signals, and define* $X'' = (X, X')$ *to be both signals. Let* S, S' *and* S'' *be the ranges of the three signals*

a. *Show directly that for all* (u, A), $X'' \succsim_{(u,A)} X$ *(hence* $X'' \succsim_{(u,A)} X'$*).*

b. *Show directly that* $\{(\beta^{X''}_{x''}, P(X'' = x''))_{x''\in S''}$ *is riskier than* $\{(\beta^X_x, P(X = x))_{x\in S}$.

c. *Interpret* X *as the result of a doctor's diagnostic test and* X' *is the result of a possible additional test. Show that if* $f^*_{(X,X')}(x, x') = f^*_X(x)$ *for a problem* (u, A), *then there is no point in doing the extra test.*

PROBLEM 9.3. *Most of the following results are in the Müller* [**15**] *article, which covers and extends the famous Rothschild and Stiglitz* [**17**], [**18**] *articles on increases in risk.*

a. *If* q *is a mean preserving spread of* p, *then* q *is riskier than* p.

b. *If* $X \sim p$ *(i.e.* $P(X \in A) = p(A)$*),* $Y \sim q$, *and there is a random variable* Z *such that* $E\,(Z\big|X) = 0$ *and* $X + Z \sim q$, *then* q *is riskier than* p.

c. *Let* $p_0 = \delta_0$, $p_1 = \frac{1}{2}\delta_{-1} + \frac{1}{2}\delta_1$. *Let* $p_2 = \frac{1}{2}\delta_{-1} + \frac{1}{4}\delta_0 + \frac{1}{4}\delta_2$, *and* $p_3 = \frac{1}{4}\delta_{-2} + \frac{1}{4}\delta_0 + \frac{1}{4}\delta_0 + \frac{1}{4}\delta_2 = \frac{1}{4}\delta_{-2} + \frac{1}{2}\delta_0 + \frac{1}{4}\delta_2$. *Continuing in this fashion,* $p_4 = \frac{1}{8}\delta_{-4} + \frac{6}{8}\delta_0 + \frac{1}{8}\delta_4$, *and so on.*
 i. *Show that* p_1 *is a mean preserving spread of* p_0.
 ii. *Show that* p_{k+1} *is a mean preserving spread of* p_k.
 iii. *Show that* $p_k \to_w p_0$.

d. *The previous problem showed that a sequence can become riskier and riskier and still converge to something that is strictly less risky. Show that this cannot happen if* $p_k([a,b]) \equiv 1$ *for some compact interval* $[a,b]$*. Specifically, show that if* $p_k([a,b]) \equiv 1$*, for all* k*,* p_{k+1} *is riskier than* p_k*, and* $p_k \to q$*, then* q *is riskier than all of the* p_k*.*

PROBLEM 9.4. *In each time period,* $t = 1, \ldots,$ *a random wage offer,* $X_t \geq 0$*, arrives. The* X_t *are iid with cdf* F*. The problem is which offer to accept. If offer* $X_t = x$ *is accepted, utility is* $\beta^t u(x)$ *where* $0 < \beta < 1$*, and* $u : \mathbb{R}_+ \to \mathbb{R}$ *is strictly monotonic, concave, and* $\int u(x)\,dF(x) < \infty$*. A "reservation wage" policy is one that accepts all offers of* $\underline{x}$ *or above for some* $\underline{x}$*.*

a. *Show that the optimal policy is a reservation wage policy, and give the distribution of the random time until an offer is expected.*
b. *In terms of* u *and* F*, give the expected utility of following a reservation wage policy with reservation wage* $\underline{x}$*.*
c. *If the offers are, instead, iid* Y_t *with cdf* G *and* G *is riskier than* F*, then, for* $\beta \simeq 1$*, the optimal reservation wage is higher, and the expected utility is also higher.*

PROBLEM 9.5. X_a *is your random income depending on your action* $a \geq 0$*, understood as money that you spend on stochastically increasing* X_a*. The distribution of* X_a *is* $R_{a,c} := cQ_a + (1-c)\mu$*,* $0 \leq c \leq 1$*. Here,* μ *does not depend on* a*, but, if* $a > a'$*, then* Q_a *first order stochastically dominates* $Q_{a'}$*. The parameter* c *is the amount of "control" that you have,* $c = 0$ *means you have no control,* $c = 1$*, means you have the most possible. This question asks you how* a^* *depends on* c*. Intuitively, increasing* c *ought to increase the optimal action, more control means that your actions have more effect.*

Let $f(a,c) = Eu(X_a - a)$ *where* u *is an increasing, concave function. Increases in* a *pull down* $X_a - a$*, hence* $u(X_a - a)$*, by increasing the direct cost, but increase* $X_a - a$ *by stochastically increasing* X_a*.*

a. *Show that* $f(\cdot,\cdot)$ *is not, in generaly, supermodular.*
b. *Suppose that* $f(a,c)$ *is smooth, that we can interchange integration and differentiation, and that the optimum,* $a^*(c)$ *is differentiable. The* $f_a := \partial f/\partial a$ *is equal to*

$$f_a = -\int u'(x-a)d\mu(x) + cd/da[\textit{messy term with } Q_a \textit{ and } \mu].$$

We let $m = [$*messy term with* Q_a *and* $\mu]$*.*
   i. *Show that if* $f_a(a,c) = 0$*, then* $\partial m(a,c)/\partial a > 0$*.*
   ii. *Show that* $f_{a,c} := \partial^2 f/\partial a \partial c = \partial m/\partial a > 0$*.*
   iii. *Show that* $da^*(c)/dc > 0$*.*

2. Elicitation

An expert has a posterior distribution over whether $Y = 0$ or $Y = 1$ with $p = P(Y = 1)$. If one offers the expert a contract which says "Tell me your best estimate, or report, r, of the probability that $Y = 1$, and if, Y turns out to be 1, I'll give you $\log(r)$, while, if it turns out to be 0, I'll give you $\log(1-r)$. Assuming risk neutrality, it's pretty easy to see that the expert will tell you their true posterior distribution.

This kind of contract is called, for obvious reasons, an **elicitation scheme**. What is nice about the one just given is that it is **proper**, that is, no matter what p is, $r^*(p) = p$. (Though the word "scheme" may not be used by anyone but me.)

What we want is the expert to report $r \in \Delta$, to design a function, called a **scoring rule**, $S(r, y)$ such that for all $p \in \Delta$,

$$\text{arg max}_{r\in\Delta} \int S(r, y)\, dp(y) = \{p\}. \tag{10}$$

There are many elicitation schemes, some improper, some proper. The proper logarithmic elicitation scheme given above leans pretty heavily on risk neutrality, $\log(p)$ and $\log(1-p)$ are pretty negative when $p \simeq 0, 1$. One way to understand that is that the value function,

$$V(p) := \max_{r\in\Delta} \int S(r, y)\, dp(y) = \{p\} \tag{11}$$

has some problems staying bounded over all of Δ.

Here are some more elicitation schemes, some proper, some improper, for the finite outcome case $Y \in \{1, \dots, Y\}$. In each case, you should find the associated value function.

A. $S_N(r, y) = r_y 1_{Y=y}$, where the "$_N$" is for "Naive,"
B. $S_L(r, y) = \log(r_y)1_{Y=y}$,
C. $S_Q(r, y) = 2r_y 1_{Y=y} - \sum_y (r_y)^2$, and
D. $S_{Sph}(r, y) = (r_y 1_{Y=y})/\sqrt{\sum_y (r_y)^2}$.

Juries try to process signals about $Y = 0$ or $Y = 1$, innocence or guilt. Their signals are pretty correlated, but not perfectly so. Voting their beliefs involves solving a problem involving maximizing their utility from jailing an innocent/freeing a guilty person. The unanimity rule has some built-in protections.

As well as being proper, a scoring rule can be **effective**, that is, for all p, $[d(q', p) > d(q, p)] \Rightarrow [\int S(q, y)\, dp(y) > \int S(q', y)\, dp(y)]$. Effective scoring rules are proper, but the reverse is not true.

Here's a cute result about effectiveness in the $L^2(A, \mathcal{B}, \lambda)$ case, A a measurable subset of $\mathbb{R}^k$, λ some σ-finite measure, positive on A (e.g. Lebesgue measure). The quadratic scoring rule $S_Q(r, y) = 2r(y) - \|r\|_2^2$ is effective because

$$[d(q, p) < d(q, p)] \Leftrightarrow [\langle q-p, q-p\rangle < \langle q'-p, q'-p\rangle] \Leftrightarrow [2\langle q, p\rangle - \langle q, q\rangle] > [2\langle q', p\rangle - \langle q', q'\rangle]. \tag{12}$$

Problem 9.6. *The patient feels ill. There is a probability β that the patient has a serious disease, $P(d = 1) = \beta = 1 - P(d = 0)$. At a cost of C, the patient can be treated for the disease, the action $a = 1$, at a cost of 0, the treatment can be skipped, $a = 0$. For each of the four (a, d) pairs, there is a random time, $\tau_{(a,d)} > 0$ until the patient recovers. Matching the treatment to the disease is good in the sense that*

- $\tau_{(0,1)}$ *first order stochastically dominates* $\tau_{(1,1)}$, *and*
- $\tau_{(1,0)}$ *first order stochastically dominates* $\tau_{(0,0)}$.

To keep the intuition straight, remember that a large value of $\tau_{(a,d)}$ corresponds to a long time spent feeling ill, so larger τ's are bad.

While ill, the patient's flow utility is s, while healthy it is h, $s < h$, and the patient discounts at a rate r. To be completely explicit, if $\tau = \tau^\circ$ and the patient spends c, the patient's utility is

$$U(\tau^\circ; s, h, r) = r\left[\int_0^{\tau^\circ} se^{-rt}\,dt + \int_{\tau^\circ}^{\infty} he^{-rt}\,dt\right] - c.$$

Assume that $E\,U(\tau_{(1,1)}; s, h, r) - C > E\,U(\tau_{(0,1)}; s, h, r)$, that is, the treatment is worth the cost if the patient is sure they have the disease.

Gregory House M.D. performs his tests and diagnoses the patient. The tests give a signal s, which is too complicated for anyone but Doctor House to understand. It gives rise to a posterior distribution β^s that the patient has the disease. House's reservation wage is $\underline{w} > 0$.

1. *Show that the agent is risk-loving with respect to distributions of τ. Explain why this is a sensible assumption.*
2. *Show that the policy that the patient would like House to implement is to recommend the treatment of $\beta^s \in (\beta^*, 1]$ for some cut-off value β^*. How does β^* vary with s? With h? With C?. With r?*
3. *Show that paying House a retainer, w_h when healthy, and $w_s = 0$ when sick, can truthfully elicit whether or not $\beta^s \in (\beta^*, 1]$. [There is rumor that this used to be a typical Chinese doctor-patient contract.]*
4. *Suppose now that Doctor House's signal is a "no-brainer," that is, a $\{0,1\}$-valued test with a false positive rate $P(s = 1 \mid d = 0) = \epsilon_{fp}$ and a false negative rate $P(s = 0 \mid d = 1) = \epsilon_{fn}$. [Most often, $1 - \epsilon_{fn}$ is called the "sensitivity" of the test, so a small ϵ_{fn} is what defines a very sensitive test.] It is clear that it is not worthwhile consulting Doctor House if $P(\beta^s \in (\beta^*, 1]) = 0$.*
   a. *Give conditions on ϵ_{fp}, ϵ_{fn}, and β such that $P(\beta^s \in (\beta^*, 1]) > 0$.*
   b. *Give conditions under which it is in the patient's interest to consult Doctor House.*

3. Blackwell's Theorem

We start with $(\Omega, \mathcal{F}, P)$ and a random variable Y having a prior distribution, call it $p \in \Delta(\mathbb{R})$. We have different signals $X : \Omega \to \mathbb{R}$ which give rise to posterior distributions $\beta_x^X(A) = P(Y \in A | X = x)$. The first observation is that for all A,

$$\int \beta^X_{X(\omega)}(A)\,dP(\omega) = p(A),$$

that is, the posteriors must integrate back to the prior. This is why conditional probabilities are sometimes called **disintegrations**. The mapping $\omega \mapsto \beta^X_{X(\omega)}$ gives a distribution $Q_X \in \Delta(\Delta(\mathbb{R}))$ with the property that $\int \beta\,dQ_X(\beta) = p$, that is, for each A, $\int \beta(A)\,dQ_X(\beta) = p(A)$.

Some examples: X is independent of Y; $X = Y + \epsilon$, $E\,(\epsilon|X) = 0$; $X = R \cdot Y$, $E(R|Y) = 1$.

Observation: All signal based information structures can be summarized by their associated distribution Q_X. We get at this through **regular conditional probabilities (rcp's)**. The usual σ-field on $\Delta(\mathbb{R})$ is defined in one of two equivalent fashions: the smallest σ-field containing the sets $\{Q : Q(E) \leq r\}$ for E a measurable subset of $\mathbb{R}$ and $r \in \mathbb{R}$; the Borel σ-field generated by the Prokhorov open sets.

Definition 9.1. *An* ***rcp*** *for Y given X is a measurable function $f : \mathbb{R} \to \Delta(\mathbb{R})$ such that for all measurable $A \subset \mathbb{R}$, the function $\omega \mapsto f(X(\omega))(A)$ is a version of $P(Y \in A|X)(\omega)$*

You've seen these before ($\mathbb{R}^2$ with densities, but densities can be very mis-leading (Borel paradox).

Theorem 9.2. *When X, Y are $\mathbb{R}$-valued random variables, regular conditional probabilities exist.*

Defining $\omega \mapsto B_\omega \in \Delta(\mathbb{R})$ by $B_\omega = f(X(\omega))$, each X gives rise to a random variable taking values in $\Delta(\mathbb{R})$. Any random variable has a distribution, call it Q_X in this case, $Q_X \in \Delta(\Delta(\mathbb{R}))$.

Define $V_{(u,A)}(\beta) = \sup_{a\in A} \int u(a, y)\, d\beta(y)$. We define $X \succsim_{(u,A)} X'$ if $E\, V_{(u,A)}(\beta^X) \geq E\, V_{(u,A)}(\beta^{X'})$, equivalently, if $\int V_{(u,A)}(\beta)\, dQ_X(\beta) \geq \int V_{(u,A)}(\beta)\, dQ_{X'}(\beta)$.

Lemma 9.3. *Let $\omega \mapsto X(\omega)$ be a random signal and $\omega \mapsto B_\omega$ the associated rcp. For any (u, A), $X \sim B$.*

Definition 9.4. *The* ***Blackwell order*** *on information structures is $Q_X \succsim_B Q_{X'}$ if for all convex $V : \Delta \to \mathbb{R}$, $\int V(\beta)\, dQ_X(\beta) \geq \int V(\beta)\, dQ_{X'}(\beta)$.*

If $X' = g(X)$, then X' can only be worse than X, never better. More generally, let Z be a random variable independent of X, X', Y. If $X'(\omega) = g(X(\omega), Z(\omega))$, then X' can, intuitively, be no better than X. When the range of X and X' is finite, these can be thought of as Markovian scrambles of the information.

Examples.

Theorem 9.5. [Blackwell] *The following are equivalent:*

(1) $Q_X \succsim_B Q_{X'}$;
(2) for all (u, A), $V_{(u,A)}(X) \geq V_{(u,A)}(X')$; and
(3) $X' = g(X, Z)$ for Z independent of (X, X', Y).

Proof. (Sketch) The first two are equivalent because every convex function on $\Delta(\mathbb{R})$ is the upper envelope of its linear support functions. That the third implies the first follows from the conditional version of Jensen's inequality. I don't know an easy argument for either the first or the second implying the third. □

4. Some Game Theory Applications

More information always makes a single person better off, $(X, X') \succsim_B X$ (e.g. by Jensen's inequality). This is not true when there are two or more people involved. We will first get at this taking a detour through some of the major concepts in cooperative game theory, the core and the Nash bargaining solution. Then we'll go back to non-cooperative game theory.

4.1. Information and the Core. The idea of the core has a very simple emotional logic that most of us learned while we were still children: if we don't like something, we'll take our toys and play by ourselves. From Micro II, every Walrasian equilibrium allocation is in the core. From a more advanced micro course, if the economy is large (in any of a number of senses), then every core allocation is (or is approximately) an equilibrium allocation.

Put Edgeworth box version of core and equilibrium convergence under replication of the economy here.

Now we introduce information into the story, assuming that everyone has the same information, first looking for core allocations when there is no info, then seeing what happens when we add info.

Core theory with differential information is a topic about which we know little except that it's hard to come up with a good starting definition.

PROBLEM 9.7. *There are two possible states of the world, $Y = 1$ and $Y = 2$, with probabilities $\frac{1}{2}$ each. State 1 is good for person A, they have 10 while person B has 6. In state 2, things are reversed, A has 6 while B has 10. Both people have strictly concave vN-M utility functions.*

*An **allocation** a is a pair of pairs, $((x_{A,1}, x_{A,2}), (x_{B,1}, x_{B,2}))$ where $x_{A,1}$ is what A gets in state 1 $x_{B,2}$ is what B gets in state 2, and so on. Thus, the initial allocation is $((10, 6), (6, 10))$. The restriction that must be satisfied $x_{A,1} + x_{B,1} \leq 10 + 6$ and $x_{A,2} + x_{B,2} \leq 6 + 10$.*

*A **coalition** is a subset of the people. An allocation a is **blocked** if some coalition can suggest an allocation a' that uses only the endowments available to people in the coalition and which makes everyone in the coalition better off. An allocation is in the **core** if it is not blocked. The idea is that we only expect core allocations to arise.*

(1) Find the allocation that maximizes A's expected utility subject to B having the expected utility of their initial allocation, $E\,u_B$.
(2) Find the allocation that maximizes B's expected utility subject to A having the expected utility of their initial allocation, $E\,u_A$.
(3) Find all of the core allocations.
(4) Suppose that there is a competitive market for state-contingent consumption. Find the equilibrium prices and allocation.

For the rest of this problem, we suppose that there is a signal, X, that contains information about Y and which is revealed to both people before they start thinking about allocations.

(5) Suppose that X is perfectly informative, e.g. $X(\omega) = [Y(\omega)]^2$.
  (a) After observing each possible value of X, what is in the core?
  (b) Show that both people prefer, ex ante, the core allocations you found above to these post information core allocations.
(6) Suppose now that X is informative, but barely, e.g. having the following joint distribution with Y for some (small) $\epsilon > 0$:

$X = 2$	$0.25 - \epsilon$	$0.25 + \epsilon$
$X = 1$	$0.25 + \epsilon$	$0.25 - \epsilon$
	$Y = 1$	$Y = 2$

  (a) Show that for small enough ϵ, there are core allocations from part 3 above that are in the core after observing each possible value of X.
  (b) Show that for all $\epsilon > 0$, the competitive equilibrium allocation is ex ante worse for both than the equilibrium allocation in part 4.

4.2. Information and Trade. The "no-trade" Theorems (about which a big deal is made) go here.

There is a random variable Y, e.g. whether or not my horse can run faster than your horse. A bet between two people, i and j, is a function $b(Y)$ with the interpretation that i gets $b(y)$ if $Y = y$ and j gets $-b(y)$ in that case. $b(y) \equiv 0$ is the no trade point.

For example, i bets \$100 that his horse can bet j's horse by at least 10 seconds has $Y = 1_{T_i - T_j \geq 10}$, T_i and T_j are the times of the two horses, $b(1) = 100$ and $b(0) = -100$. If i offers 3-to-1 odds, then $b(1) = 100$ and $b(0) = -300$. And so on and so forth.

Before betting, i receives a signal X_i and j receives a signal X_j that contains information about Y. After they see their information, they look for a bet that both would accept. Now we are after a bet $b(X_1, X_2, Y)$ that both would accept. For example, both might be willing to bet \$100 on their horse when and only when they get the signal that their horse was particularly $f$risky this morning. In which case the bet is $b(f_i, f_j, 1) = 100$, $b(f_i, f_j, 0) = -100$ and $b(X_i, X_j, y) = 0$ if $(X_i, X_j) \neq (f_i, f_j)$.

Theorem 9.6. *If u_i and u_j have strictly concave vN-M utility functions and random wealths W_i, W_j that are independent of X_i, X_j, Y, then there is no bet other than the no trade bet that both would be willing to accept.*

If they are both risk neutral, there are many bets that both would be willing to accept, but no bet can be found that generates a positive surplus.

One intuition comes from thinking through as follows: Suppose that i tells himself "I got a really good signal that my horse is going to win, therefore I'd be willing to bet a lot that she's going to win." Pick a bet that i is willing to take and think about j's reaction. "Hmmmm, he's willing to bet \$100 at 3-to-1, the only way that that could be is if he knows something. That changes my willingness to say 'Yes' to the bet."

Intuition is not a proof.

Proof. Suppose that $E\, u_i(W_i + b(X_i, X_j, Y)) \geq E\, u_i(W_i)$ and $E\, u_j(W_j - b(X_i, X_j, Y)) \geq E\, u_j(W_j)$. Since the vN-M utility functions are strictly concave, by Jensen's inequality (yet again) b being non-constant requires that $E\, b > 0$ and $E\, b < 0$, which is hard to arrange. □

4.3. The Nash Bargaining Solution. Let X be the set of options available to a household. A point $x \in X$ may specify an allocation of the rights and duties to the household members. Let $u_i(x)$ be i's utility to the option x, $i = 1, 2$. Let $V = \{(u_1(x), u_2(x)) : x \in X\} \subset \mathbb{R}^2$ be the set of possible utility levels V. Let e be a point in $\mathbb{R}^2$. For $v \in V$, let $L_i(v)$ be the line $L_i(v) = \{v + \lambda e_i : \lambda \in \mathbb{R}\}$, e_i the unit vector in the i'th direction.

Definition 9.7. *A bargaining situation (V, e) is a set $V \subset \mathbb{R}^2$ and a point e satisfying*

(1) V is closed,
(2) V is convex,
(3) $V = V + \mathbb{R}^2_-$, and
(4) for all $v \in V$, $L_1(v) \not\subset V$, and $L_2(v) \not\subset V$.
(5) e is a point in the interior of V

Lemma 9.8. *If $V \subset \mathbb{R}^2$ is convex and $V = V + \mathbb{R}^2_-$, then if there exists $v' \in V$ and $L_i(v') \not\subset V$, then for all $v \in V$, $L_i(v) \not\subset V$.*

Problem 9.8. *Prove this Lemma.*

The interpretation of $e = (e_1, e_2)$ is that e_i is i's reservation utility level, the utility they would get by breaking off the bargaining. This gives a lower bound to what i must get out of the bargaining situation in order to keep them in it. By assuming that e is in the interior of V, we are assuming that there is something to bargain about.

Definition 9.9. *The Nash bargaining solution is the utility allocation that solves*

$$\max\ (v_1 - e_1) \cdot (v_2 - e_2) \ \text{ subject to } \ (v_1, v_2) \in V,\ v \geq e.$$

Equivalently,

$$\max_{x \in X}\ (u_1(x) - e_1)(u_2(x) - e_2) \text{ subject to } \ (u_1(x), u_2(x)) \geq e.$$

It is worthwhile drawing a couple of pictures to see what happens as you move e around. Also check that the solution is invariant to affine positive rescaling of the players' utilities.

Problem 9.9. *Let $s^*(e) = (s_1^*(e_1, e_2), s_2^*(e_1, e_2))$ be the Nash bargaining solution, i.e. solves*

$$\max (v_1 - e_1) \cdot (v_2 - e_2) \ \text{ subject to } \ (v_1, v_2) \in V.$$

Suppose also that

$$V = \{(v_1, v_2) : f(v_1, v_2) \leq 0\}$$

where f is a differentiable, convex function with $\partial f / \partial v_i \neq 0$.

(1) *Where possible, find whether the following partial derivatives are positive or negative:*

$$\frac{\partial s_1^*}{\partial e_1},\ \ \frac{\partial s_1^*}{\partial e_2},\ \ \frac{\partial s_2^*}{\partial e_1},\ \ \frac{\partial s_2^*}{\partial e_2}.$$

(2) *Where possible, find whether the following partial derivatives are positive or negative:*

$$\frac{\partial^2 s_1^*}{\partial e_1^2},\ \ \frac{\partial^2 s_1^*}{\partial e_1 \partial e_2},\ \ \frac{\partial^2 s_2^*}{\partial e_2}.$$

(3) *Consider the following variant of the Nash maximization problem,*

$$\max\ ((av_1 + b) - (ae_1 + b)) \cdot (v_2 - e_2) \ \text{ subject to } \ (v_1, v_2) \in V$$

where $a > 0$. Show that the solution to this problem is $(as_1^* + b, s_2^*)$ where (s_1^*, s_2^*) is the Nash bargaining solution we started with. In other words, show that the Nash bargaining solution is independent of affine rescalings. (You might want to avoid using calculus arguments for this problem.)

It is remarkable that this solution is the only one that satisfies some rather innocuous-looking axioms. We'll need

Definition 9.10. *A* **bargaining solution** *is a mapping $(V, e) \mapsto s(V, e)$, $s \in V$, $s \geq e$. The solution is* **efficient** *if there is no $v' \in V$ such that $v' > s$.*

We'll also need

Definition 9.11. *For $(x_1, x_2) \in \mathbb{R}^2$, a* ***positive affine rescaling*** *is a function $A(x_1, x_2) = (a_1 x_1 + b_1, a_2 x_2 + b_2)$ where $a_1, a_2 > 0$.*

Here are some reasonable looking axioms for efficient bargaining solutions:

(1) Affine rescaling axiom: The solution should be independent of positive affine rescalings of the utilities. That is, $s(AV, Ae) = A(s(V, e))$ for all positive affine rescalings A.
(2) Midpoint axiom: If $V \cap \mathbb{R}^2_+ = \{(u_1, u_2) : u_1 + u_2 \leq 1\}$ and $e = (0, 0)$, then the midpoint of the line, $(1/2, 1/2)$, is the solution.
(3) Independence of irrelevant alternatives axiom: If $s(V, e) \in V' \subseteq V$, then $s(V', e) = s(V, e)$.

Theorem 9.12 (Nash). *There is only one efficient bargaining solution that satisfies these three axioms, and it is the solution to the problem*

$$\max\ (v_1 - e_1) \cdot (v_2 - e_2) \ \ \textit{subject to} \ \ (v_1, v_2) \in V.$$

Problem 9.10. *Prove this theorem.*

Back to property rights, the household problem when facing a set of options X is now modeled as

$$\max_{x \in X,\ u_i(x) > e_i,\ i=1,2} (u_1(x) - e_1)(u_2(x) - e_2).$$

In effect, $w(x) = (u_1(x) - e_1)(u_2(x) - e_2)$ is the household utility function. Khan[1] argues, in the context of patenting activity as married women in the U.S. gained the right to sign legally binding contracts, that changing the property laws does not change X. Therefore, changes in the property laws can only affect the optimal behavior in the above problem if they change the e_i. This may be a reasonable way to understand the legal changes – they gave women a better set of outside options, which is captured by increasing the womens reservation utility level.

Never one to let good enough alone, I'd like to look at a different but still reasonable set of axioms. This one leads to the Kalai-Smorodinsky bargaining solution.[2] For one of our bargaining problems (V, e), let ∂V denote the (upper) boundary of V, and let $\bar{u}_i^V = \max\{u_i : (u_i, e_i) \in V\}$.

(1) Affine rescaling axiom: The solution should be independent of affine rescalings of the utilities, that is, $s(AV, Ae) = A(s(V, e))$ for all positive affine rescalings A.
(2) Box axiom: If $V \cap \mathbb{R}^2_+ = \{(u_1, u_2) : u_i \leq u_i^\circ\}$, then $s(V, e) = (u_1^\circ, u_2^\circ)$.
(3) Proportional increases axiom: Suppose that $s(V, e) \in \partial V'$ and that $(\bar{u}_1^V, \bar{u}_2^V)$ and $(\bar{u}_1^{V'}, \bar{u}_2^{V'})$ are proportional. Then $s(V, e) = s(V', e)$.

Geometrically, to find the Kalai-Smorodinsky bargaining solution, one shifts so that $e = (0, 0)$, solves the problem $\max\{\lambda : \lambda \geq 0,\ \lambda(\bar{u}_1^V, \bar{u}_2^V) \in V\}$ for λ^*, and set $s(V, e) = \lambda^*(\bar{u}_1^V, \bar{u}_2^V)$.

[1]Zorina Khan, "Married Women's Property Laws and Female Commercial Activity: Evidence from United States Patent Records, 1790-1895," *Journal of Economic History*, **56**(2), 356-388 (1996).
[2]For more detail, see the paper by Nejat Anbarci, "Simple Characterizations . . . ," (1995).

PROBLEM 9.11. *This is a directed compare and contrast problem:*

(1) *Give two* (V, e) *where the Nash solution is the same as the Kalai-Smorodinsky solution.*
(2) *Give two* (V, e) *where the Nash solution is different than the Kalai-Smorodinsky solution.*
(3) *Let* $s^{KS}(V, e)$ *denote the Kalai-Smorodinsky solution. If possible, find whether or not* s_i^{KS} *is increasing or decreasing in* e_j, $i, j \in \{1, 2\}$.
(4) *Let* $s^{KS}(V, e)$ *denote the Kalai-Smorodinsky solution. If possible, find whether or not* s_i^{KS} *is increasing or decreasing in* $\bar{u}_j^V$, $i, j \in \{1, 2\}$.

PROBLEM 9.12. *Nash's bargaining solution let us "explain" the effect of changes in property laws as increases in womens' reservation utility levels. What would the corresponding "explanation" be for the Kalai-Smorodinsky solution?*

There is another, non-cooperative way to look at the Nash bargaining solution, this one due to Binmore following a suggestion of Nash's (*Nash bargaining theory II*, London School of Economics, (1981)). To see why it is so striking requires passing through the basics of (one of) Schelling's insights.

Consider the problem of dividing a pie between two players, let x and y denote 1 and 2's payoffs, (x, y) is **feasible** if $x \geq e_1$ and $y \geq e_2$ and $g(x, y) \leq 1$ where g is a smooth function w/ everywhere strictly positive partial derivatives. Let V denote the set of feasible utility levels and assume that V is convex. Consider the simultaneous move game where the players suggest a division, if their suggestions are feasible, that's what happens, if they are not feasible, e is what happens.

LEMMA 9.13 (Schelling). *Any efficient feasible division is an equilibrium.*

This is Schelling's basic insight, for something to be an equilbrium in a game of division, both have to believe that it is an equilibrium, but not much else needs to happen. (In this part of Schelling's analysis he lays the framework for much of the later work on common knowledge analyses of games.) Especially in a game like this, the conditions for a Nash equilibrium seem to need something more before you want to believe in them.

Back to Binmore's version (as suggested by Nash). Suppose that the game has a random feasible set, that is, suppose that the agents pick their offers, the game proceeds as above, except that a pair of offers in infeasible when $g(x, y) > z$, and the distribution of z has a cdf with $F(\underline{z}) = 0$, $F(\bar{z}) = 1$ and $\underline{z} < 1 < \bar{z}$. Binmore's observation is

LEMMA 9.14. *If* F^n *is a sequence of full support distributions converging weakly to point mass on* 1*, then the equilibrium outcomes of this game converge Nash's axiomatic bargaining solution.*

It is worth proving this result for yourself.

4.4. Changing Information in Noncooperative Games. We've seen spying, there are communication eq'a, info leakage, repeated game issues.

Revisiting Beer-Quiche with different signals, on both sides, about the private information.

5. Second Order Uncertainty

Consider, if you will, the question, "What is the probability that the 5'th brightest star in the Pleiades star cluster has a planet that is conceivably habitable by humans?" Intuitively, we understand the answers "I haven't the faintest flying fart of a clue," "50-50 plus or minus 49.9999," or, somewhat more implausibly, "p where p is distributed according to a $\beta(r, s)$ distribution," and many other indications of high levels of uncertainty. I offer you a proper, or even an effective scoring rule. And just to make matters clear, I commit to shooting you in the foot if you don't give a report.

. . .

Here is a sketch of an approach to preferences when there is ambiguity, modeled in a specific way, and to Nash equilibrium with these kinds of descriptions.

5.1. Ellsberg Urns. An (Ellsberg) urn is known to contain 90 balls, 30 of which are Red, each of the remaining 60 can be either Green or Blue. The decision maker is faced with the urn, the description just given, and two pairs of choice situations.

(1) Single ticket choices:
   (a) The choice between the Red and the Green ticket.
   (b) The choice between the Red and the Blue ticket.
(2) Pairs of ticket choices:
   (a) The choice of the R&B or the G&B pair.
   (b) The choice of the R&G or the B&G pair.

In each situation, after the DM makes her choice, one of the 90 balls will be picked at random. If the ball's color matches the color of (one of) the chosen ticket(s), the decision maker gets $\$1,000$, otherwise they get nothing.

Typical preferences are

$$R \succ G \text{ and } R \succ B,$$
$$R\&B \prec G\&B \text{ and } R\&G \prec B\&G.$$

It is hard to find a probability-based explanation for these preferences because it would yield both

$$P(R) > P(G) \text{ and } P(R) > P(B),$$

as well as

$$P(R) + P(B) < P(G) + P(B),\ P(R) + P(G) < P(B) + P(G).$$

This argument does not encompass all stories about beliefs about the distribution of the numbers of Blues and Greens. It does not encompass the idea that someone or something is cheating. If someone or something changes the number of balls after you've picked a color, the additions above are irrelevant.

Will return to this hostile universe/experimenter idea.

5.1.1. *Preferences over sets of probabilities.* Choices lead to outcomes. Outcomes are sets of probabilities, possibly singleton sets. Preferences over sets of probabilities induce preferences over choices.

The probability that the Red ticket wins is $\frac{1}{3}$. That is, the action "choose Red" is risky, with the associated probability $\frac{1}{3}$. The actions "choose Blue" and "choose Green" are ambiguous, leading to the interval of probabilities $[0, \frac{2}{3}]$.

Choosing the Blue&Green pair is risky, $\frac{2}{3}$, choosing the other two pairs is ambiguous, $[\frac{1}{3}, 1]$.

As noted, the typical preferences are

$$\{\tfrac{1}{3}\} \succ [0, \tfrac{2}{3}] \text{ and } \{\tfrac{2}{3}\} \succ [\tfrac{1}{3}, 1].$$

People prefer the center of the interval to the interval itself.

5.1.2. *Linear preferences on intervals of probabilities.* We add sets, A, B, in a vector space with

$$A + B = \{a + b : a \in A,\ b \in B\}.$$

We multiply by non-negative constants with

$$\alpha A = \{\alpha a : a \in A\}.$$

Convex and other linear combinations simply paste these two operations together.

Every interval $[q, s]$ is of the form $[c - r, c + r]$ for $c = (q + s)/2$, $r = (s - q)/2 \geq 0$. Adding intervals involves adding the centers and the radii, multiplying an interval by $\alpha \geq 0$ multiplies the center and the radius by α.

5.1.3. *Linear preferences on intervals of probabilities.* Under study are non-trivial, continuous, linear preferences, $\succsim$, on, $\mathbb{K} = \mathbb{K}([0, 1])$, the closed convex subsets of $[0, 1]$.

The singleton subsets of $[0, 1]$ will be denoted either $[p, p]$ or $\{p\}$. Restricted to the class of singleton sets, continuous linear preferences are vN-M preferences.

Non-trivial, continuous, affine preference can be represented by a function $U : \mathbb{K} \to \mathbb{R}$ that is

(1) continuous, $c^n \to c$ and $r^n \to r$, implies $U([c^n - r^n, c^n + r^n]) \to U([c - r, c + r])$,
(2) linear, $U(\alpha[c-r, c+r] + (1-\alpha)[c'-r', c'+r']) = \alpha U([c-r, c+r]) + (1-\alpha)U([c'-r', c'+r'])$, and
(3) normalized, $U(\{0\}) = 0$ and $U(\{1\}) = 1$.

Lemma 9.15. *The unique representation of preferences satisfying* (1), (2) *and* (3) *is* $U([c - r, c + r]) = c - vr$.

Any $v > 0$ represents the urn preferences above.

If $U([c - r, c + r]) = c - vr$, then

(1) v measures the tradeoff between risk and uncertainty,
(2) asking that $U(\{a\}) \leq U[a, b] \leq U(\{b\})$ is the same as requiring $|v| \leq 1$, and
(3) v can be elicited by observing choices between risky and ambiguous outcomes,

The condition

$$U(\{q\}) \leq U[q, s] \leq U(\{s\})$$

is called **balance**. In this two outcome case, it asks only that people

(1) would not be willing to pay to learn that they are surely facing the worst of the possible risky situations, nor

(2) be willing to pay to move from knowing the probability of the good outcome is s to knowing that the probability of the good outcome is bounded above by s.

Balance may be violated if there is a crucial probability level, call it c, with two properties: $q < c < s$, and some crucial decision about (say) preparations for future disasters changes depending on whether or not the probability is above or below c.

5.2. Games with ambiguity. A game with ambiguity is $\Gamma_{Amb}(u, U) = (A_i, (u_i, U_i))_{i \in I}$. Each A_i is finite here. There is a redundancy in notation because each U_i is an extension of the corresponding u_i.

Let $K_i \in \mathbb{K}_i := \mathbb{K}(\Delta(A_i))$ be a non-empty, compact, convex subset of $\Delta(A_i)$.

A Nash equilibrium without ambiguity can be understood as a vector, $\sigma^* = (\sigma_i^*)_{i \in I}$, of **beliefs**, which, when held by all players, belongs, componentwise, to the set of best responses to the beliefs. That is,

$$\text{for all } i,\ \{\sigma_i^*\} \subset Br_i(u_i, \sigma^*).$$

5.2.1. *Equilibria with ambiguity.* An equilibrium with ambiguity is a vector of sets of beliefs, $K^* = (K_i^*)_{i \in I}$. There are (at least) two possible ways to write down equilibrium conditions on the K_i^*. Both use the best response correspondence.

The best response correspondence has the usual properties, being the set of probabilities putting mass 1 on the pure strategy best responses.

Each i evaluates the set of distributions $\{\mu_i\} \times \Pi_{j \neq i} K_j^*$ using the utility function $U_i(\cdot)$. For fixed K^*, the mapping

$$\mu_i \mapsto U_i(\{\mu_i\} \times \Pi_{j \neq i} K_j^*)$$

is continuous and linear.

Nash equilibrium

$$\text{for all } i,\ \{\sigma_i^*\} \subset Br_i(u_i, \sigma^*).$$

Potential equilibrium conditions:

$$\text{for all } i,\ Br_i(U_i, K^*) \subset K_i^*.$$

This requires that players' beliefs about others include the best response set. This allows people to be wrong, but asks that their beliefs include the true best responses.

Observe: $(K_i^*)_{i \in I} = (\Delta_i)_{i \in I}$ is always an equilibrium if we use this loose condition. This is not an attractive property for a solution concept.

Nash equilibrium

$$\text{for all } i,\ \{\sigma_i^*\} \subset Br_i(u_i, \sigma^*).$$

Definition 9.16. $K^* = (K_i^*)_{i \in I} \in \times_{i \in I} \mathbb{K}_i$ *is an* **equilibrium set for** $\Gamma_{Amb}(u, U)$ *if for all* $i \in I$,

$$K_i^* \subset Br_i(U_i, K^*).$$

This requires that players' beliefs about others be a subset of what optimization might make the others choose to do.

5.2.2. *Generalities about equilibria with ambiguity.* There are minimally ambiguous equilibria, we've studied them before,

LEMMA 9.17. $(K_i^*)_{i\in I} = (\{\sigma_i^*\})_{i\in I}$ *is an equilibrium set for* $\Gamma_{Amb}(u, U)$ *if and only if* $(\sigma_i^*)_{i\in I}$ *is a Nash equilibrium.*

There are maximal sets of ambiguous equilibria (I believe),

CONJECTURE 9.18. *For every game* $\Gamma_{Amb}(u, U)$*, there is an ambiguous equilibrium set* $(K_i^*)_{i\in I}$ *with the property that no ambiguous equilibrium set* $(K_i')_{i\in I}$ *satisfies* $K_i^* \subset K_i'$ *for all* i *with strict inequality for at least one* $i \in I$.

5.2.3. *Interpreting equilibria with ambiguity.* We often think of equilibrium as a point where "things have settled down." Ambiguity is incomplete knowledge about the probabilities over consequences associated with different actions. There is a tension between "incomplete knowledge" and having "settled down." In games, an equilibrating notion is providing bounds on what people know.

It may be possible to interpret ambiguous sets of equilibria as sets of distributions that happen on the way to an equilibrium along some kind of learning dynamic. That would be nice.

Another possible interpretation of equilibria with ambiguity is that the ambiguity part of the preferences represents something about how future interactions will happen.

This is all too vague. Some examples.

5.3. Examples of 2×2 games with ambiguity. Preferences are assumed to be balanced, that is, ambiguity aversion cannot be so strong as to make someone prefer knowing that the distribution is given by the worst in a set. In this context, let us first look at U_i-dominance solvable games

5.4. U_i-dominance solvable games. Bluntly, ambiguous equilibrium in these games requires that there not be any ambiguity.

DEFINITION 9.19. *An action* a_i *is* U_i*-dominated if for all* K*, there exists a* b_i *such that* $U_i(K\backslash a_i) < U_i(K\backslash b_i)$.

Let us consider the following game:

	Left	Right
Top	$(8, 1)$	$(6, 0)$
Bottom	$(4, 0)$	$(2, 1)$

Note that this is not a complete description of the game because it does not specify either player's attitude toward ambiguity. To describe their attitudes, only two numbers are needed per player. For 1 they are v_{Top} and v_{Bottom}, which give the two tradeoffs between the radius and the center of the set of probabilities describing 1's beliefs the probability that 2 is playing Left or Right.

If the v_{Top} and a v_{Bottom} for 1's yield balanced preferences, then for all K_2,

$$6 \leq U_1(\text{Top} \times K_2) \leq 8 \text{ and}$$

$$2 \leq U_1(\text{Bottom} \times K_2) \leq 4.$$

This yields $Br_1(U_1, K) \equiv \{\text{Top}\}$. Therefore, in any equilibrium $K_1^* = \{\text{Top}\}$. Equilibrium requires that 2's beliefs be that 1 will surely play Top.

Given 2's surety, 2's preferences reduce to expected utility preferences. The unique best response is Left, and the unique ambiguous equilibrium for this U_i-dominance solvable game is the unique Nash equilibrium.

The equilibrium requirement, $K_1^* \subset Br_1(U_i, K^*)$, has informational implications for 2. Equilibrium requires that 2 have thoroughly learned something. Bluntly, ambiguous equilibrium in this game requires that there not be any ambiguity.

The following game is u_i-dominance solvable for a Nash equilibrium, but is not U_i-dominance solvable.

	Left	Right
Top	$(8, 1)$	$(5, 0)$
Bottom	$(7, 0)$	$(4, 1)$

For many different intervals $[q, s]$, $0 < q < s < 1$ there exist $v_{\text{Top}} > v_{\text{Bottom}}$ such that 1's preferences are balanced and

$$U_1(\text{Top} \times [a, b]) = U_1(\text{Bottom} \times [q, s]).$$

Equal likelihood of Top and Bottom makes 2 indifferent. If $v_{\text{Left}} = v_{\text{Right}}$, then for any $0 < r < \frac{1}{2}$,

$$K^* = ([\tfrac{1}{2} - r, \tfrac{1}{2} + r], [a, s])$$

is an ambiguous equilibrium.

The u_1-dominance argument for Top loses its force if $v_{\text{Top}} > v_{\text{Bottom}}$ and 1's beliefs about 2 are a (widish) interval. The inequality between v_{Top} and v_{Bottom} says that ambiguity is more tolerable to 1 when 1 chooses Bottom.

Attitudes toward ambiguity about other's actions may be dependent on own actions. This is a special case of including future consequences from the present game. More generally, this is a way of supposing that the people who would choose the different actions are different in ways not captured by the vN-M utilities given in the matrix.

Suppose that 1 is a parent who could (Top) make sure their own child knew about the perils and pleasures of sex and drugs and rock and roll, or could (Bottom) try to protect them from such things by telling them about the evils inherent in worldy pleasures.

2 is a friend of the child who could (Left) go to a party with the sex and drugs and rock and roll crowd or (Right) go to a party with the chaperoned church crowd. The choices made by the friends of one's child(ren) affect a parent's utility.

Not knowing the distribution over what the child's friend is doing is more bothersome to a parent choosing Top because ... , and less bothersome to a parent choosing Bottom because ... .

5.4.1. *Monitoring games with ambiguity.* Suppose that a monitoring agency has a budget devoted to monitoring for violations of worker safety laws. The budget is sufficient to monitor b of the M firms for violations in any given period. Suppose that the profits are π_v when violating worker safety laws, π when operating according to the law, and that the fine for being caught is f. Suppose also that

$$\pi_v > \pi > \pi_v - f,$$

so that for some critical probability, p^*, the firm is just indifferent between violating and not violating worker safety laws.

If the budget is large enough that $b/M > p^*$, then random monitoring of all firms achieves complete compliance. For a smaller budget, the strategy that maximizes the amount of compliance picks a set $M' \subset M$ of the firms, and monitors them at random, where $\#M'$ is the largest integer satisfying $b/(\#M') > p^*$.

Another possibility is to introduce ambiguity about the size of b or f into the firms' decision problems.

With ambiguity about b, the size of the monitoring budget, and ambiguity aversion, one can get deterrence with a smaller budget, effectively stretching the budget of the monitoring agency. One does not expect this to be a permanent, or an equilibrium, solution.

One can keep the legal system in a turmoil, say by regular Congressional re-writing of penalties on non-compliant firms so that the distribution of punishments associated with conviction is not known. This has a larger deterrent effect on the ambiguity averse.

A caveat: making penalties random undercuts the perceived legitimacy of the legal system, and it is surely this legitimacy that keeps many obeying the laws.

5.4.2. *Coordination games with ambiguity.* Suppose that the two players are romantically involved, and must decide, after at least one cell phone dies on a Friday evening, where to meet. They know it will be in one of two places, the wine bar or the sports bar. The players disagree about which bar is more enjoyable, receive 0 utility from going to their least favorite, 1 utility from their most favorite, and, in either place, an additional jolt of pleasure worth 2 utils if they can be with their sweetie. Combining, this gives the coordination game

	Wine	Sports
Wine	$(3, 2)$	$(1, 1)$
Sports	$(0, 0)$	$(2, 3)$

	Wine	Sports
Wine	$(3, 2)$	$(1, 1)$
Sports	$(0, 0)$	$(2, 3)$

The Nash equilibria are (Wine, Wine), (Sports, Sports), and $((\frac{3}{4}, \frac{1}{4}), (\frac{1}{4}, \frac{3}{4}))$, with utilities $(3, 2)$, $(2, 3)$, and $(1\frac{1}{2}, 1\frac{1}{2})$ respectively. Mis-coordination is the worst of the equilibria, the risk of not meeting their sweeties hurts both players.

We will now see that ambiguity aversion can make the players even more miserable than they are when all they face is risk.

	Wine	Sports
Wine	$(3, 2)$	$(1, 1)$
Sports	$(0, 0)$	$(2, 3)$

$((\alpha^*, 1-\alpha^*), (\beta^*, 1-\beta^*)) = ((\frac{3}{4}, \frac{1}{4}), (\frac{1}{4}, \frac{3}{4}))$ is an mixed equilibrium with a minimal amount of ambiguity. It yields utilities $(u_1^*, u_2^*) = (1\frac{1}{2}, 1\frac{1}{2})$.

Suppose that $[c_\beta - r_\beta, c_\beta + r_\beta]$ is 1's belief about 2's choice. Then 1's utility to Wine and Sports are

$$U_1(\{\text{Wine}\} \times [c_\beta - r_\beta, c_\beta + r_\beta]) = 3 \cdot c_\beta + 1 \cdot (1 - c_\beta) - v_{\text{Wine}} r_\beta,$$
$$U_1(\{\text{Sports}\} \times [c_\beta - r_\beta, c_\beta + r_\beta]) = 0 \cdot c_\beta + 2 \cdot (1 - c_\beta) - v_{\text{Sports}} r_\beta.$$

1 being indifferent between the bars requires

$$3 \cdot c_\beta + 1 \cdot (1 - c_\beta) - v_{\text{Wine}} r_\beta = 0 \cdot c_\beta + 2 \cdot (1 - c_\beta) - v_{\text{Sports}} r_\beta, \text{ or}$$
$$c_\beta = \beta^* + \tfrac{1}{4} r_\beta \left(v_{\text{Wine}} - v_{\text{Sports}}\right)$$

where $\beta^* = \frac{1}{4}$ is 2's unambiguous, mixed equilibrium strategy.

$[c_\beta - r_\beta, c_\beta + r_\beta] \subset [0, 1]$ requires $r_\beta \leq c_\beta \leq 1 - r_\beta$. Combining, one can graph all of the possible (c_β, r_β) pairs that are part of some ambiguous equilibrium.

Recall that equilibrium utilities for the mixed strategy unambiguous equilibrium are $(u_1^*, u_2^*) = (1\frac{1}{2}, 1\frac{1}{2})$. In the ambiguous equilibria, 1's utility is

$$u_1^* - \tfrac{1}{2} r_\beta \left(v_{\text{Wine}} + v_{\text{Sports}}\right),$$

with a similar expression for 2's utility.

Here, the players suffer not only the risk of missing their sweetie, they suffer the pangs of not even being able to do a very good job guessing how likely their sweetie is to go to which bar.

CHAPTER 10

Repeated Games Without Private Information

Here we take a game $\Gamma = (A_i, u_i)_{i\in I}$ and play it once at time $t = 1$, reveal to all players which $a_i \in A_i$ each player chose, then play it again at time $t = 2$, reveal, etc. until N plays have happened, $N \leq \infty$.

Notice that there is no private information in these games. In the next chapter we will cover repeated games in which there is private information, with special emphasis on the contributions in the paper by Abreu, Pearce, and Stacchetti.

The basic observation is, roughly, "repeating games can greatly expand the set of equilibria." This section of the course is devoted to making this statement meaningful and qualifying it.

There are four basic kinds of reasons to study what happens in repeated games, they are not mutually exclusive. First, it delivers an aesthetically pleasing theory. Second, it has the benefit of making us a bit more humble in our predictions. (Humility is not a universally admired quality.) Third, we believe that many of the most interesting economic interactions are repeated many many times, it is good to study what happens in these games. Fourth, economics, and equilibrium based theories more generally, do best when analyzing routinized interactions. In game theory models, routinized interactions makes it easier to believe that each i has figured out not only that they are solving the problem

$$\max_{\sigma_i\in\Delta(A_i)} u_i(\sigma^*\backslash\sigma_i),$$

but that the solution is σ_i^*. Don't take the last reason too seriously, the theory of repeated games we will look at first is not particularly good for analyzing how people might arrive at solving the equilibrium maximization problem.

1. The Basic Set-Up and a Preliminary Result

When playing the game $N < \infty$ times, the possible history space for the game is H^N, the product space

$$H^N = \underbrace{S \times \ldots \times S}_{N \text{ times}}.$$

For any $h^N = (s^1, \ldots, s^t, s^{t+1}, \ldots, s^N) \in H^N$, i's payoffs are

$$U_i^N(h^N) = \tfrac{1}{N}\sum_{t=1}^N u_i(s^t).$$

When playing the game "infinitely often", the possible history space is

$$H^\infty = (s^1, s^2, \ldots) \in S \times S \times \cdots,$$

and the payoffs are discounted with discount factor δ,

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

The important point about these payoffs is that they are on the same scale as u, specifically, for all N and all δ,

$$u(S) \subset U^N(H^N),\ u(S) \subset U_\delta^\infty(H^\infty), \text{ and}$$
$$U^N(H^N) \subset \mathbf{co}\,(u(S)),\ U_\delta^\infty(H^\infty) \subset \mathbf{co}\,(u(S)).$$

These are true because in all cases, the weights on the $u_i(s^t)$ add up to 1, $1 = 1$, $\frac{1-\delta}{\delta}\sum_{t=1}^\infty \delta^t = 1$, and $\underbrace{1/N + \cdots + 1/N}_{N \text{ times}} = 1$. The following will be important several times below.

Problem 10.1. *For all $v \in \mathbf{co}\,(u(S))$ and for all $\epsilon > 0$,*

$$(\exists N')(\forall N \geq N')(\exists h^N \in H^n)\|U^N(h^N) - v\| < \epsilon, \text{ and}$$
$$(\exists \underline{\delta} < 1)(\exists h^\infty \in H^\infty)(\forall \delta \in (\underline{\delta}, 1)\|U_\delta^\infty(h^\infty) - v\| < \epsilon.$$

[Hint for the averaging case: if $v = \sum_k \alpha_k u(s_k)$, α_k rational, then one can construct a cycle of length M so that the average over each cycle is exactly v. When the number, k, of cycles is large, the difference between v and the N-average for $kM < N \leq (k+1)M$ is less than ϵ. Hint for the discounted case: take the same cycle, and define what you are doing at t as $x(t)$. Look at the random variable T with $P(T = t) = \frac{1-\delta}{\delta}\delta^t$. What you want to show is that $E\,u(x(T)) = \sum_{t=1}^\infty \frac{1-\delta}{\delta}\delta^t u(x(t)) \simeq v$. Let $S_k = \{kM+1, kM+2, \ldots, (k+1)M\}$. Conditioning on $T \in S_k$, find $E\,(x(T)|T \in S_k)$. This is independent of k. Therefore, $E\,x(T)$ is equal to $E\,(x(T)|T \in S_k)$. This goes to y as $\delta \uparrow 1$.]

As always, strategies are complete contingent plans. For completeness, we define H^0 as a one-point set, $H^0 = \{h^0\}$. A behavioral strategy for i is, for every $t \in \{1, \ldots, N\}$, a mapping

$$\sigma_i^t : H^{t-1} \to \Delta_i,$$

so that a strategy for i is a sequence $\sigma_i = (\sigma_i^t)_{t=1}^N$, and Σ_i^N is the set of all behavioral strategies. Each vector behavioral strategy $\sigma = (\sigma_i)_{i\in I}$, specifies an outcome distribution over H^N, denoted by $\mathbb{O}(\sigma)$. Playing the strategy σ starting from a history $h^{t-1} \in H^{t-1}$ gives the outcome $\mathbb{O}(\sigma|h^{t-1})$.

Summarizing, for $N < \infty$,

$$\Gamma^N = (\Sigma_i^N, U_i^N)_{i\in I},$$

for $N = \infty$,

$$\Gamma_\delta^\infty = (\Sigma_i^\infty, U_i^\delta)_{i\in I}.$$

A vector σ^* is an equilibrium if the usual conditions hold, and the set of equilibria is $Eq(\Gamma^N)$ or $Eq(\Gamma_\delta^\infty)$ as $N < \infty$ or $N = \infty$. A vector σ^* is a sub-game perfect equilibrium if it is a Nash equilibrium given any starting history h^{t-1}, $t \in \{1, \ldots, N\}$. The set of sub-game perfect equilibria is $SGP(\Gamma^N)$ or $SGP(\Gamma_\delta^\infty)$ as $N < \infty$ or $N = \infty$.

Since the strategy sets are very different in Γ, Γ^N, and Γ_δ^∞, the way that we will be comparing the equilibrium sets is to compare $u(Eq(\Gamma))$, $U^N(Eq(\Gamma^N))$, $U^N(SGP(\Gamma^N))$, $U_\delta^\infty(Eq(\Gamma_\delta^\infty))$ and $U_\delta^\infty(SGP(\Gamma_\delta^\infty))$. The starting point is

Lemma 10.1. *If* $\sigma^* \in Eq(\Gamma)$, *then* $\sigma_i^t \equiv \sigma_i^* \in SGP(\Gamma^N)$, $i \in I$, $t = 1, \ldots, N$, *and* $\sigma_i^t \equiv \sigma_i^* \in SGP(\Gamma_\delta^\infty)$, $i \in I$, $t = 1, 2, \ldots$.

Since every SGP is an equilibrium and $Eq(\Gamma) \neq \emptyset$, immediate corollaries are

$$\emptyset \neq u(Eq(\Gamma)) \subset U^N(SGP(\Gamma^N)) \subset U^N(Eq(\Gamma^N)), \text{ and}$$
$$\emptyset \neq u(Eq(\Gamma)) \subset U_\delta^\infty(SGP(\Gamma_\delta^\infty)) \subset U_\delta^\infty(Eq(\Gamma_\delta^\infty)).$$

In this sense, we've "rigged the game," all that can happen is increase in the set of equilibria when the game is repeated.

2. Prisoners' Dilemma finitely and infinitely

To get a flavor of what will happen in this section, we will look at repeating the Prisoners' Dilemma game Γ from above.

	Squeal	Silent
Squeal	$(-B+r, -B+r)$	$(-b+r, -B)$
Silent	$(-B, -b+r)$	$(-b, -b)$

Problem 10.2. *Show that* $\mathbb{O}(Eq(\Gamma^N))$ *contains only one point when* $N < \infty$. *Show that* $SGP(\Gamma^N)$ *contains only one point when* $N < \infty$.

One way to work the next problem uses "Grim Trigger Strategies," that is, $\sigma^1 =$ (Silent, Silent), and for $t \geq 2$, $\sigma_i^t(h^{t-1}) =$ Silent if h^{t-1} is all Silent, and is equal to Squeal for all other h^{t-1}.

Problem 10.3. *Show that there exists a* $\underline{\delta} < 1$ *such that for all* $\delta \in (\underline{\delta}, 1)$, $\mathbb{O}(SGP(\Gamma_\delta^\infty))$ *contains the history in which each player plays* Silent *in each period.*

The grim trigger strategies are a special case of what are called "Nash reversion" strategies — pick a Nash equilibrium τ for Γ and a vector $s \in S$. Nash reversion strategies are $\sigma^1 = s$, and for $t \geq 2$,

$$\sigma^t(h^{t-1}) = \begin{cases} s & \text{if } h^{t-1} = (s, s, \ldots, s) \\ \tau & \text{otherwise} \end{cases}.$$

In this game, the only $\tau \in Eq(\Gamma)$ is (Squeal, Squeal).

Summarizing what we have seen in the repeated prisoners' dilemma so far, for all $N < \infty$ and for all δ sufficiently close to 1,

$$u(Eq(\Gamma)) = U^N(Eq(\Gamma^N)) = U^N(SGP(\Gamma^N)) \subsetneq U_\delta^\infty(SGP(\Gamma_\delta^\infty)).$$

What is a bit puzzling is the question, "Why is there so large a distinction between Γ^N for large N and Γ_δ^∞ for δ close to 1?" This is puzzling because both types of games are supposed to be capturing interactions that are repeated many many times.

Roy Radner had a solution for the puzzle in the case of the repeated Prisoners' Dilemma. His solution was later (considerably) generalized by Fudenberg and Levine. Both papers worked with Herbert Simon's satisficing. Radner worked with the definition of satisficing having to do with solving optimization problems to within some $\epsilon > 0$ of the optimum achievable utility. Fudenberg and Levine replaced complicated optimization problems by simpler ones. This is Simon's other notion of satisficing. Fudenberg and Levine then showed that, in games more general than the repeated games we're looking at, this gives Radner's version of satisficing.

Definition (Radner): For a game $\gamma = (T_i, v_i)_{i \in I}$ and an $\epsilon \geq 0$, a strategy vector σ is an **ϵ-equilibrium** if

$$(\forall i \in I)(\forall t_i \in T_i)[u_i(\sigma) \geq u_i(\sigma \backslash t_i) - \epsilon].$$

If $\epsilon = 0$, an ϵ-equilibrium is an equilibrium. One can (and you should) write down the definition of an ϵ-SGP.

Radner showed that for every $\epsilon > 0$, there exists an N' such that for all $N \geq N'$, there exists strategies $\sigma \in SGP(\Gamma^N)$ with the property that $\mathbb{O}(\sigma)$ involves (Silent, Silent) at all time periods. One part of Fudenberg and Levine's work considered a subclass of strategies for a repeated game Γ^∞_δ. The subclass consisted of strategies of the form "stop thinking about what to do after N periods." They showed that the set of limits of ϵ-SGP in these strategies, limits being taken as $\epsilon \to 0$ and $N \to \infty$, give exactly the SGP of Γ^∞_δ. Further, equilibria within these subclasses are ϵ-SGP's.

In any case, using either logic, and variants the trigger strategies discussed above it is possible to show that, for Γ being the Prisoners' Dilemma,

Theorem 10.2. *If $v > u(\text{Squeal, Squeal})$ and $v \in \mathbf{co}\,(u(S))$, then for all $\epsilon > 0$,*

(1) exists N' such that for all $N \geq N'$, $B(v, \epsilon) \cap U^N(SGP^\epsilon(\Gamma^N)) \neq \emptyset$, and
(2) there exists a $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$, $B(v, \epsilon) \cap U^\infty_\delta(SGP(\Gamma^N)) \neq \emptyset$.

So, those are the major patterns for the results, the equilibrium sets expand as N grows large or as $\delta \uparrow 1$, and the utilities of the approximate equilibrium set for Γ^N, N large, look like the utilities for Γ^∞_δ, $\delta \simeq 1$.

Radner offered a couple of rationales for ϵ-equilibria with $\epsilon > 0$, and it worth the time to reiterate his points.

(1) First, actually optimizing is quite hard and we might reasonably suppose that people only optimize their payoffs to within some small ϵ. Herbert Simon distinguished between approximately optimizing in this sense, and in exactly optimizing a simplified version of the problem. In this game, we might consider looking for equilibria in some more limited class of strategies, e.g. in a class that contains the one called "tit-for-tat" — start by playing cooperatively and otherwise match your opponent's last play, and other simpleminded strategies. (I told you game theory often captured the subtle dynamics of a kindergarden classroom.) This second approach has been modeled extensively by assuming that players are limited in a number of aspects, most prominently by Abreu and Rubinstein who assumed that they are "finite automata".
(2) Second, and more intriguingly, Radner argued that we might believe that the players in this game understand that if both optimize all the way, if both squeeze out the last little bit of surplus, then a disaster will befall them both. It seems to me that this is an argument about socialization towards Pareto dominant outcomes — for the common good I am willing to sacrifice a little bit. This seems reasonable, but one might argue that if this is true, then it ought to show up in the payoffs. A really rigid version of this argument would say that if it is hard to put this into the payoffs convincingly (it is) then this indicates that we don't really understand what's behind this argument.

> One way to understand what is going on is to think about the advantages of having a reputation for being a good citizen. In other words, any game should be thought of as being imbedded in a larger social context.

3. Some results on finite repetition

Now let Γ be the tragedy of the Commons game as given (a long ways) above. The two claims we will look at are

Claim: If $h \in \mathbb{O}(Eq(\Gamma^N))$, then at the N'th time period, h specifies the static Nash equilibrium.

This is a special case of

Lemma 10.3. *In the last period of equilibrium play of Γ^N, a static Nash equilibrium is played.*

This implies that complete efficiency can<u>not</u> be achieved as part of a Nash equilibrium of Γ^N, though one could get the efficient outcome in periods $t = 1, 2, \ldots, N-1$, which is pretty close. One could have all countries (players) playing the strategy

> start by playing $1/I$ times the efficient fleet size, continue playing this so long as everyone else does and $t \leq N-1$, if at $t \leq N-1$, some country i has deviated from $1/I$-times the efficient fleet size at any point in the past, then identify the first deviator breaking ties by alphabetizing (in Swahili), have everyone who is not the first deviator play $1/(I-1)$ times the total fleet size that yields 0 profits at $t+1$, have the deviator play fleet size of 0, and at $t = N$, if there has been no deviation from $1/I$ times the efficient fleet size, play the static game Nash equilibrium fleet size.

Problem 10.4. *Give conditions on the payoffs U^N for which the strategies just described belong to $Eq(\Gamma^N)$.*

Something peculiar is happening, the Tragedy of the Commons and the Prisoners' Dilemma, played once, have just one equilibrium, however, when played $N < \infty$ times, the Tragedy has many equilibria, the Dilemma just one. It is possible to explain this using the language of "threats."

4. Threats in finitely repeated games

Put equilibrium considerations in the back of your mind for just a little bit, and cast your memory back (fondly?) to those school year days spent in terror of threats by older, larger kids. Consider the question, "What is the worst that players $j \neq i$ can do to player i?" Well, they can call him/her a booby or a nincompoop, but these are not sticks nor stones, and given that we are now adults, we are not supposed to believe that this hurts too much. However, stealing your lunch money and stuffing you into a garbage can, now that hurts. What $j \neq i$ can do is get together and agree to take those actions that make i so miserable as possible. This is a threat with some teeth to it. Now, i has some protection against this behavior — knowing that the others are ganging up, i can plan accordingly to maximize i's utility against the gang-up behavior. There are three "safe" utility levels that one might imagine i being able to guarantee $i$self,

$$\underline{v}_i^{pure} = \min_{a_{-i}\in\times_{j\neq i}A_j} \max_{t_i\in A_i} u_i(a_i, s_{-i}),$$
$$\underline{v}_i^{mixed} = \min_{\sigma_{-i}\in\times_{j\neq i}\Delta(A_j)} \max_{t_i\in A_i} u_i(a_i, \sigma_{-i}), \text{ and}$$
$$\underline{v}_i^{corr} = \min_{\mu_{-i}\in\Delta(\times_{j\neq i}A_j)} \max_{t_i\in A_i} u_i(a_i, \mu_{-i}),$$

Since $\times_{j\neq i}A_j \subset \times_{j\neq i}\Delta(A_j) \subset \Delta(\times_{j\neq i}A_j)$,

$$\underline{v}_i^{pure} \geq \underline{v}_i^{mixed} \geq \underline{v}_i^{corr}.$$

Problem 10.5. *Give games where the two inequalities are strict.*

The first corresponds of the worst that dolts who do not understand randomization can do to i, the second corresponds of the worst that enemies who do understand independent randomization can do to i, the third corresponds of the worst that fiends who completely understand randomization can do to i. The three $\underline{v}_i$'s are called "safety levels." Here is one of the reasons.

Lemma 10.4. *For all $i \in I$ and for all N (for all δ), if σ is an equilibrium for Γ^N (for Γ^∞_δ), then $U_i^N(\sigma) \geq \underline{v}_i^{mixed}$ ($U^\infty_{\delta,i}(\sigma) \geq \underline{v}_i^{mixed}$).*

This lemma is ridiculously easy to prove once you see how. Suppose that other players are playing some strategy σ_{-i}. In period 1, have i play a myopic, that is, one period best response to the distribution over A_{-i} induced by σ^1_{-i}, $\sigma_i \in Br_i(\sigma^1(h^0))$. More generally, after any h^{t-1}, have i play $\sigma_i^t \in Br_i(\sigma(h^{t-1})$. In each period, it must be the case that $u_i(s^t) \geq \underline{v}_i^{mixed}$.

The following is a pair of reasons to call the $\underline{v}_i$'s safety levels, neither proof is particularly easy.

Theorem (Benoit & Krishna): Suppose that for each $i \in I$ there is a pair of equilibria $\sigma^*(i)$ and $\sigma'(i)$ for Γ such that $u_i(\sigma^*(i)) > u_i(\sigma'(i))$. Suppose also that the convex hull of $u(S)$ has non-empty interior. Let v be a vector in the convex hull of $u(S)$ such that for all $i \in I$, $v_i > \underline{v}_i^{pure}$. Then for all $\epsilon > 0$, there exists an N' such that for all $N \geq N'$, there is a subgame perfect Nash equilibrium σ^* of Γ^N such that $\|u(\sigma^*) - v\| < \epsilon$. If the words "subgame perfect" are deleted, then change $\underline{v}_i^{pure}$ to $\underline{v}_i^{mixed}$.

It is intuitive for two reasons, one obvious and one a bit more subtle, that more things are possible when we look at equilibria rather than subgame perfect equilibria. First, there are more equilibria than there are subgame perfect equilibria, this is obvious. Second, some of the strategies that go into the proof require players to min-max someone else, and this can be rather costly. In an equilibrium, one can threaten to min-max someone and never have to seriously consider carrying through on it. But for an equilibrium to be subgame perfect, it must only consider min-max threats that are seriously considered as possibilities. Let us look at both these points in the following 2×2 game,

	L	R
T	(2,9)	(-20,-80)
B	(10,0)	(-30,-100)

For this game, $Eq(\Gamma) = (B, L)$ and $u(Eq(\Gamma)) = (10, 0)$.

Claim: $\mathbb{O}(Eq(\Gamma^2))$ contains the history $h = ((T, L), (B, L))$.

It is important to note that (T, L) is nothing like the unique equilibrium of Γ. The claim can be seen to be true by considering the strategies $\sigma^1(h^0) = (T, L)$, $\sigma_1^2(h^1) \equiv B$, and

$$\sigma_2^2(h^1) = \begin{cases} L & \text{if } h^1 = (T, L) \\ R & \text{otherwise} \end{cases} .$$

These are in fact Nash equilibria, just check the mutual best response property. They are not subgame perfect equilibria, just check that they call for play of a dominated strategy in the case of "otherwise." That is the obvious reasoning.

Show that the $\underline{v}_i^{pure}$ are $(-20, 0)$ for this game. The more subtle observation is that for 2 to min-max 1, 2 must suffer a great deal. To have a subgame perfect equilibrium in which 1's utility is held down, we must have strategies in which it regularly happens that some s_i^t giving 2 $u_2(s^t) < \underline{v}_2^{pure}$ happens. Therefore, the strategies for the Benoit and Krishna result must also threaten the threateners. In subgame perfect equilibrium strategies, 2 must be threatened with dire consequences, and it must be an equilibrium threat, after s/he has avoided receiving a period's worth of -80 or -100. In particular, s/he must be threatened by something even worse that what s/he was getting by going along with the strategies. In equilibrium strategies, 2 must be threatened with dire consequences, but it needn't be an equilibrium threat.

Claim: Let v be a vector in the convex hull of $u(S)$, $\mathbf{co}\,(u(S))$. If $v \gg (-20, 0)$, then for any $\epsilon > 0$, there exists $\underline{\delta} < 1$ such that if for all $i \in I$, $\delta_i \in (\underline{\delta}, 1)$, then

$$(\exists v' \in u(\mathbb{O}(Eq(\Gamma^\infty(\delta))))[\|v' - v\| < \epsilon].$$

5. Threats in infinitely repeated games

The third reason to call the $\underline{v}_i$'s safety levels appears in the following result, which we will not prove, though we will talk about it.[1]

Folk Theorem: Suppose that $\mathbf{co}\,(u(S))$ has non-empty interior. Let v be a vector in $\mathbf{co}\,(u(S))$ such that for all $i \in I$, $v_i > \underline{v}_i^{mixed}$. For all $\epsilon > 0$, there exists a $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$,

$$B(v, \epsilon) \cap U^\delta(SGP(\Gamma_\delta^\infty)) \neq \emptyset.$$

Before discussing how the proof works, let look at an example violating the condition that $\mathbf{co}\,(u(S))$ have non-empty interior, in particular, let us look at the Matching Pennies game,

	H	T
H	$(+1, -1)$	$(-1, +1)$
T	$(-1, +1)$	$(+1, -1)$

Claim: For all $\delta \in (0, 1)^I$, if $\sigma^\infty \in Eq(\Gamma^\infty(\delta))$, then $\mathbb{O}(\sigma^\infty)$ is the i.i.d. distribution putting mass $\frac{1}{4}$ on each point in S in each period. In this game, there is no $v \in \mathbf{co}\,(u(S))$ that is greater than the threat point vector $(0, 0)$.

Three person example violating the interiority condition goes here.

Now let us turn to a discussion of how the proof works, a discussion that gives a preview of the Simple Penal Codes below. Pick your v satisfying the assumptions of the Folk Theorem

[1] We will not even state the most general version of this result, for that see Lones Smith (19??, Econometrica).

above except that we want the extra condition that v be in the interior of $\mathbf{co}\,(u(S))$. Pick an arbitrary $\epsilon' > 0$. Pick a strictly positive $\epsilon'' \leq \epsilon$ such that $B(v, 2\epsilon'') \gg \underline{v}^{mixed}$ and $B(v, 2\epsilon'') \subset \mathbf{co}\,(u(S))$. Set $\epsilon = \min\{\epsilon', \epsilon''\}$. Suppose that we can fix some cyclical history $q^0 = (q_1^0, q_2^0, \ldots)$ such that $\lim_{\delta\uparrow 1} U^\delta(q^0) = v$. (We can surely find an h such that $\lim_{\delta\uparrow 1} U^\delta(q^0) \in B(v, \epsilon/3)$, but let's pretend we nailed v exactly.) Now, for each $i \in I$, fix a history q^i such that $\lim_{\delta\uparrow 1} U_i^\delta(q^i) = v_i - \epsilon$, and $\lim_{\delta\uparrow 1} U_j^\delta(q^i) = v_j + \epsilon$, $j \neq i$. (Again, we can get within $\epsilon/3$ for sure, let us pretend we have nailed this vector exactly.) Let us further arrange the q^i so that i's utility for the first periods is their min-max payoff until they regret having deviated and then it goes to a cycle c^i satisfying

$$\lim_{\delta\uparrow 1} U^\delta(c^i) = v + \epsilon\tilde{1}.$$

returns to the q^0 cycle. Consider the strategies of roughly the form

> Start out playing the history q^0; so long as nobody deviates, continue to play q^0; if i deviates from q^0 at time t, start playing q^i from time $t+1$ onwards; if we've started playing q^i, $i \neq 0$, and j deviates from q^i at some point in time, start playing q^j from that period on; if two or more people have deviated, then pick the lowest numbered i among the deviators, and start playing q^i from the next period onwards.

These are only rough specifications of strategies since they are a little bit unclear about what happens after many people deviate. We'll clean this up later. The important point to note is the following, if $\delta \simeq 1$, then everyone wants to go along with q^0, i.e. we have an equilibrium. If i doesn't go along, then i and everyone will start in on playing q^i. The history q^i is front loaded with vectors that i does not like, and since $\delta \simeq 1$, i would prefer to go along with q^0 than to get q^i. Now, would anyone want to deviate from q^i? Well, any $j \neq i$ would not, they are trading $v_j + \epsilon$ against $v_j - \epsilon$ if they do. Would i ever want to deviate from q^i? If i did deviate, it would merely prolong the amount of time everyone else is min-max'ing i.

Problem 10.6. *This question concerns infinitely repeated Cournot competition between two firms with identical, constant marginal costs, $c > 0$, identical discount factors, $0 < \delta < 1$, and a linear demand curve with intercept greater than c.*

1. *For what range of δ's can* **monopoly** *output be a subgame perfect equilibrium with Nash reversion strategies?*
2. *Show that as $\delta \uparrow 1$, prices arbitrarily close to the* **competitive** *price c also arises as part of the equilibrium price path.*

6. Dynamic programming and the value function

The basic idea is to start at the end of a problem and work backwards towards the beginning of the problem. The device that allows us to accomplish this is the value function.

7. Examples

As is mostly true, we'll start with examples and build the theory from them.

7.1. NIM. M units of money are put on the table in front of you. You and your opponent take turns removing either 1 or 2 units of money. The person removing the last unit wins the entire pile of M units. Question: Do you want to go first, or second, or does it matter?

Solve this first for small M and then work backwards using the value of being at M and it being your turn or it being your opponent's turn.

7.2. Twain's travels. Mark Twain's problem is to get from node O to node z in a least cost fashion. This is formulated using a set of nodes

$$\mathcal{N} = \{O, a, b, c, d, e, f, g, h, z\},$$

a set of directed paths,

$$\mathcal{P} = \{(O, a), (O, \quad b), (a, c), (a, d), (b, c), (b, d), (b, e), (c, f), (c, g), (d, f), (d, h), (e, g), (e, h), (f, z), (g, z), (h, z)\}$$

and a set of costs

(s,t)	$u(s,t)$	(s,t)	$u(s,t)$	(s,t)	$u(s,t)$	(s,t)	$u(s,t)$
(O,a)	19	(O,b)	23	(a,c)	11	(a,d)	6
(b,c)	6	(b,d)	19	(b,e)	3	(c,f)	16
(c,g)	23	(d,f)	13	(d,h)	29	(e,g)	3
(e,h)	22	(f,z)	12	(g,z)	21	(h,z)	2

We solve this problem by giving a tree representation of the the set of directed paths in the problem and then moving backwards through the tree. At each step we find the value to being at a particular node as the optimum of the one-stage rewards plus the value of where we end up.

7.3. Some formal stuff. Stages, nodes, links or branches, all captured by the related notions of trees and acyclical, directed graphs with origin. Nodes are a finite[2] set $\mathcal{N}$, branches, $\mathcal{P}$ are a subset of the ordered pairs of elements of $\mathcal{N}$.

A **branch of length** n **joining** s **to** t is a sequence $s = t_0, t_1, \ldots, t_n = t$ such that for $i = 0, 1, \ldots, n-1$, $(t_i, t_{i+1}) \in \mathcal{P}$. A branch joining s to t is denoted $s \rightsquigarrow t$. If there is a branch of length 1 joining s to t, then s is a parent of t. Define s to be an **ancestor of** t if there is a branch of some length joining s to t.

The rules for **trees** are 1, 2, 3, and 4.a, while the rules for **acyclical, directed graphs with origin** are 1, 2, 3, and 4.b.

1. For each $s \in \mathcal{N}$, $(s, s) \notin \mathcal{P}$.
2. There is exactly one node, call it O (for origin), with the property that there is no s such that $(s, O) \in \mathcal{P}$, that is, O is the only node without a parent, the primal node if you will.
3. For each $s \in \mathcal{N}$, if $s \neq O$, then O is an ancestor of s. In particular, each $s \neq O$ has a parent.

[2]Finiteness will disappear soon.

(4) The first of the next two gives us classical decision and game theoretic trees, the second is a weaker condition, and gives us acyclical, directed graphs with origin that are used for dynamic programming.
  (a) For each $t \in \mathcal{N}$, $t \neq O$, there is exactly one node $s \in \mathcal{N}$ such that $(s,t) \in \mathcal{P}$.
  (b) No node is its own ancestor.

Trees are the basic building blocks for extensive form games, so we will see them again.

PROBLEM 10.7. *Show that every tree is an acyclical, directed graphs with origin but that the reverse is not true.*

Some comments and definitions: (1) Let $C(s)$ denote the **children of the node** s, that is, the set of $t \in \mathcal{N}$ such that $(s,t) \in \mathcal{P}$. We can divide the nodes of the tree or acyclical, directed graph with origin into stages, the first stage, A_1 is the set of nodes $C(O)$, the second stage is $C(A_1)$, the $n+1$'th stage is $C(A_n)$. This can be very helpful when each node s is in a unique stage, and sometimes even when this is not true. (2) A **terminal node** is a node that has no children, and Z denotes the set of terminal nodes. (3) For a note t, let $S(t)$ denote the **successors of** t, that an element of $C(t)$ or of $C(C(t))$ or of .... Formally, $S(t) = \cup_n C^n(t)$ where C^n is the n-fold application of the mapping C from $\mathcal{N}$ to the set of subsets of $\mathcal{N}$. (4) In a tree, every terminal node z is connected to the origin by a unique branch. This is not true in an acyclical, directed graph with origin.

For dynamic programming, payoffs u depend on the branch taken from O to the terminal nodes. For each pair $(s,t) \in \mathcal{P}$, there is a payoff $u(s,t)$, and the payoff to a branch $O \rightsquigarrow z = (O = t_0, t_1, \ldots, t_n = z)$ is $\sum_{i=0}^{n-1} u(t_i, t_{i+1})$. More generally, define $u(s \rightsquigarrow z)$ to be the part of the payoffs accrued on the $s \rightsquigarrow z$ branch. Then the value function is defined by

$$V(t) := \text{opt}_{z \in S(t) \cap Z}\{u(t \rightsquigarrow z)\},$$

where "opt" can be either "max" or "min." Rather arbitrarily, we'll use "max." The basic recursion relation is known as Bellman's equation or Pontryagin's formula, it is

$$V(s) = \max{}_{t \in C(s)}\{u(s,t) + V(t)\}.$$

When the stage structure is an important part of the problem, this is sometimes expressed as

$$V_n(s) = \max{}_{t \in C(s)}\{u(s,t) + V_{n+1}(t)\}$$

where s is in stage n and t in stage $n+1$.

The basic device is to start at the terminal nodes z, figure out the value for all parents of terminal nodes that have nothing but terminal children, call these nodes T_1, then figure out the value for all parents of T_1 nodes that have nothing but T_1 children, etc. until you have worked all the way back through the tree.

Let $T_0 = Z$. Let

$$T_1 = \{s \in P(t_0) : t_0 \in T_0, \text{ and } C(s) \subset T_0\},$$

inductively

$$T_{n+1} = \{s \in P(t_n) : t_n \in T_n, \text{ and } C(s) \subset \cup_{0 \leq m \leq n} T_m\}.$$

PROBLEM 10.8. *Show that there exists an n^* such that $T_{n^*} = \{O\}$ and that for all $n < n^*$, $T_n \neq \emptyset$.*

For each $s \in T_1$,

$$V(s) = \text{opt}_{z \in C(s)}\{u(s, z)\},$$

for each $s \in T_n$, $n \geq 2$,

$$V(s) = \text{opt}_{t \in T_{n-1}}\{u(s, t) + V(t)\}.$$

Essentially, we work backwards for terminal nodes all the way to O. Furthermore, let

$$t^*(s) \in \arg\max_{t \in T_{n-1}}\{u(s, t) + V(t)\}.$$

Then $t^*(t^*(\cdots(t^*(O))\cdots))$ is an optimal path through the tree or acyclical, directed graph with origin.

Problem 10.9. *Work through the Twain problem as a maximization rather than a minimization problem.*

There are computational advantages to using the value function.

Problem 10.10. *Suppose you have to solve a dynamic programming problem of crossing the country by passing through each of* 11 *regions. There are* 11 *cities in each region. From each of the cities in a region, you may travel to any of the cities in the next region.*

(1) How many possible paths are there?
(2) If you must calculate the values of each terminal node by adding the payoffs to the branches, how many additions must you perform?
(3) How many additions must you perform to solve the problem using dynamic programming? What (roughly) is the ratio of the two numbers of additions?

Problem 10.11. *Any two problems from Gibbons' §2.1 problems (p. 130-132).*

Problem 10.12. *Any two problems from Gibbons' §2.2 problems (p. 133-134).*

Problem 10.13. *Any two problems from Fudenberg and Tiroles' 3.3, 3.6, 3.7, or 3.8 (p. 100-102).*

Problem 10.14. *The professor in a game theory class offers to auction a ten-dollar bill. There are only two students allowed to bid. Ann gets to bid first. After Ann's bid, Bob can either quit or raise Ann's bid. If he raises her, she can either quit or raise his bid in turn. The game continues in this fashion until one player quits. When a player quits, the other player wins and receives the ten-dollar bill. However, and this is crucial,* **both** *players must pay their bids to the professor. Bids must be in multiples of a quarter. Find the backwards induction equilibria of this game, being sure to explicitly make any further assumptions you need.*

7.4. Investment under uncertainty. When actions taken in one stage result in distributions over the state in the next period, the Bellman equation changes to reflect this.

Here is a classical problem: Initial capital is $A_0 > 0$. At the beginning of each period $t \in \{0, 1, \ldots\}$, A_t must be allocated between consumption, C_t, investment, I_t, that will surely return $r \cdot I_t$, $r > 1$, in one period, and investment, J_t, in a risky investment that will return $Y_t \cdot J_t$ in one period, where the $Y_t \geq 0$ are i.i.d. random variables. Felicity is derived from consumption in any given period according the increasing function $c \mapsto u(c)$, and the objective is to maximize the expected discounted value of $\sum_{t=0}^{\infty} \delta^t u(C_t)$, $0 < \delta < 1$. We are going to discuss this problem using some results that we will prove later. We will

(1) Set this up as a stochastic dynamic programming problem and give the Bellman equation.
(2) Show that the value function, $V(s)$ is concave if $u(s)$ is concave using iterative approximations to the value function.
(3) Assume that $u(c) = c^r$ for some $r \in (0, 1)$, and show that the optimal policy allocates a constant proportion of A_t to each of the three alternatives and that the value function $V(s)$ is a multiple of $u(s)$.
(4) Show that if $u(\cdot)$ is concave and $EY < r$ then $J_t \equiv 0$.

8. Rubinstein-Ståhl bargaining

Two people, 1 and 2, are bargaining about the division of a cake of size 1. They bargain by taking turns, one turn per period. If it is i's turn to make an offer, she does so at the beginning of the period. The offer is α where α is the share of the cake to 1 and $(1 - \alpha)$ is the share to 2. After an offer α is made, it may be accepted or rejected in that period. If accepted, the cake is divided forthwith. If it rejected, the cake shrinks to δ times its size at the beginning of the period, and it becomes the next period. In the next period it is j's turn to make an offer. Things continue in this vein either until some final period T, or else indefinitely.

Suppose that person 1 gets to make the final offer. Find the unique subgame perfect equilibrium. Suppose that 2 is going to make the next to last offer, find the unique subgame perfect equilibrium. Suppose that 1 is going to make the next to next last offer, find the subgame perfect equilibrium. Note the contraction mapping aspect and find the unique solution for the infinite length game in which 1 makes the first offer.

PROBLEM 10.15. *The Joker and the Penguin have stolen* 3 *diamond eggs from the Gotham museum. If an egg is divided, it loses all value. The Joker and the Penguin split the eggs by making alternating offers, if an offer is refused, the refuser gets to make the next offer. Each offer and refusal or acceptance uses up* 2 *minutes. During each such* 2 *minute period, there is an independent, probability* r, $r \in (0, 1)$, *event. The event is Batman swooping in to rescue the eggs, leaving the two arch-villains with no eggs (eggsept the egg on their faces, what a yolk). However, if the villains agree on a division before Batman finds them, they escape and enjoy their ill-gotten gains.*

Question: What does the set of subgame perfect equilibria look like? [Hint: it is not in your interest to simply give the Rubinstein bargaining model answer. That model assumed that what was being divided was continuously divisible.]

9. Optimal simple penal codes

Here we are going to examine the structure of the subgame perfect equilibria of infinitely repeated games.

10. Abreu's example

The following is Abreu's (1988 E'trica) original example of using simple penal codes. It is a 3×3 game, a two player simultaneous discrete, quantity-setting duopoly game Γ in which each firm may choose a Low, Medium, or High output level:

	L	M	H
L	$10, 10$	$3, 15$	$0, 7$
M	$15, 3$	$7, 7$	$-4, 5$
H	$7, 0$	$5, -4$	$-15, -15$

Consider the Nash reversion simple penal code, $F(q^0, (q^1, q^2))$ with q^0 being $((L, L), (L, L), \ldots)$ and $q^1 = q^2$ being $((M, M), (M, M), \ldots)$. For discount factor $\delta \geq 5/8$, the Nash reversion strategies are subgame perfect, and for $\delta < 5/8$, they are not. Now consider the simple penal code $F(q^0, (q^1, q^2))$ with q^0 being $((L, L), (L, L), \ldots)$, q^1 being $((M, H), (L, M), (L, M), \ldots)$ and q^2 being $((H, M), (M, L), (M, L), \ldots)$. For $\delta = 4/7 < 5/8$, this vector of strategies is subgame perfect.

11. Harris' formulation of optimal simple penal codes

Time starts at $t = 1$, at each stage the simultaneous game $\Gamma = (A_i, u_i)_{i \in I}$ is played, A_i is sequentially compact with metric ρ_i, and u_i is jointly continuous with respect to any of metrics ρ inducing the product topology. For a history (read vector) $h \in H = \times_{t=1}^{\infty} S$, h_t denotes the t'th component of h, payoffs are given by

$$U_i(h) = \sum_{t=1}^{\infty} \delta_i^t u_i(h_t),$$

where $0 < \delta_i < 1$. The product topology of H can be metrized by

$$d(h, h') = \sum_{t=1}^{\infty} \min\{2^{-n}, \rho(h_t, h'_t)\}.$$

Problem 10.16. *The set H is sequentially compact in the metric d, and each U_i is continuous.*

Player i's strategy for period t is a function $\sigma_i^t : H^{t-1} \to A_i$ where $H^{t-1} := \times_{k=1}^{t-1} S$, and H^0 is, by convention, some one point set. An initial history is a vector $h^t \in H^t$. Player i's strategy is the vector $\sigma_i = (\sigma_i^1, \sigma_i^2, \ldots)$. A profile of strategies for the players is then $\sigma = (\sigma^i)_{i \in I}$, and a profile of strategies at time t is $\sigma^t = (\sigma_i^t)_{i \in I}$.

Let $\mathbb{O}(\sigma, h, t)$ denote the outcome h for the first t time periods followed by the outcome determined by play of σ. Thus, $(\mathbb{O}(\sigma, h, t))_k = h_k$ for $1 \leq k \leq t$, and

$$(\mathbb{O}(\sigma, h, t))_k = \sigma^k((\mathbb{O}(\sigma, h, t))_1, (\mathbb{O}(\sigma, h, t))_2, \ldots, (\mathbb{O}(\sigma, h, t))_{k-1})$$

for $k > t$.

Definition: A strategy combination σ is a **subgame perfect equilibrium of the repeated game** if $(\forall i, t, h)$ and for all strategies γ^i for i,

$$U_i((\mathbb{O}(\sigma, h, t))) \geq U_i((\mathbb{O}(\sigma \backslash \gamma^i, h, t))).$$

The assumption is that the repeated game has a subgame perfect equilibrium in pure strategies.

Definition: For a strategy vector σ, the history $q \in H$ **comes into force in period** $t+1$ **after a given initial history** $h^t = (h_1, \ldots, h_t)$ of $\mathbb{O}(\sigma, h, t) = (h_1, \ldots, h_t, q_1, q_2, \ldots)$.

Given histories in H, q^0 and $(q^i)_{i \in I}$, the following recursive construction of a strategy vector $F(q^0, q^1, \ldots, q^I)$ will be used many times below:

(1) q^0 comes into force in period 1.
(2) If q^j, $0 \leq j \leq I$ came into force in period k, if q^j isfollowed in all periods up to but not including period $t \geq k$, and if player i deviates againts q^j_{t-k+1} in period t, then q^i comes into force in period $t+1$.
(3) If q^j came into force in period k, and more than 1 player deviated against q^j_{t-k+1} in period $t \geq k$, then q^i comes into force in period $t+1$ where i is the lowest numbered amongst the deviating players.

Definition: A **simple penal code** is a vector of histories $(q^i)_{i \in I}$. A simple penal code is **perfect** if there exists $q^0 \in H$ such that $F(q^0, q^1, \ldots, q^I)$ is a subgame perfect equilibrium.

Problem 10.17. *If $(q^i)_{i \in I}$ is perfect, then $(\forall i \in I)[F(q^i, q^1, \ldots, q^I)]$ is a subgame perfect equilibrium.*

Let $P \subset H$ denote the set of outcomes associated with subgame perfect equilibria. Let $\underline{U}_i$ denote $\inf\{U_i(q) : q \in P\}$.

Definition: A simple penal code $(q^i)_{i \in I}$ is **optimal** if it is perfect and if

$$(\forall i \in I)[U_i(q^i) = \underline{U}_i].$$

Lemma: Let $f = F(q^0, q^1, \ldots, q^I)$. If no single period deviation against f is profitable for any player in any subgame, then f is a subgame perfect equilibrium.

Problem 10.18. *Show that $(q^i)_{i \in I}$ is a perfect simple penal code if and only if for all i, j, t, and all $a_i \in A_i$,*

$$\sum_{k=t}^{\infty} \delta_i^k u_i(q^j_k) \geq \delta_i^t u_i(q^j_t \backslash a_i) + \delta_i^t U_i(q^i).$$

Proposition: There exists an optimal simple penal code.

Proposition: If $(q^i)_{i \in I}$ is an optimal simple penal code, then

$$(\forall q^0 \in P)[F(q^0, q^1, \ldots, q^I) \text{ is a subgame perfect equilibrium}].$$

12. "Shunning," market-place racism, and other examples

Again, we will work through a problem.

PROBLEM 10.19. *There are N identical, incumbent firms Bertrand (price) competing in an industry. The industry (inverse) demand curve is $p = 1 - q$. Each incumbent firm has the constant marginal cost function $C(q_i) = cq_i$ for some $0 < c \ll 1$. The consumers react instantly to price changes, going to the lowest cost supplier, evenly dividing themselves among firms tied at the lowest price. For the repeated versions of this game, the firms have the same discount factor, $\beta \in (0, 1)$.*

(1) Find the one-shot Bertrand-Nash equilibrium prices, quantities, and profits for this industry.

(2) (Nash reversion) As a function of β, find the highest price (i.e. most collusive) q^0 such that $F(q^0, q^1, \ldots, q^I)$ is a subgame perfect equilibrium where $(q^1, \ldots, q^I)$ is the Nash equilibrium.

(3) (Temporary Nash reversion) Let p_m denote the price that a monopolist would charge in this industry and let $\pi_m > 0$ denote industry monopoly profits. For each $i \in I$, let q^i be the history

$$q^i = (\underbrace{\vec{c}, \ldots, \vec{c}}_{T \text{ periods}}, \vec{p_m}, \vec{p_m}, \ldots)$$

where $\vec{x}$ is an I-length vector with x in each component. Let q^0 be the history

$$q^0 = (p_m, p_m, \ldots).$$

Show that for β close to 1 and T large, $F(q^0, q^1, \ldots, q^I)$ is a subgame perfect equilibrium. (Verbally, these strategies are "play c for T periods if there is ever a price gouger.")

(4) (Discriminatory equilibria) Now suppose that any of the firms could lower their costs to $C(q_i) = rcq_i$ with $0 < r < 1$ by hiring, at a lower wage, members of an oppressed racial, ethnic, gender, or religious group. Let $h_{i,t}$ denote firm i's hiring practices in period t, with $h_{i,t} = 0$ indicating no hires from the oppressed group and $h_{i,t} = 1$ denoting hires from the oppressed group. Actions by i in period t are $(h_{i,t}, p_{i,t})$. Let q^0 denote the history

$$q^0 = ((0, \vec{p}_m), (0, \vec{p}_m), \ldots).$$

Let a^i denote the length-I vector with i'th component $(0, sc)$, $0 \le s < r$, and all other components $(0, rc)$. Let q^i denote the history

$$q^i = (\underbrace{a^i, \ldots, a^i}_{T \text{ periods}}, q^0).$$

Find conditions on s, T, and β such that $F(q^0, q^1, \ldots, q^I)$ is a subgame perfect equilibrium. Verbally, if someone either cheats on the cartel price or hires a member of the oppressed group, they must do penance for T periods before anyone can go back to the profitable discriminatory status quo.

(5) *Subgame perfect equilibria must survive single person deviations. This is definition and can therefore be changed. One can sometimes find equilibria that survive deviations by groups. Formulate discriminatory equilibria that survive coordinated deviations by "small" subsets of the firms in this game.*

Bibliography

1. Robert M. Anderson and William R. Zame, *Genericity with infinitely many parameters*, Adv. Theor. Econ. **1** (2001), Art. 1, 64 pp. (electronic). MR MR2002579 (2004g:91038)
2. David Blackwell, *Comparison of experiments*, Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, 1950 (Berkeley and Los Angeles), University of California Press, 1951, pp. 93–102. MR MR0046002 (13,667f)
3. ———, *Equivalent comparisons of experiments*, Ann. Math. Statistics **24** (1953), 265–272. MR MR0056251 (15,47b)
4. Tilman Börgers, *Weak dominance and approximate common knowledge*, Journal of Economic Theory **64** (1994), no. 1, 265–277.
5. Christopher Chambers and Federico Echenique, *Supermodularity and preferences*, 2007.
6. John Conlisk, *Costly predation and the distribution of competence*, American Economic Review **91** (2001), no. 3, 475–484.
7. Eddie Dekel and Drew Fudenberg, *Rational behavior with payoff uncertainty*, Journal of Economic Theory **52** (1990), no. 2, 243–268.
8. Drew Fudenberg and Jean Tirole, *Game theory*, MIT Press, Cambridge, MA, 1991. MR MR1124618 (92m:90001)
9. Sergiu Hart and Andreu Mas-Colell, *Uncoupled dynamics do not lead to nash equilibrium*, American Economic Review **93** (2003), no. 5, 1830–1836.
10. Aviad Heifetz, Chris Shannon, and Yossi Spiegel, *What to maximize if you must*, 2004.
11. Levent Koçkesen, Efe A. Ok, and Rajiv Sethi, *The strategic advantage of negatively interdependent preferences*, J. Econom. Theory **92** (2000), no. 2, 274–299. MR MR1765925 (2001f:91028)
12. Paul Milgrom and John Roberts, *Rationalizability, learning, and equilibrium in games with strategic complementarities*, Econometrica **58** (1990), no. 6, 1255–1277. MR MR1080810 (91j:90101)
13. Paul Milgrom and Chris Shannon, *Monotone comparative statics*, Econometrica **62** (1994), no. 1, 157–180. MR MR1258667 (95d:90022)
14. Russell Mokhiber, *Corporate crime and violence: big business power and the abuse of the public trust*, Sierra Club Books, San Francisco, 1988.
15. Alfred Müller, *Comparing risks with unbounded distributions*, J. Math. Econom. **30** (1998), no. 2, 229–239. MR MR1652641 (99m:90049)
16. Eric Rasmusen, *Games and information*, third ed., Blackwell Publishing, Malden, MA, 2001, An introduction to game theory. MR MR2185321 (2006g:91002)
17. Michael Rothschild and Joseph E. Stiglitz, *Increasing risk. I. a definition*, J. Econom. Theory **2** (1970), 225–243. MR MR0503565 (58 #20284a)
18. ———, *Increasing risk. II. Its economic consequences*, J. Econom. Theory **3** (1971), 66–84. MR MR0503567 (58 #20284c)
19. John Van Huyck, Raymond Battalio, and Richard Beil, *Tacit coordination games, strategic uncertainty, and coordination failure*, American Economic Review **80** (1990), no. 1, 234–248.