Notes on Formal Political Analysis 1

Nolan McCarty and Adam Meirowitz
ameirowi@princeton.edu

October 1, 2002

Contents

1 Prologue

To establish some perspective we begin with some academic history:

-In their 1943 treatise *Theory of Games and Economic Behavior*, two Princeton scholars, John von Neumann and Oscar Morgenstern began the discipline called *game theory* and invented the technology to analyze games involving two players with opposed preferences.

-In 1950 John Nash, a mathematics graduate student at Princeton, provided the tools to extend the theory of von Neumann and Morgenstern to n-person games with arbitrary preferences.[1]

-In 1951 Kenneth Arrow established that there is no means of aggregating individual preferences into social preferences satisfying a few normatively desirable properties. The implications of his argument for democratic theory are far reaching. Arrow's axiomatic approach elucidates a mathematical structure of preference aggregation that scholars dating back to at least the French revolution had only loosely understood. This work redirected the field of *social choice theory.*

-In 1957 what was to become the most widely cited book by living political scientists, Anthony Downs' *An Economic Theory of Democracy* was published. In his dissertation (advised by Ken Arrow), Downs initiated an important research program (later named Downsian competition) aimed at understanding electoral competition. Basic ideas like candidates "running to the middle" or "alienating the extremes" stem from this program.

-In the early 1960's William Riker began, at the University of Rochester's department of political science, the research program that now goes by the names positive political theory and formal analysis of politics. This program spread first to Carnegie Mellon University and then to California Institute of Technology, before establishing a foothold in most top universities. The Politics department at Princeton has faculty that were leaders in the development of this program as well as faculty that were students of leaders of this development. The Economics department at Princeton has some of the best game theorists in the world.

Today the terms "formal analysis" or "game theory" appear in 750 published articles in just the political science journals that are indexed by jstor. Likewise of the economics journals indexed by jstor a search of "electoral competition" or "Downs" yields over 800 hits. Game theoretic models regularly appear in the study of electoral politics, parliamentary politics, congressional politics, presidential politics, separation of powers, bureaucratic politics, judicial politics and crisis avoidance.

[1] Nash's contribution was considered only an acceptable dissertation.

2 Introduction

These notes will serve as a guide to POL575, Formal Political Analysis 1. While Jim Morrow's *Game Theory for Political Scientists* presents many of the topics covered in class, my presentation will depart from his on many topics. Accordingly, these notes will be made available despite the fact that they are in an early state of preparation. A first course in Formal Political Analysis faces a difficulty. On the one hand anything more than superficial coverage of the literature using formal methods in politics requires solid preparation in choice theory, and game theory. On the other hand studying choice theory and game theory without simultaneously considering the applicability of these tools may seem too costly an endeavor for the busy graduate student.

We compromise by studying choice and game theory while considering applications. Material is more-or-less organized as a class on theory and not topics. We begin with the basic theory of individual choice, consider important results from social choice theory, and then move to game theory. Along the way, models and results from the literature are introduced. In some cases simplified versions will be presented, and in others we will devote a fair amount of time to working through journal articles. The cost of this organization is that we will only consider a small subset of the literature using formal methods, and the applications are not organized in a convenient substantive manner (like international relations, elections, courts, the bureaucracy, congress, etc...).

A second unfortunate consequence of this organization is that the beginning of the course has less political science in it than the latter part of the class. The reason is simple, few scholars have been able to learn subtle things about politics with only the most basic concepts. However, the work that does this successfully may capture a very important phenomena. Accordingly since most work in formal analysis (certainly recent work) uses relatively sophisticated tools, it is more difficult to select interesting articles that use the material presented very early in the course. The one exception being that social choice theory is presented early in the class, and a very large and exciting literature exists on this topic. The shortage of applications is most evident in the early chapters on game theory. While these chapters present material that may dramatically reshape intuition about strategic interaction, the set of direct published applications is small.

The goal of these notes is to present the theoretical tools, definitions, and results on individual choice theory, social choice theory, and game theory in the order presented in class. A few notes on exposition are needed. A few sections that cover interesting but technical topics in a more straightforward manner than existing texts are included. These chapters are marked with a *.. Throughout problems or exercises are included. These questions are usually not difficult, but stopping to think may assist digestion of the material. Do not worry, I will not expect you to turn in exercises, unless they are explicitly assigned.

3 Mathematical prequisites

Throughout sets are defined in the following manner $\{a \in A : \text{something is true about } a\}$. This standard notation is to be read as the set of points a contained in set A for which "something is true about a". An example is $\{i \in \{1,2,3\} : i \text{ is even}\}$, which is just the set $\{2\}$. The notation $A^c = \{x : x \notin A\}$ denotes the complement of A. The notation $A \backslash B = \{x : x \in A \text{ and } x \notin B\}$denotes the set containing everything in A but not B. The notation $|A|$ is used to denote the cardinality (number of elements) of a finite set. By $A \cup B = \{x : x \in A \text{ or } x \in B\}$ we denote the union of A and B. By $A \cap B = \{x : x \in A \text{ and } x \in B\}$ we denote the intersection of A and B. The subser operation $\subset$ has the following meaning $A \subset B$ iff $x \in A$ implies ($\Longrightarrow$) $x \in B$. Two sets A and B are equivalent ($A = B$) if $A \subset B$ and $B \subset A$.

We will also use logical relations about conditions or assumptions. For instance the implication

$$\text{conditon A} \Rightarrow \text{conditon B}$$

is equivalent to $\{x : \text{condition A is true of } x\} \subset \{x : \text{condition B is true of } x\}$. Simmiliarly the if and only if (iff) statement

$$\text{conditon A} \Leftrightarrow \text{conditon B}$$

is equivalent to

$$\begin{aligned} \text{conditon A} &\Rightarrow \text{conditon B} \\ \text{conditon B} &\Rightarrow \text{conditon A.} \end{aligned}$$

Standard sets (spaces) from analysis will be used. By $\mathbb{R}$ we denote the real line. Given two sets X and Y the product set $X \times Y$ consists of vectors with first coordinate in X and second coordinate in Y. By $\mathbb{R}^n$ we denote $n-$dimensional Eculidean space. The set $\mathbb{Z}$ refers to the integers. Occasionally, definitions and topics from real analysis will be presented. The following texts are good starting places for additional reading on analysis: Gaughan 1993, *Introduction to Analysis 4ed.* Brooks/Cole; Komogorov and Fomin1970, *Introductory Real Analysis*, Dover.

4 Choice Theory

Much of formal political analysis deals with the behavior of collections of rational individuals that interact in a particular political institution. In this chapter we make explicit a weak notion of rationality.

4.1 Finite Choice Spaces

We consider an agent facing a **finite** set $X = \{x_1, ..., x_n\}$ of elements from which she is to make a choice. The set X contains distinct policies that the agent can select. We say the set if finite because there are only a finite number of elements in the set. An arbitrary element of this set x_i is a particular policy. The notation $x_i \in X$ means that x_i is an element of the set X. One example is the set of integers between 1 and 3. In this case $X = \{1, 2, 3\}$. Another example is $X = \{$send in the troops, try negotiating, do nothing$\}$.

An agent has preferences over the set X. The notation $x_i R x_j$ states that policy x_i is weakly preferred to policy x_j. We will now be very explicit about these preferences. The next two paragraphs may seem opaquely abstract, but one we have made the meaning of a binary relation explicit things will be clearer. A binary preference relation R on X is a list of pairs of policies in X that are ordered by weak preference. We will use the terms binary relation, and binary ordering, interchangeably. Returning to our example with $X = \{1, 2, 3\}$ one possible binary ordering is $R = \{(1, 2), (1, 1), (2, 2), (1, 3), (3, 3), (2, 3)\}$. The interpretation of this notation is that if the pair (x_i, x_j) is in R then $x_i R x_j$. The fact that $(1, 1)$ is included in R in this example may seem odd, but thinking about R as the weak preference relation, this just means that the policy 1 is weakly preferred to itself (some people refer to this as indifference). In fact the ordering exhibited above is equivalent to the commonly used ordering $\leq$ on the set X.

Exercise 1 *Demonstrate that if $X = \{1, 2, 3\}$ then the ordering $\leq$ is represented as* $\{(1, 2), (1, 1), (2, 2), (1, 3), (3, 3), (2, 3)\}$.

In a slightly more complicated jargon, we can say that given a set X a binary ordering on X is a subset of the set $X \times X$. The set $X \times X$ is just the set of ordered pairs (x, y) where $x \in X$ and $y \in X$. In this sense if $(x, y) \in R$ than xRy. Rethinking the example of $\leq$ on the set X we see that the order matters as $1 \leq 2$ but it is not the case that $2 \leq 1$.

It is customary to start with a set X and then define a particular binary relation–weak preference– on X. Given the weak preference relation R on X we can define two other important relations: strict preference and indifference.

Definition 1 *Given any $x, y \in X$ we say xPy if and only if xRy and not yRx. We say xIy iff xRy and yRx.*

Accordingly P denotes strict preference and I denotes indifference. Returning to the example of $\leq$ on X, the strict preference relation derived from $\leq$ is equivalent to the relation $<$ and the indifference relation is equivalent to the relation $=$. Note that both P and I are also orderings on X.

Exercise 2 *To illustrate these concepts in the example with* $X = \{1, 2, 3\}$ *and* $R = \{(1, 2), (1, 1), (2, 2), (1, 3), (3, 3), (2, 3)\}$ *express* P *and* I *as a set of pairs. —Hint* $(1, 1)$ *is in* I *but not* P*.*

So far our mathematization of preferences doesn't satisfy any reasonable conditions. Our goal is to understand what conditions preference orderings need to satisfy so that given a set X there is an optimal choice for the agent. By optimal we mean an $x^* \in X$ for which x^*Ry for every $y \in X$ We offer the following definition

Definition 2 *Given a set* X *and binary relation* R *on* X *the maximal set* $M(R, X) \subset X$ *is defined as follows* $M(R, X) = \{x \in X : xRy \ \forall \ y \in X\}$

For now our **working definition of rationality is choice that is consistent with the selection of maximal policies induced by preferences.** This concept only makes sense if $M(R, X) \neq \emptyset$ (meaning $M(R, X)$ is not empty). We spend the rest of this chapter determining what conditions on X and R imply that $M(R, X) \neq \emptyset$. We now consider two conditions that will move us in this direction.

Definition 3 *A binary relation* R *on* X *is*
(i) complete if for all $x, y \in X$ *with* $x \neq y$ *either* xRy *or* yRx *or both.*
(ii) reflexive if for all $x \in X$ xRx*.*

The general approach is to take as a primitive a complete and reflexive weak preference relation (which is a binary relation). This then implies that the strict preference relation P is not reflexive, and the indifference relation I is reflexive. An alternative way of constructing these concepts is to define strict preference and indifference for a given set and then derive the weak preference relation from these two primitives.

Exercise 3 *Demonstrate that given any finite set* X *and two binary orderings* P *and* I *with* I *reflexive, and* P *not reflexive on any* $x \in X$ *there is a unique weak preference relation* R *on* X *for which (i)* xPy *iff* xRy *and* $\tilde{}yRx$ *and (ii)* xIy *iff* xRy *and* yRx*. Additionally show that* R *is reflexive. Finally show that if* I *is complete then* R *is complete.*

While completeness and reflexivity get us closer to a "rational" preference relation they are not sufficient. We need to rule out problems like xRy, yRZ and zRx. The problem with this state of affairs is that there is no reasonable choice for an agent with these preferences–Why choose y when you can choose x, why choose x when you can choose z, and why choose z when you can choose y. The following conditions represent axioms that resolve this problem.

Definition 4 *A binary relation R on X is*

(1) Transitive if for all $x, y, z \in X$ if xRy and yRz then xRz.

(2) Quasi-transitive if for all $x, y, z \in X$ if xPy and yPz then xPz.

(3) Acyclic if for all$\{x, y, z, ..., a, b\} \in X$ if xPy and yPz .. .and aPb then xRb

In most applied settings it is reasonable to expect preferences to satisfies transitivity. After this section we will assume that preferences are transitive.

Definition 5 *Given a set X a weak ordering is a binary relation that is complete, reflexive and transitive.*

It is not difficult to see that Transitivity rules out exactly the cycle considered above. In fact we can now state our first result.

Theorem 1 *If X is finite and R is a weak ordering then $M(R, X) \neq \emptyset$.*

> *Proof:* Assume that X is finite and R is complete, reflixive, and transitive. We establish the result by induction on the number of elements in X.
>
> **Step 1:** If X has 1 element (ie $X = \{x\}$), then by reflexivity xRx and thus $M(R, X) = \{x\}$.
>
> **Step 2:** We show that if it true that for any X' with n elements R' a weak ordering impies that $M(R, X') \neq \emptyset$ then for any X with $n + 1$ elements when R is a weak ordering on X, $M(R, X) \neq \emptyset$.
>
> **-Proof of step 2:** assume that for any X' with n elements R' a weak ordering on X' impies that $M(R', X') \neq \emptyset$. Now consider a set X with $n + 1$ elements. For arbitrary $x \in X$ it is true that $X = X' \cup \{x\}$ with X' a set having n elements. By assumption $M(R', X') \neq \emptyset$. So for an arbitrary $y \in M(R', X')$ either yRx or xRy or both by completeness.
>
> –If yRx then since $y \in M(R', X')$ we have yRz for all $z \in X' \cup \{x\}$ and thus $y \in M(R, X)$ and the step 2 result is established.
>
> –If it is not the case that yRx then we have xRy. Since $y \in M(R', X')$ we have yRz for any $z \in X'$. Thus for any $z \in X'$ we have xRy and yRz.Since R is transitive this implies that we have xRz for any $z \in X'$. This and xRy imply that for any $w \in X'$ we have xRw and thus $x \in M(R, X)$ and the step 2 result is established.
>
> **Step 3:** By steps 1 and 2 for any finite sized X if R is a weak order on X then $M(R, X) \neq \emptyset$.■

It turns out that a weak ordering is not needed for $M(R, X)$.

Theorem 2 *Assume X is finite and R is a complete and reflexive binary relation on $\dot{X}$. $M(R, S) \neq \emptyset$ on any $S \subset X$ (except $S = \emptyset$) iff R is acyclic.*

Exercise 4 *Prove this.*

Even with a finite choice space and no uncertainty the theory of choice is fairly rich. Austen-Smith and Banks (1999) is a good first source for students interested in going further. Many economists and psychologists, have been concerned about the assumption of completeness and a theory of choice without this condition has been derived.

4.2 Continuous Choice Spaces*

4.2.1 Non-emptyness of $M(R, X)$

Examination of the argument for claim 1 demonstrates that the fact that the choice space was finite was useful. When agents are to choose from a choice space that is like the continuum (e.g. the set of real numbers denoted $\mathbb{R}$ or the set $[0, 1] = \{x \in \mathbb{R} : x \geq 0 \text{ and } x \leq 1\}$) more structure on preferences is needed to insure that we have a model of choice in which $M(R, S) \neq \emptyset$. Two simple examples demonstrate where things can go awry.

Example 1 *Let $X = (0, 1)$ (or let $X = \mathbb{R}^1$ and let R be equivalent to $\geq$ so that xRy iff $x \geq y$ In this case the set $M(R, X)$ is empty. To see this note that if $x \in X$ then there exists some $y \in X$ for which $y > x$ and thus it is not the case that xRy.*

In the example the fact that $(0, 1)$ has no biggest element results in the emptiness of the maximal set.

Example 2 *Let $X = [0, 1]$ and define R as follows: xRy if $x, y \in [0, \frac{1}{2}]$ and $x \geq y$ or if $x, y \in (\frac{1}{2}, 1]$ and $x \leq y$ or if $x \in (\frac{1}{2}, 1]$ and $y \in [0, \frac{1}{2}]$. In this case $M(R, X)$ is empty.To see this note that nothing in $[0, \frac{1}{2}]$ can be in $M(R, X)$ but that on the set $(\frac{1}{2}, 1]$ a problem analogous to that of the previous example occurs.*

To diagnose the exact nature of this problem and determine reasonable axioms that avoid the problem we need a few mathematical concepts.[2] We begin with the assumption that the choice set is a subset of n-dimensional Euclidean space, $X \subset \mathbb{R}^n$. A point in such a space can be written as a vector $x = (x^1, x^2, ..., x^n)$ where each coordinate x^i is a point in $\mathbb{R}^1$. The logic of this section can be grasped with the simplest example of $X \subset R^1$. In this case a set $A \subset X$ is termed **open** if for every point $x \in A$ there is some $\varepsilon > 0$ s.t. for any $y \in X$ satisfying $|x - y| < \varepsilon$ it is the case that $y \in A$. The point is, a set is open if for any point in the set there is space around the point that is also in

[2]More precisely we need a few Topological concepts. Students interested in further study of choice theory would be well served by examining a text on Real Analysis. An approachable introductory text is: Gaughan, Edward. 1993. *Introduction to Analysis, 4ed.* Brooks/Cole Publishing Company. A more complete text is: Kolmogorov, A.N. and S.V. Fomin. 1970. *Inroductory Real Analysis.* Dover.

the set. To define the concept of openness in n-dimensional Euclidean space we use the norm

$$\|x-y\| = \sqrt{\sum_{i=1}^{n} (x^i - y^i)^2}.$$

The quantity $\|x-y\|$ is the distance between points x and y. This distance generalizes the absolute value used in $\mathbb{R}^1$.

Definition 6 *An open ball of radius $\varepsilon > 0$ and center $x \in X$, is denoted $B(x,\varepsilon) = \{y \in X : \|x-y\| < \varepsilon\}$.*

Definition 7 *A set $A \subset \mathbb{R}^n$ is open if for every $x \in A$ there is some $\varepsilon > 0$ s.t. $B(x,\varepsilon) \subset A$.*

We call a set closed if its complement is open.

Definition 8 *A set $A \subset \mathbb{R}^n$ is closed if its complement $B = \mathbb{R}^n \backslash A$ is an open set.*

Another relevant condition characterizes sets that have bounds.

Definition 9 *A set $A \subset \mathbb{R}^n$ is bounded if there exists a finite number b s.t. for every $x \in A$ it is the case that $\|x-0\| < b$.*

Our first pathology (example x) involved an X that was either not closed or not bounded. In Euclidean space $\mathbb{R}^n$ the sets that are both closed and bounded are called compact.

Definition 10 *A set $A \subset \mathbb{R}^n$ is compact if it is closed and bounded.*

In arbitrary spaces (for example infintite dimensional spaces) the equivalenve between compactness and closed and bounded does not hold. In this book all examples or problems will deal with subsets of Euclidean space. However it is instructive to see how concepts and results may be generalized to arbitrary spaces. The standard approach is to consider a space and a collection of open sets that satisfy a small number of requirements rendoring it a Topological space. Here way introduce an alternative definition of compactness because it makes the proof of our next result substantially easier (and more elegant). The detour into more abstract analysis is worthwhile as subsequent results will also involve the assumption that relevant spaces are compact.

The more general (or Topological) definition of compactness deals with open covers.

Definition 11 *Given a set A an open covering of A is a collection of sets $\{O_\theta\}_{\theta\in\Theta}$ where Θ is an arbitrary index set and O_θ is open for every $\theta \in \Theta$ s.t. $A \subset \{\cup_{\theta\in\Theta} O_\theta\}$ (in other words if $x \in A$ then there is some $\theta \in \Theta$ s.t. $x \in O_\theta$).*

A set is compact if every open covering has a finite sub covering.

Definition 12 *A set A is compact if $\{O_\theta\}_{\theta\in\Theta}$ an open covering of A implies that for some finite set $B \subset \Theta$, $\{O_\theta\}_{\theta\in B}$ is a covering of A.*

We now consider three conditions a preference relation R on $X \subset \mathbb{R}^n$ can satisfy. These concepts make use of a concept called the upper contour (or more preferred set). Given a set X and a binary relation R on X the (strict) upper contour set of a point $x \in X$ is $P(x) := \{y \in X : yPx\}$. The (strict) lower contour set of point x is $P^{-1}(x) := \{y \in X : xPy\}$. So the upper contour set of x contains the points that are preferred to x and the lower contour set of x contains the points that x is preferred to. Similarly, the level set of x is $I(x) = \{y \in X : yRx \text{ and } xRy\}$.

Definition 13 *A binary relation R on X is*
(i) upper continuous if for all $x \in X$, $P(x)$ is open
(ii) lower continuous if for all $x \in X$, $P^{-1}(x)$ is open
(iii) continuous if is both lower and upper continuous.[3]

The intuition behind these conditions is illuminating. When preferences are complete for any policy x any point y that is very close to x is either in $P(x), P^{-1}(x)$, or $I(x)$. When preferences are continuous if $y \in P(x)$ or $y \in P^{-1}(x)$ then points sufficiently close to y will also be in that set. The jump in preferences exhibited in the second example is ruled out when preferences are lower continuous. An important results can now be stated.

Theorem 3 *If X is non-empty and Compact, and R on X is complete, reflexive, transitive and lower continuous, then $M(R,X) \neq \emptyset$*

The proof of this result is more technical than most other sections of these notes. The result also holds on an arbitrary topological space (allowing it to apply to choice problems in which x is say a infinite sequence of consumptions, or a function, or a probability distribution.

> **Proof:** Assume that X is non-empty and Compact, and R on X is complete, reflexive, transitive and lower continuous. By way of contradiction assume that $M(R,X) = \emptyset$. This means that every point in X is contained in some set $P^{-1}(\alpha)$ for some $\alpha \in X$. Since R is lower continuous every such $P^{-1}(\alpha)$ is open. This means that $\{P^{-1}(\alpha)\}_{\alpha\in X}$ is an open covering of X. Since X is compact there exists a finite set of points $B \subset X$ for which the collection $\{P^{-1}(\alpha)\}_{\alpha\in B}$ is also a covering of X. (That is if $x \in X$ then $x \in P^{-1}(\alpha)$ for some $\alpha \in B$). But we know from a previous result that $M(R,B) \neq \emptyset$, since B is finite and R is complete, reflexive and transitive. Thus a point $x^* \in M(R,B)$ exists. Now any point $y \in X$

[3] A very careful student may note that the term "open" in this definition should be modified to "open relative to X".

is either in $M(R, B)$ or it isn't. By definition if $y \in M(R, B)$ then $x^* R y$. If $y \notin M(R, B)$ since $\{P^{-1}(\alpha)\}_{\alpha \in B}$ covers X there is some $\alpha \in B$ s.t. $y \in P^{-1}(\alpha)$. But this means that $\alpha R y$. Recall that by assumption $x^* \in M(R, B)$ meaning that $x^* R \alpha$. Since R is transitive on X these two facts mean that $x^* R y$. Thus we have shown that for all $y \in X$, $x^* R y$. This means that $x^* \in M(R, X)$, contradicting the assumption. Thus we have established the non-emptyness of this set.■

4.2.2 Uniqueness of $M(R, X)$

When the choice space is not finite, we can impose sufficient structure to insure that $M(R, X)$ has only one element. We will need one condition on X and one condition on R to attain this result.

Definition 14 *We say the set $X \subset R^n$ is convex if for any $x, y \in X$ and the point $\lambda x + (1 - \lambda) y$ is in X for every $\lambda \in [0, 1]$*

The point $\lambda x + (1 - \lambda) y$ is often called the convex combination (or a weighted average) of x and y. As an example the set $[0, 1]$ is convex because for any two points in the set, any point in between these two points is also in the set. Convexity requires that there are no holes in a space. When the space is more than 1 dimensional convexity also requires that its surface not have any appendages.

We now consider a property of preferences on convex sets.

Definition 15 *We say the preference R on the convex set X is strictly convex if for any distinct points $x, y \in X$ if $x R y$ then $[\lambda x + (1 - \lambda) y] P y$ for any $\lambda \in (0, 1)$.*

Strict convexity of preferences requires that the sets $P^{-1}(x)$ are convex, and that the boundary of these sets contain no flat edges. The following powerful result is very easy to establish.

Theorem 4 *If X is convex and R on X is strictly convex, then if $M(R, X)$ is non-empty it contains a single element.*

Proof: By way of a contradiction assume that X is convex, R is strictly convex, and two distinct policies x, y are both in $M(R, X)$. For arbitrary $\lambda \in (0, 1)$ the point $[\lambda x + (1 - \lambda) y]$ is in X since X is convex. But since R is strictly convex, $[\lambda x + (1 - \lambda) y] P y$. But this contradicts the assumptions that $y \in M(R, X)$. Thus the result is established.■

The import of the last two theorems is straightforward. When the choice set is compact and a weak order is also lower continuous rationality is a meaningful concept –a rational or optimal choice/decision exists. Whent the choice set is convex and the prefernece ordering is strictly convex if an optimal choice exists it is unique. In otherwords if rationality is meaningful (in the sense that an optimal choice exists) that we have a tight point prediction of what the agent will choose.

5 Utility Theory

So far our understanding of choice and rationality is based on the use of binary preferences and the maximal set. In many cases a richer theory of choice is needed. In this chapter we consider a theory of choice based on utility functions. Recall that given a set X a real-valued function is an assignment that gives each element $x \in X$ a value $f(x) \in \mathbb{R}^1$. We can express such a function as $f : X \to \mathbb{R}^1$. The set X is generally referred to as the domain of the function and the set $\mathbb{R}^1$ is generally referred to as the range or image. Examples of a functions defined on the domain $X = \mathbb{R}^1$ are $f(x) = x^2$; $f(x) = 3x-1$; $f(x) = x$; $f(x) = 3$. We are interested in representing a notion of choice based on a utility function $u : X \to \mathbb{R}^1$, where $u(x)$ is a metric of the level of satisfaction associated with the choice x. The perspective taken here is that given a choice space X, the underlying primitive is the preference R on X. We are interested in the possibility of representing such preferences with utility functions,

Definition 16 *Given X and R on X we say the utility function $u : X \to \mathbb{R}^1$ represents R if for all $x, y \in X$ $u(x) \geq u(y)$ iff xRy.*

Exercise 5 *Show that if $u : X \to \mathbb{R}^1$represents R then $u(x) > u(y)$ iff xPy and $u(x) = u(y)$ iff xIy.*

When X is finite and R is complete, reflexive and transitive, the existence of a utility representation of R on X is obvious.

Exercise 6 *For X finite and R a weak ordering on X exhibit an algorithm for constructing a utility function representing R.*

Just as we saw previously that when X is finite, and R is a weak ordering $M(R, X)$ is non-empty, it can be shown that if X is finite and $u : X \to \mathbb{R}^1$ a maximizer exists. Note that x is a maximizer of $u : X \to \mathbb{R}^1$ if $u(x) \geq u(y)$ for all $y \in X$. If X is not finite further conditions on X and the function are needed to insure the existence of maximizers.

A nice property of functions is continuity.

Definition 17 *We say a function $f : X \to \mathbb{R}^1$ is continuous if for every $x \in X$ the following is true: For every $\varepsilon > 0$ there exists some $\delta > 0$ s.t. if $\|x - y\| < \delta$ $|f(x) - f(y)| < \varepsilon$.*

A very useful result is:

Theorem 5 *(Debreu 1959) If $X \subset R^n$ and R is complete, reflexive, transitive, and continuous, then there exists a continuous utility function $u : X \to \mathbb{R}^1$ that represents R.*

We will not undertake the proof of this claim. However the converse is not difficult to establish.

Exercise 7 *Show that if $X \subset R^n$ and the continuous utility function $u : X \to \mathbb{R}^1$ represents the binary ordering R on X then R is complete, reflexive, transitive, and continuous.*

A result analogous to theorem 3 is the following.

Theorem 6 *If $X \subset \mathbb{R}^n$ is compact and $u : X \to \mathbb{R}^1$ is continuous then a maximizer exists.*

This result is sometimes known as the Weierstrass Theorem (after a famous mathematician). It is important to note that the utility functions here are arbitrary. Some texts call this an ordinal notion of utility as opposed to a cardinal notion of utility. There is nothing interesting about the particular value of a utility function at a specific point $x \in X$. All that matters is the ordering of $u(x)$ and $u(y)$ for any two $x, y \in X$. We say that $f : R^1 \to R^1$ is a strictly increasing function if for all $x, y \in X$ $x > y$ implies that $f(x) > f(y)$. A more precise way of expressing the indeterminacy or ordinal nature of utility functions is: Utility functions are defined only up to a strictly increasing transformation. This means that if $u : X \to \mathbb{R}^1$ represents R on X then $f \circ u : X \to \mathbb{R}^1$ represents R on X. Note that $f \circ u : X \to \mathbb{R}^1$ is a fancy way of writing $f(u(x))$. Thus scaling a utility function is of no consequence. In the next section we will deal with utility functions defined over lotteries on the choice space. In this section we will need something stronger than ordinal utility functions.

Thus far we have not established any results that allow us to find the maximizer of a utility function. Fortunately, if we assume that utility functions are differentiable, the tools of calculus will allow us to characterize extrema. We do not address these results here.

6 Expected Utility Theory

6.1 Choice Under Uncertainty

For the theory of games that von-Neumann and Morgenstern constructed ordinal utility is insufficient. We need to be able to analyze the choices of rational agents when they make choices in risky or uncertain environments. We begin with a brief review of probability theory.

Our model of uncertainty involves a state space Ω, a set of events $\mathcal{F}$ and a function indicating the probability of each event in $\mathcal{F}$, $p : \mathcal{F} \to [0,1]$. An element $\omega \in \Omega$ is a random state. An element $A \in \mathcal{F}$ is actually a subset of Ω, and $p(A)$ is the probability that the true state is in the set A. We require that

$$\text{If } A \in \mathcal{F} \text{ then } A^c \in \mathcal{F} \text{ where } A^c = \Omega \backslash A$$
$$\text{If } A_\alpha \text{ is a finite collection of subsets in } \mathcal{F} \text{ then } \cup_\alpha A_\alpha \in \mathcal{F}$$

A few conditions must be satisfied, by the function that indicates the probability of an event:

$$\text{If } A, B \in \mathcal{F} \text{ then } p(A \cup B) = p(A) + p(B) - p(A \cap B)$$

$$p(\emptyset) = 0,\ p(\Omega) = 1.$$

If Ω is finite then for any $A \in \mathcal{F}$

$$p(A) = \sum_{\omega \in A} p(\omega).$$

As an example consider the toss of a fair 3-sided die. $\Omega = \{1,2,3\}, \mathcal{F} = \{\emptyset, \{1\}, \{2\}, \{3\}, \{1,2\}, \{1,3\}, \{2,3\}, \{1,2,3\}\}$ and $p(A) = \frac{1}{3}$ if A has two elements, $p(A) = \frac{2}{3}$ if A has two elements, and $p(\{1,2,3\}) = 1$. We will only consider finite state spaces in this section.

Our goal is to consider an agent choosing from a choice space X, facing uncertainty of the form $(\Omega, \mathcal{F}, p)$ and a utility function $u : X \times \Omega \to \mathbb{R}^1$. With a utility function of this form, the utility of a particular choice $x \in X$ depends on the unknown state $\omega \in \Omega$, and is written $u(x, \omega)$. Since the agent does not know the true state, she must choose the policy $x \in X$ which maximizes her expected utility, where the expected utility of a choice x is of the form

$$U(x) = \sum_{\omega \in \Omega} u(x, \omega) p(\omega).$$

The expected utility of x is thus the weighted average value of $u(x, \omega)$ where the weights are $p(\omega)$. Thus low probability events, will effect the expected utility less than high probability events. The remainder of this section is focused on establishing the existence of an expected utility function from primitives on preferences over lotteries. We now consider a slightly simpler model. Given a finite state space $\Omega = \{\omega_1, ..., \omega_n\}$ let $\Delta(\Omega)$ denote the set of lotteries on Ω. Thus an element $L \in \Delta(\Omega)$ is vector $L = (p_1, ..., p_n)$ where $p_n \in [0, 1]$ is the probability that state ω_n occurs under lottery L. Note that for any $L \in \Delta(\Omega)$ it must be the case that $\sum_{i=1}^{n} p_i = 1$. We want to consider preferences defined over this space of lotteries, that is preferences on the space $\Delta(\Omega)$. Since $\Delta(\Omega) \subset \mathbb{R}^n$ the concepts of section (x) apply. More precisely our notions of completeness, reflexivity, and transitivity apply. We now consider two additional conditions that a binary relation R on $\Delta(\Omega)$ can satisfy.

Definition 18 *We say that the relation R on $\Delta(\Omega)$ is continuous if for any $L, L', L'' \in \Delta(\Omega)$*
LPL' and $L'PL''$ imply that for some $\lambda \in (0, 1)$ $[\lambda L + (1 - \lambda)L'']\, IL'$.

The compound lottery $\lambda L + (1 - \lambda)L''$ places probability $\lambda p_i + (1 - \lambda)p_i''$ on ω_i. When preferences are continuous on $\Delta(\Omega)$ then if $[\lambda L + (1 - \lambda)L']\, RL''$ it is also the case that $[\lambda' L + (1 - \lambda')L']\, RL''$ when λ' is sufficiently close to λ. Similarly if $L''R\, [\lambda L + (1 - \lambda)L']$ then $L''R\, [\lambda' L + (1 - \lambda')L']$ when λ' is sufficiently close to λ. Thus small changes in lotteries result in small enough changes in the lotteries don't effect the binary preference much. The second condition is:

Definition 19 *We say that the relation R on $\Delta(\Omega)$ satisfies the independence axiom if for any $L, L', L'' \in \Delta(\Omega)$ and $\lambda \in [0, 1]$*
LRL' implies that $[\lambda L + (1 - \lambda)L'']R\ [\lambda L' + (1 - \lambda)L'']$.

Preferences on $\Delta(\Omega)$ satisfy the independence axiom when the preference over two lotteries is independent of the mixture of a third lottery to the first two. von Neumann and Morgenstern established:

Theorem 7 *(Expected Utility Theorem) If R on $\Delta(\Omega)$ (with Ω containing n elements) is a weak ordering (complete, reflexive and transitive) and satisfies continuity and the independence axiom then there is a utility function $U : \Delta(\Omega) \to \mathbb{R}^1$ that represents R. Moreover there exists a vector of n numbers $(u_1, ..., u_i, .., u_n)$ with each $u_i \in \mathbb{R}^1$ s.t. for any $L \in \Delta(\Omega)$*

$$U(L) = \sum_{i=1}^{n} p_i u_i.$$

The expected utility theorem establishes the existence of a linear representation of utility where the expected utility of any lottery $U(L)$ is equal to the weighted sum of the utilities of each state, where the weights correspond to the state probabilities induced by the lottery L. Recall that our theory of utility functions over choice spaces was an ordinal theory–utility functions were invariant to strictly increasing transformations. The expected utility representation is more sensitive. In some sense these utility functions are cardinal. While a given R on X is represented by many distinct utility functions, the set of such functions is smaller. Here utility functions are defined up to affine or linear transformations. This means that if $U(\cdot)$ represents R on $\Delta(\Omega)$ then for any

$a \in \mathbb{R}^1$ and $b > 0$, the function $V(\cdot) = a + bU(\cdot)$ also represents R. The vector $(u_1, ..., u_i, .., u_n)$ is sometimes called a vector of Bernoulli utility functions or state-utility functions. In geometric terms if you think about $(u_1, ..., u_i, .., u_n)$ as a point in $\mathbb{R}^n$ then any linear operation on $(u_1, ..., u_i, .., u_n)$ yields another vector of Bernoulli utility functions $(a + bu_1, ..., a + bu_i, .., a + bu_n)$ that represent the same preferences. We will interchange the notation u_ω and $u(\omega)$ throughout.

When analyzing game-theoretic models involving uncertainty we will assume that agents have preferences that are representable by von Neumann-Morgenstern utility functions.

6.2 Risk-Preferences

The expected utility representation is very convenient. A simple analysis will allow us to gain some intuition for choice under uncertainty.

Example 3 *We now consider an example of an agent choosing between lotteries over money on the very simple state space $\Omega = \{-1, 0, 1\}$, So an element of $\Delta(\Omega)$ is a pair (p_{-1}, p_0, p_1) with $p_i \in [0, 1]$ and $p_{-1} + p_0 + p_1 = 1$. It may be helpful to think about lotteries over paying a dollar, neither paying nor receiving a dollar, and receiving a dollar respectively. Now given the expected utility theorem we can represent any R (satisfying the conditions stated in the theorem) on $\Delta(\Omega)$ with the function that assigns expected utility to L of $U(L) = p_{-1}u_{-1} + p_0u_0 + p_1u_1$. An important point to note is that the theorem does not say that vector of Bernoulli utility functions $u = (-1, 0, 1)$ will work. If this particular vector can be used then the preferences are said to be risk-neutral. Such preferences exhibit the property the agent is indifferent between a lottery and its certainty equivalent.*

Definition 20 *The certainty equivalent of a lottery L over money is a number $cer(L)$ satisfying*

$$u(cer(L)) = \sum_{\omega \in \Omega} p_\omega u(\omega).$$

Definition 21 *The expected value of a lottery L over money is*

$$E(L) = \sum_{\omega \in \Omega} p_\omega \omega.$$

A risk neutral individual is indifferent between the lotteries $L = (\frac{1}{2}, 0, \frac{1}{2})$ and $L' = (0, 1, 0)$. In fact a risk neutral agent is indifferent (by the independence axiom) between any two lotteries of the form $L'' = (\lambda, 1 - 2\lambda, \lambda)$ and $L''' = (\lambda', 1 - 2\lambda', \lambda')$.

Definition 22 *Preferences exhibit risk-neutrality if for any $L \in \Delta(\Omega)$ $U(L) = U(E(L)$. Preferences exhibit risk aversion if for any $L \in \Delta(\Omega)$ $U(L) \leq U(E(L))$. Preferences exhibit risk acceptance if for any $L \in \Delta(\Omega)$ $U(L) \geq$*

$U(E(L))$. The last two conditions are termed strict risk aversion and strict risk acceptance if the inequalities hold strictly for every non degenerate. $L \in \Delta(\Omega)$.

There is a direct relationship between the risk attitude of preferences and the concavity of the Bernoulli utility functions.

Exercise 8 *Show that if preferences exhibit risk neutrality then $cer(L) = E(L)$, if preferences exhibit risk aversion $cer(L) \leq E(L)$, and if preferences exhibit risk acceptance $cer(L) \geq E(L)$.*

Definition 23 *A function $f : X \to R^1$ is linear if for any $x, y \in X$ and any $\lambda \in (0, 1)$ $f(\lambda x + (1 - \lambda)y) = \lambda f(x) + (1 - \lambda)f(y)$*

Definition 24 *A function $f : X \to R^1$ is concave if for any $x, y \in X$ and any $\lambda \in (0, 1)$ $f(\lambda x + (1 - \lambda)y) \geq \lambda f(x) + (1 - \lambda)f(y)$. If the inequality is always strict it is strictly concave.*

Definition 25 *A function $f : X \to R^1$ is convex if for any $x, y \in X$ and any $\lambda \in (0, 1)$ $f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y)$. If the inequality is always strict it is strictly convex.*

Exercise 9 *In the setting of the above example if $u_{-1} = -1$ and $u_1 = 1$ find the set of values that u_0 may take to make the agent (1) risk neutral, (2) risk acceptant and (3) risk averse.*

The following result makes clear the relationship between risk attitudes and the curvature of the Bernoulli utility functions

Theorem 8 *Consider preferences that satisfy the conditions of the expected utility theorem defined on the set $\Delta(\Omega)$ and a Bernoulli utility function $u = (u_1, ..., u_n)$ s.t. the preferences are represented by the expected utility function $U(L) = \sum_{i=1}^{n} p_i u_i$: 1. Preferences exhibit risk neutrality iff the Bernoulli utility function is linear. 2. Preferences exhibit risk aversion iff the Bernoulli utility function is concave. 3. Preferences exhibit risk acceptance iff the Bernoulli utility function is convex*

Exercise 10 *Prove theorem 8. Hint– show 2 and 3 and then note that the only functions which are both convex and concave are those that are linear.*

The assumptions of the expected utility theorem may seem straightforward but many examples of choice inconsistent with the independence axiom seem to have be recovered in experimental settings. Morrow exhibits the Allais paradox on page 45 which demonstrates that some individuals exhibit behavior that seems to violate the independence axiom. Much work in decision theory is focused on the construction and testing of theories that make weaker assumptions than those of the classical expected utility model.

7 Social Choice Theory

So far we have constructed a theory of individual choice based on preferences (generally represented as weak orderings). We now construct a theory of preference aggregation. Social choice theory seeks an understanding of the properties of mechanisms that aggregate individual preferences into collective or social choices.

We will take as primitives a finite collection of agents $N = \{1, 2, ..., n\}$ ($n > 2$) a choice space X and a list of n weak orderings $\rho = \{R_1, R_2, ..., R_n\}$. We call such a list of orderings a preference profile. A typical agent i has preference ordering R_i on X. Given X the set of weak orderings is denoted $\mathcal{R}$. The set of profiles (or n-tuples) is denoted $\mathcal{R}^n$. By $\mathcal{B}$ we denote the set of complete and reflexive orderings on X.

Definition 26 *A preference aggregation rule is a function* $f : \mathcal{R}^n \rightarrow \mathcal{B}$.

Examples like majority rule, and the Borda rule will be described below. An uninteresting but simple example is a constant preference aggregation rule that assigns the same complete and reflexive ordering to any $\rho \in \mathcal{R}^n$.

7.1 Properties of Preference Aggregation Rules

Individual ordering are subscripted, R_i, the social ordering $R_{f(\rho)}$ is denoted R. We consider three conditions that a preference aggregation rule can satisfy:

Definition 27 *A preference aggregation rule f is:*
i. Transitive if for every $\rho \in \mathcal{R}^n$ the ordering R is transitive.
ii. Non dictatorial if there does not exist an $i \in N$ s.t. for every $\rho \in \mathcal{R}^n$ for every $x, y \in X$, xP_iy implies xPy.
iii. Weakly Paretian if, for any $x, y \in X$, if xP_iy for every $i \in N$ then xPy.
iv. Independent of irrelevant alternatives if, for any pair of policies $x, y \in X$ and any two profiles $\rho, \rho' \in \mathcal{R}^n$ with xR_iy iff xR'_iy for all $i \in N$, xRy iff $xR'y$.

A rule is transitive when for every preference profile the collective ordering is a weak ordering. A rule is dictatorial if there is some agent whos preferences are always satisfied by the social ordering – regardless of the alternative, her preferences and the preferences of the other agents. A rule is weakly Paretian if whenever all the agents agree on a pairwise comparison, the collective preference also agrees. A rule satisfies independence of irrelevant alternatives (IIA) if the social preference on any two alternatives is not affected by the ranking of a third alternative. Two examples of preference aggregation rules are:

-majority rule: for all $x, y \in X$ xPy iff $|i \in N : xP_iy| > \frac{n}{2}$.

Assuming each agent has complete strict preferences and letting $r_i(x) = |z : zP_ix|$,

-Borda count rule: for all $x, y \in X$ xPy iff $\sum_{i \in N} r_i(x) < \sum_{i \in N} r_i(y)$.

Exercise 11 *Demonstrate that Majority rule is not transitive.*

Exercise 12 *Demonstrate that the Borda count rule is not IIA.*

We need one more definition before turning to Arrow's Theorem.

Definition 28 *Given a preference aggregation rule f a set $L \subset N$ is semidecisive for x against y if for every $\rho \in \mathcal{R}^n$ with xP_iy (all $i \in L$) and yP_jx (all $j \in L^c = N \backslash L$) we have xPy. A set L is decisive for x against y if for every $\rho \in \mathcal{R}^n$ with xP_iy (all $i \in L$) we have xPy. A set L is decisive if for every $x, y \in X$ it is decisive for x against y.*

We can now state and prove Arrow's Theorem.

Theorem 9 *(Arrow); If X is finite and has at least three alternatives, then there is no preference aggregation rule $f : \mathcal{R}^n \rightarrow \mathcal{B}$ that is transitive, non dictatorial, weakly Paretian and independent of irrelevant alternatives.*

Before proving the theorem we need to state and prove a useful lemma.

Lemma 1 *Assume f is a transitive preference aggregation rule that is independent of irrelevant alternatives and weakly Paretian. If $L \subset N$ is semidecisive for x against y for some $x, y \in X$ then L is decisive.*

Proof: Assume that $L \subset N$ is semidecisive for x against y and that under the profile $\rho \in \mathcal{R}^n$ xP_iz for all $i \in L$. Consider a profile $\rho' \in \mathcal{R}^n$ s.t. for all $i \in L$ $xP_i'yP_i'z$ and for all $j \in L^c$ $yP_j'x$ and $yP_j'z$ with $z \notin \{x, y\}$ and for all $i \in L^c$ xR_iz iff $xR_i'z$. Since L is semidecisive for x against y $xP'y$. Since f is weakly Paretian $yP'z$. Since f is transitive $xP'z$. But since the preferences of L^c on x, z have not been specified in ρ' and both ρ and ρ' agree on x and z (i.e. xR_iz iff $xR_i'z$), the fact that f is IIA implies xPz. Thus L **is decisive for** x **against** z. This of course implies that L is semidecisive for x against z and an analogous argument demonstrates that L **is decisive for** x **against** y. We now verify that L is decisive for y against z. Consider a profile $\rho^0 \in \mathcal{R}^n$ with yP_i^0z for all $i \in L$ and $\rho^+ \in \mathcal{R}^n$ s.t. for all $i \in L$ $yP_i^+xP_i^+z$ and for all $j \in L^c$ zP_j^+x and yP_j^+x and for all $i \in L^c$ yR_i^0z iff yR_i^+z. Since we have already shown that L is decisive for x against z we have xP^+z. Since f is weakly Paretian we have yP^+x. Since f is transitive we have yP^+z. Since the preferences of only members of L have been specified on $\{y, z\}$ by ρ^+ and both ρ^0 and ρ^+ agree on y and z, IIA implies yP^0z. Thus L **is decisive for** y **against** z. This of course implies that L is semidecisive for y against z. Relabeling the first step and using this fact implies that L **is decisive for** y **against** x. Combining these (boldfaced) conclusions leads to the claim that L is decisive.■

We now prove the theorem

> **Proof of Arrow's Theorem:** Assume that X is finite and has at least three alternatives. By way of a contradiction assume that we have a preference aggregation rule that is transitive, non dictatorial, weakly Paretian and independent of irrelevant alternatives. Given the lemma, for any set $L \subset N$ either L is decisive or there is no pair $x, y \in X$ for which L is semidecisive for x against y. Consider two disjoint sets $A, B \subset N$ (disjoint means that $A \cap B = \emptyset$) which are not semidecisive for any x and y (and thus not decisive) Let $C = N \backslash \{A \cup B\}$. Since $n > 2$, and no singleton set $\{i\}$ is decisive, three such sets A, B, C exist. Now consider the profile $\rho^- \in \mathcal{R}^n$ with $xP_i^- yP_i^- z$ for $i \in A$; $zP_j^- xP_j^- y$ for $j \in B$; and $yP_t^- zP_t^- x$ for $t \in C$. Since A and B are not semidecisive for any pairs, we must have $zP^- x$ and $yP^- z$. Since f is transitive we must have $yP^- x$. This implies that the set $A \cup B$ is not semidecisive for x against y. This means that the set is not decisive. Thus the union of two disjoint sets which are not decisive is also not decisive. Since f is not dictatorial no singleton set is decisive. This conclusion means that no (finite) union of individuals is decisive. But this implies that N is not decisive. This contradicts the assumption that f is weakly Paretian. Thus the result is established.■

It is important to note that a preference aggregation rule has as its domain the space of profiles of weak orderings. The theorem does not state that for a particular preference profile no desirable means to aggregate the preferences exists. One response to Arrow's theorem is to consider restrictions to the domain $\mathcal{R}^n$ and consider whether there are preference aggregation rules that satisfy the normative axioms on this smaller domain An important restriction is termed single-peakedness. In the definition we will use the notion of a one-to-one and onto function (often termed a bijective function). A function $q : X \to X$ is one-to-one if for every $y \in X$ the set $q^{-1}(y) = \{x \in X; q(x) = y\}$ is a singleton. The function is onto if for every $y \in X$ there is some $x \in X$ s.t. $q(x) = y$.

Definition 29 *Given a set N and a choice space X a preference profile $\rho \in \mathcal{R}^n$ is single-peaked if there exists some function $q : X \to X$ that is one-to-one and onto such that for every $i \in N$ there is some $t_i \in X$ s.t. if $q(x) < q(y) < q(t_i)$ then $t_i P_i y P_i x$ and if $q(t_i) < q(b) < q(c)$ then $t_i P_i b P_i c$. The set of single-peaked profiles is denoted $\mathcal{S} \subset \mathcal{R}^n$.*

In the definition the policy t_i is interpreted as i's ideal policy, and the further $q(y)$ is from $q(t_i)$ the less the agent prefers y. It is crucial to note that we are permuting the space X with one function q and given this transformation each agent has an ideal point t_i and the ranking of alternatives decreases as the alternatives are farther from t_i in terms of $q(\cdot)$.

Exercise 13 *Show that the following preferences on* $X = \{1, 2, 3\}$ *are single-peaked,* $1P_1 2P_1 3$; $1P_2 3P_2 2$; $3P_3 1P_3 2$.

Exercise 14 *Show that if* ρ *is single-peaked and* $\{t_i\}_{i \in N}$ *is a list of* n *ideal points that are defined in the previous definition then for any* $i \in N$ $t_i \in M(R_i, X)$

Exercise 15 *Show that if* ρ *is single-peaked each agent has a unique ideal point* t_i.

A comforting result is:

Theorem 10 *Given* $\rho \in \mathcal{S}$ *majority rule is transitive, weakly Paretian, IIA and non dictatorial.*

Exercise 16 *Prove the theorem*

Since we are interested in collective choice, we are ultimately interested in the set of policies that are maximal for a preference aggregation rule.

Definition 30 *Given* X, $\rho \in \mathcal{R}^n$ *and a preference aggregation rule, the core is defined as* $C_{f(\rho)}(X) = M(R, X)$ *where* $R = f(\rho)$.

Applying theorem 1, we see that if X is finite and the collective preference is a weak ordering then the core is non-empty. However our analysis of Arrow's theorem indicates that transitivity of preference aggregation rules is not always satisfied.

7.2 Properties of the Majority Rule Core

In the last exercise you showed that if $\rho \in \mathcal{S}$ then majority rule yields a weak ordering. We can establish a stronger result than just the non-emptyness of the core in this case. The famous median voter theorem of Duncan Black can be stated as follows:

Theorem 11 *(Black) If* $n > 2$ *is odd and* $\rho \in \mathcal{S}$ *then letting* $f(\cdot)$ *be majority rule,* $C_{f(\rho)}(X) = \{t_i : |j \in N\backslash i : t_j \leq t_i| = |k \in N\backslash i : t_k \geq t_i|\}$. *That is the core is the median voter's ideal point.*

> *Proof:* Without loss of generality assume that with $q(x) = x$ for all $x \in X$ there is some $t_i \in X$ s.t. if $q(x) < q(y) < q(t_i)$ then $t_i P_i y P_i x$ and if $q(t_i) < q(b) < q(c)$ then $t_i P_i b P_i c$. Now consider the set $\{t_i : |j \in N\backslash i : t_j \leq t_i| = |k \in N\backslash i : t_k \geq t_i|\}$. Since n is odd this set is a singleton. The point in thise set is the median ideal point, we will call this set m.

We first show that $m \in C_{f(\rho)}(X)$. Clearly $|j \in N : t_j \leq m| > \frac{n}{2}$. Now if $y > m$ then $\{j \in N : t_j < m\} \subset \{i \in N : mP_iy\}$ so $|i \in N : mP_iy| > \frac{n}{2}$, thus mPy. Similarly $|j \in N : t_j \geq m| > \frac{n}{2}$. If $y < m$ then $\{j \in N : t_j > m\} \subset \{i \in N : mP_iy\}$ so $|i \in N : mP_iy| > \frac{n}{2}$ thus mPy. These two facts imply that $m \in C_{f(\rho)}(X)$.

We now show that there cannot be some $y \neq m$ with $y \in C_{f(\rho)}(X)$. By way of a contradiction assume that such y exists. The fact that either $y < m$ or $y > m$ implies that $|i \in N : mP_iy| > \frac{n}{2}$ and thus mPy contradicting the assumption.■

Exercise 17 *Characterize the majority rule core when $n > 3$ is even.*

These results indicate that if we are willing to assume that preferences satisfy the restrictive assumption of single-peakedness the majority rule core is well defined, and a notion of the "will of the majority" makes sense. However the restriction may not be appropriate for some settings. When the dimensionality of the policy space is 2 or more a natural generalization of the single-peakedness exists. The well studied spatial model is one such generalization.

Definition 31 *In the spatial model we assume $X \subset \mathbb{R}^d$ (d finite) is convex and agents have strictly convex, continuous preferences on X.*

If X is compact then theorems 3 and 4 imply that each agent has a unique ideal point y_i in X. Instead of assuming that X is compact we will assume directly that each agent has an ideal point. The assumption that preferences are strictly convex requires that the upper contour sets are convex sets with no flat boundaries. The classic special case of Euclidean preferences, $u_i(x) = \|x - y_i\|$ is convenient as in this case the upper contour sets are spheres. A pair of examples will illustrate the possibilities of preference aggregation.

Example 4 *An example with $X = \mathbb{R}^2$ and a non empty majority rule core is constructed by considering 5 agents with Euclidean preferences and ideal points of*

$$\begin{aligned} y_0 &= (0, 0) \\ y_1 &= (1, 0) \\ y_2 &= (0, 1) \\ y_3 &= (-1, 0) \\ y_4 &= (0, -1). \end{aligned}$$

In this example the point $y_0 = (0, 0)$ is a core point. Any movement from this point makes exactly three agents worse off, and thus y_0Px for any $x \in X \backslash y_0$.

To demonstrate this fact, consider a point in the first quadrant (both coordinates are positive) this point is further from the ideal points of agents 0, 3 and 4 then the point (0,0). Any point in a different quadrant will also be further from the ideal points of 3 agents than the point $(0,0)$. *This same logic implies that no other point can be in the core as it is beaten by the point* $(0,0)$.

While this example may seem promising as it demonstrates how the core can be non-empty, it is not a representative example. This point is demonstrated by the following example

Example 5 *Consider the example above with the one modification that* $y_0 = (\varepsilon, \delta)$ *with* $\varepsilon > 0$ *and* $\delta > 0$ *but small. Now the point* $(0,0)$ *is not a core point because the voters 0,1,2 all prefer the point* y_0 *to the point* $(0,0)$. *But it is also the case that the point* y_0 *is not a core point because the voters 2, 3, 4 all prefer the policy* $(0,\delta)$ *to the policy* (ε,δ). *In this example the core is empty as any point can be beaten by another point.*

It turns out that the positive example satisfies a knife-edged necessary condition for the majority rule core to be non-empty. Given Euclidean preferences for any point $x \neq y_i$ it is clear which direction agent i would like policy to move.

Definition 32 *If preferences are Euclidean then for any* $x \in X$ *the gradient vector* $\nabla u_i(x) = y_i - x$.

The gradient vector is a directed vector or line segment that points away from the origin in the direction that agent i would most like policy to move from point x. For a finite set A we will call a mapping $p : A \to A$ a pairing if it is one-to-one. This means that each i in A is paired with exactly one j in A.

Definition 33 *In the spatial model with Euclidean preferences the Plott conditions are satisfied at a policy* $x \in X$ *if there exists a pairing* $p(\cdot)$ *on the set* $L = \{j \in N : y_j \neq x\}$ *s.t. for every* $i \in L$ $\nabla u_i(x) = -\lambda_i \nabla u_{p(i)}(x)$ *for some* $\lambda_i > 0$.

The intuition is that when the Plott conditions are satisfied at x the set of agents L that do not have x as their ideal point, can be paired so that each agent that wants to move in a particular direction is offset by a particular agent that wants to move in exactly the opposite direction. The example with a core satisfies the Plott conditions at the point $(0,0)$.The following result characterizes the relationship between the Plott conditions in the spatial model with Euclidean preferences and the majority rule core.

Theorem 12 *(Plott) In the spatial model with Euclidean preferences and n odd the point* x *in the interior of* X *is in the core* $C_{f(\rho)}(X)$ *iff the Plott conditions are satisfied at* x.

Recall that x is in the interior of X if there is an open ball $B(x, \varepsilon)$ that is contained in X. Even in the restrictive case of preferences that are Euclidean it is clear that the Plott conditions will not in general be satisfied. More precisely if we think of the space $\mathbb{R}^{dn}$ as the space of possible ideal points of n agents with Euclidean preferences on the choice space $\mathbb{R}^d$ then the subset of $\mathbb{R}^{dn}$ for which the Plott conditions are satisfied at some $x \in \mathbb{R}^d$ is incredibly small. Specifically it contains no open sets and thus has an empty-interior. Informally stated, if one imagined randomly picking an arbitrary profile from this space the probability of selecting one that satisfy the Plott conditions for some point would be 0.The following exercise makes clear the logic.

Exercise 18 *Show that if $\rho \in \mathbb{R}^{dn}$ is a profile of ideal points for which the Plott conditions are satisfied at some $x \in \mathbb{R}^d$ then for every $\varepsilon > 0$ there exists a profile $\rho^\varepsilon \in B(\rho, \varepsilon)$ for which the Plott conditions are not satisfied at any point for the profile ρ^ε.*

While the set of profiles with a core point is very small, for any such profile there is another profile that is arbitrarily close and also has a core point. The next exercise demonstrates that if one consider small perturbations that also yield a core point, then the core point is only perturbed a little.

Exercise 19 *Show that if $\rho \in \mathbb{R}^{dn}$ (n odd) is a profile of ideal points for which the Plott conditions are satisfied at some $x \in \mathbb{R}^d$ then for every $\varepsilon > 0$ there exists a $\delta > 0$ s.t. if $\rho^\delta \in B(\rho, \delta)$ and the Plott conditions are satisfied for some point at the profile ρ^δ then the Plott conditions are satisfied for a point $x' \in B(x, \varepsilon)$ by the profile ρ^δ.*

The assumption that preferences are Euclidean can be replaced by a differentiability condition, which is more general. We may conjecture that even though the core is generally empty, there is some other subset of the policy space which possesses normatively desirable properties and is therefore a reasonable prediction. One such concept is the following.

Definition 34 *For a set X a profile $\rho \in \mathcal{R}^n$ and a preference aggregation rule f the top cycle set $T_{f(\rho)}$ is the set*

$$T_{f(\rho)} = \{x \in X : \forall y \in X \backslash x, \exists \{a_0, ..., a_t\} \subset X \text{ s.t. } a_0 = x, a_t = y \ t < \infty \text{ and } \forall z < t \ a_{t-1} P a_t\}.$$

The top cycle set is the set of points that can be reached from any other point via a finite chain of strict preferences. That is if $x \in T_{f(\rho)}$ then for every $y \in X \backslash x$ we can select a finite number of policies $\{a_1, a_2, ..., a_t\}$ for which $x P a_1 P a_2 P....P a_t P y$. The following result indicates that either the Plott conditions are satisfied or the top cycle set covers the policy space.

Theorem 13 *(McKelvey) In the spatial model either $C_{f(\rho)}(X)$ is non-empty or $T_{f(\rho)} = X$.*

The intuition behind the result can best be obtained in a simple picture. The implications of the last two theorems are striking. In the spatial model with Euclidean preferences unless a knife-edged condition holds (Plott conditions) any policy can be reached by any other policy in a finite chain of strict preferences.

Two interpretations of the negative social choice results are appropriate – one is positive and the other is normative. First, as a positive methodology the study of preference aggregation rules does not offer clear predictions. This is best exemplified by the result from McKelvey showing that it is generally the case that any policy can beat any other policy through a finite agenda. Some have interpreted this result as a prediction of chaos, whereby the theory predicts that politics should be chaotic with observable cycles. This interpretation is naive, as it attributes a positive prediction to results that state quite clearly the theory of social choice does not generally offer predictions. A more reasonable interpretation is that the results demonstrate the need to investigate the political institutions within which collective choice is made. Under this interpretation the conclusion is that a model that takes as primitives only preferences and a preference aggregation rule may be underspecified. The tools of non-cooperative game theory will allow us to construct richer theories of collective choice.

The second, normative conclusion is that the concept of the "will of the majority" is not well defined. In general there is no such thing as the will of the majority at best there is the policy chosen by the institution. Rethinking Arrow's theorem we find that it is not just the case that there is not a will of the majority, but there is not a will of any decisive set (induced by a particular preference aggregation rule).

8 Game Theory I: Normal form games

8.1 Introduction

The results of social choice theory have been termed negative results because the demonstrate that an account of collective choice that takes as primitives only the choice space, a population and individual preferences is not very satisfying. The analysis of preference aggregation rules in this setting does not lead to unambiguous answers to interesting positive and normative questions. One approach has been to incorporate more primitives into the model. The development of non-cooperative game theory offers the opportunity to build richer models of collective choice that include some of the instititutional features present in real collective choice settings. In the following chapters on game theory we present concepts that are quite useful in the analysis of particular political institutions. At the end of these notes we return to the broader question of how the results from game theoretic models of political institutions relate to the results of social choice theory.

8.2 Normal form games

A normal form game of complete and perfect information (we drop the qualifier complete and perfect information in this section) contains the following elements:

1. A set of players N. An arbitrary player is denoted $i \in N$.

2. A set of strategies for each player, S_i $(i \in N)$. An arbitrary strategy is denoted $s_i \in S_i$. By $S := \times_{i \in N} S_i$ we denote the space of strategy profiles. An arbitrary profile is then a vector $s = (s_1, ..., s_i, ..., s_n) \in S$. *By* $S_{-1} := \times_{j \in N \setminus i} S_j$ we denote the space of strategies for every player except i. We often represent s as (s_i, s_{-i})

3. A von Neumann-Morgenstern utility function for each player, $u_i(s) : S \to \mathbb{R}^1$. . Sometimes the utility function for i is denoted $u_i(s_i, s_{-i})$. To be consistent with our notation from previous chapters the functions $u_i(\cdot)$ are Bernoulli utility functions, and given any lottery over S the agent calculates her expected utility under the lottery.

One interpretation of a normal form game is that at period 1 each player submits their strategy $s_i' \in S_i$ to a computer or referee and then in period 2 the agents receive $u_i(s')$ where $s' = (s_1', ..., s_n')$. We will see that games with more periods can actually be reinterpreted as very large normal form games.

Accordingly a normal form game is completely defined by the data: $\langle N, \{S_i, u(\cdot, ..., \cdot)\}_{i \in n} \rangle$. We sometimes use the shorthand $\langle N, S, u \rangle$ to represent a game where u without a subscript represents the vector of utility functions $(u_1(\cdot), ..., u_n(\cdot))$. Some simple but quite interesting games involving two players can be represented as matrices.

Example 1.

As an example the well-studied Prisoners dilemma (PD) involving two arrested individuals who are sent to separate rooms. Since the two players are seasoned criminals they know that if one confesses and the other does not the person that confessed will get a deal and the other will have the book thrown at him; if both confess they will be punished but their cooperation will be rewarded with a moderate jail term; if neither confesses there is not enough evidence and so no punishment will be forthcoming. This game can be represented with the matrix:

$$\begin{bmatrix} \text{player 1}\backslash\text{ player 2} & \tilde{}\ \text{confess} & \text{confess} \\ \tilde{}\ \text{confess} & 4,4 & 0,6 \\ \text{confess} & 6,0 & 2,2 \end{bmatrix}$$

The entries in the matrix represent u_1, u_2 for the corresponding strategy profile. To demonstrate the notation above this game has the following: $N = \{$player 1, player 2$\}$, $S_i = \{$confess, $\tilde{}$ confess$\}$,

$$u_i(s_i, s_{-i}) = \begin{cases} 2 \text{ if } s_i = s_{-i} = \text{ confess} \\ 4 \text{ if } s_i = s_{-i} = \tilde{}\ \text{confess} \\ 6 \text{ if } s_i = \text{confess } \& \ s_{-i} = \tilde{}\ \text{confess} \\ 0 \text{ if } s_i = \tilde{}\ \text{confess } \& \ s_{-i} = \text{confess} \end{cases}.$$

Example 2.

Another popular game is chicken. Two youths drive cars at each other. If one swerves she is called a "chicken" and subjected to years of shame. The bold youth that does not swerve in this case is treated as a hero. If both swerve they are neither shamed nor immortalized. If neither swerves they die–a fate even worse than years of shame. This game is represented by the matrix

$$\begin{bmatrix} \text{player 1}\backslash\text{ player 2} & \tilde{}\ \text{swerve} & \text{swerve} \\ \tilde{}\ \text{swerve} & 0,0 & 6,2 \\ \text{swerve} & 2,6 & 4,4 \end{bmatrix}.$$

Example 3.

A third example comes from the popular childhood game (which resurfaced in Seinfeld as Jerry and George debated who would get a desirable apartment) of "once, twice, three shoot." Each player can either show one finger or two. Player one wins if the two players use different strategies, and player two wins if the players use the same strategy. This game is represented by the matrix.

$$\begin{bmatrix} \text{player 1}\backslash\text{ player 2} & \text{one} & \text{two} \\ \text{one} & 0,1 & 1,0 \\ \text{two} & 1,0 & 0,1 \end{bmatrix}.$$

Example 4.

Of course we can consider games where there are more than two players, and the strategy spaces are not finite. One example is the following game: $N = \{1, 2\}$, $S_i = \mathbb{R}$,

$$u_1(s_1, s_2) = \begin{cases} 1 \text{ if } |s_1| < |s_2| \\ 0 \text{ if } |s_1| > |s_2| \\ \frac{1}{2} \text{ otherwise} \end{cases}$$

$$u_2(s_1, s_2) = \begin{cases} 1 \text{ if } |s_1| > |s_2| \\ 0 \text{ if } |s_1| < |s_2| \\ \frac{1}{2} \text{ otherwise} \end{cases}$$

This example is a reduced form representation of a model of elections in which the median voter has ideal point 0, symmetric single peaked preferences, and flips a coin when indifferent between candidates.

We now consider the question of what profiles of strategies $s \in S$ are reasonable predictions for how agents should play a normal form game. The motivation for non-cooperative game theory is to consider settings where agents cannot constrain themselves by contractual obligations. Accordingly a reasonable requirement for a prediction is that agent i's strategy s_i be a best response to the the strategies played by the other players s_{-i} and that this is true for every $i \in N$.

To formalize this intuition we need to extend the concept of functions to set valued mappings. A correspondence is a mapping that has as its image a set of sets. That is a function $f : A \to B$ maps points in A into points in B, and a correspondence $c : A \rightarrow\rightarrow B$ maps points in A into subsets of B. We define the best response correspondence.

Definition 35 *The best response correspondence for agent $i \in N$ is a mapping $b_i(s_{-i}) : S_{-i} \rightarrow\rightarrow S_i$ defined as*

$$b_i(s_{-i}) = \arg\max_{s_i' \in S_i} \{u_i(s_i', s_{-i})\}$$

for every $s_{-i} \in S_{-i}$.

The argument maximizer is defined as follows:

$$\arg\max_{s_i' \in S_i} \{u_i(s_i', s_{-i})\} = \{s_i' \in S_i : u_i(s_i', s_{-i}) \geq u_i(s_i'', s_{-i}) \text{ for every } s_i'' \in S_i\}.$$

In the PD the best response correspondences are:

$$b_1(\text{confess}) = \{\text{confess}\}$$
$$b_1(\tilde{}\ \text{confess}) = \{\text{confess}\}$$
$$b_2(\text{confess}) = \{\text{confess}\}$$
$$b_2(\tilde{}\ \text{confess}) = \{\text{confess}\}.$$

In principal the best response for a given strategy may be a non-singleton set. This is why we use the term best response correspondence instead of best response function. Inspecting the best response correspondences exhibited above we see that both players have an incentive to play a particular strategy (here "confess") regardless of what there opponent does. This state of fairs is called the existence of a dominant strategy. Accordingly a reasonable prediction is that the strategy profile (confess, confess) will be played. Now consider the game of chicken. In this case the best response correspondences are:

$$\begin{aligned} b_1(\text{swerve}) &= \{\tilde{}\ \text{swerve}\} \\ b_1(\tilde{}\ \text{swerve}) &= \{\text{swerve}\} \\ b_2(\text{swerve}) &= \{\tilde{}\ \text{swerve}\} \\ b_2(\tilde{}\ \text{swerve}) &= \{\text{swerve}\}. \end{aligned}$$

In this case neither player has a dominant strategy. If player i is playing "swerve" then player j will want to play "˜ swerve" and if player i is playing "˜ swerve" then player j will want to play "swerve" . So what is a reasonable prediction for this game? If we just seek to have strategy profiles in wich each player is playing a best response to the other, then two such profiles exist: (swerve,˜ swerve) and (˜ swerve,swerve). The formalization of this equilibrium concept is as follows.

Definition 36 *A Nash equilibrium (in pure strategies) to a normal form game is a strategy profile* (s^*) *satisfying the condition: for every* $i \in N$

$$s_i^* \in b_i(s_{-i}^*)$$

This condition can be restated as:

Definition 37 *A Nash equilibrium (in pure strategies) to a normal form game is a strategy profile* (s^*) *satisfying the condition: for every* $i \in N$

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i', s_{-i}^*) \text{ for every } s_i' \in S_i.$$

Exercise 20 *Verify that the previous two definitions are equivalent. Hint show that if a strategy profile satisfies the first then it must satisfy the second, and then that if it satisfies the second it must satisfy the first.*

The concept of a Nash equilibrium (NE) is deceptively simple. We require that players **correctly conjecture what the other players will do and that they play a best response to this conjecture**. An alternative interpretation is that at a strategy profile which is a NE **no player has an incentive to unilaterally change her strategy**. As the examples of the PD and chicken demonstrate in very simple games Nash equilibria may be unique or non-unique. In the fourth example (the reduced form election) it is not difficult to see that the unique NE is $(s_1^* = 0, s_2^* = 0)$. To see this we note that for $i \in \{1, 2\}$ and $j \in \{1, 2\} \backslash \{i\}$ the best response correspondence is:

$$b_i(s_j) = \begin{cases} \{x \in X : |x| < |s_j|\} \text{ if } |s_j| > 0 \\ \{0\} \text{ if } |s_j| = 0 \end{cases}.$$

It is clear that since $|0| = 0$ $b_i(0) = 0$ which means that $(s_1^* = 0, s_2^* = 0)$ is a NE. To see that no other NE exists consider an arbitrary profile (s_1, s_2) and without loss of generality assume that $|s_1| \geq |s_2| \geq 0$ with at least one of the inequalities strict. This means that $s_1 \notin b_1(s_2)$ and so (s_1, s_2) cannot be a NE. In other words at least one player will have an incentive to unilaterally change her strategy to a number closer to 0.

Analysis of the third example demonstrates that NE may not exist. In this game the best response correspondences are:

$$\begin{aligned} b_1(\text{one}) &= \{\text{two}\} \\ b_1(\text{two}) &= \{\text{one}\} \\ b_2(\text{one}) &= \{\text{one}\} \\ b_2(\text{two}) &= \{\text{two}\}. \end{aligned}$$

Here no strategy profile can satisfy the required condition. If the players are using different strategies player 2 will want to deviate. If the players are using identical strategies player 1 will want to deviate. These example demonstrate that the following remark is true:

Remark 14 *In an arbitrary Normal form game with finite strategy spaces: (1) there may be a unique NE; (2) there may be multiple NE; (3) there may be no NE.*

In applications one generally defines a game and seeks to characterize the set of NE (or some other set of strategy profiles). Accordingly in characterizing the equilibrium set, one is interested in existence and uniqueness. From the perspective of an applied scholar it is generally the case that a unique NE is most desirable as it means that we are analyzing a well specified model that makes clean predictions. The case of multiple equilibria is less desirable as it may mean that the model yields ambiguous predictions. The case of no equilibria may be very unsatisfactory as the model makes no predictions, and therefore does not limit our expectations of possible phenomena.

One solution is to focus only on games in which the strategy spaces are convex and compact subsets of $\mathbb{R}^d$ and the utility functions are continuous and strictly quasi concave mappings $u_i(s) : S \to \mathbb{R}^1$. We have not yet defined strict quasi concavity of functions but it is analogous to strict convexity of preferences in an important way.

Definition 38 *A function $f(x) : X \to \mathbb{R}^1$ with X a convex set is strictly quasi concave if for any $t \in \mathbb{R}^1$, $x \neq y \in X$ and $\lambda \in (0, 1)$ with $f(x) \geq t$ and $f(y) \geq t$ it is the case that $f(\lambda x + (1 - \lambda)y) > t$.*

Exercise 21 *Assume X is convex and $u : X \to \mathbb{R}^1$ represents the weak ordering R on X. Then R is strictly convex iff $u(\cdot)$ is a strictly quasi concave function.*

As we have already shown strictly convex preferences have singleton maximal sets (when the sets are non empty). Given the above exercise, this means that in a game having S a convex set and utility functions $u_i(s_i, s_{-i})$ that are strictly quasi concave in s_i for each $s_{-i} \in S_{-i}$ the best response correspondence will be a function (a single valued correspondence).

In the next subsection we prove the following result.

Theorem 15 *If the normal form game $\langle N, S, u\rangle$ satisfies the following conditions:*

(1) S_i is a convex and compact subset of a Euclidean space for each $i \in N$.

(2) $u_i(s_i, s_{-i}) : S \to \mathbb{R}^1$ is a continuos function for each $i \in N$

(3) for every $i \in N$ and every $s'_{-i} \in S_{-i}$ the function $u_i(s_i, s'_{-i}) : S_i \to \mathbb{R}^1$ is strictly quasi concave

a NE exists.

A second solution to the problem of not having NE is to modify the game to allow players to randomize over their pure strategies. Accordingly, we define the mixed extension of a normal form game as follows.

Definition 39 *Given a normal form game $\Gamma = \langle N, S, u\rangle$ the mixed extension game $\Gamma^m = \langle N, \Delta, u^m\rangle$ is constructed as follows: $\Delta_i = \Delta(S_i)$ with an arbitrary strategy $\sigma_i \in \Delta_i$ for all $i \in N$ and $\Delta = \times_{i \in N}\Delta_i$ with $\sigma_i(s_i)$ denoting the probability that mixed strategy σ_i assigns to pure strategy s_i. The expected utility function, $U_i(\sigma_i, \sigma_{-i}) : \Delta \to \mathbb{R}^1$ is defined as*

$$U_i(\sigma_i, \sigma_{-i}) = \sum_{s_{-i} \in S_{-i}} \sum_{s_i \in S_i} u_i(s_i, s_{-i})\sigma_i(s_i), \sigma_{-i}(s_{-i}) \text{ for all } i \in N.$$

For a given normal form game the mixed extension allows players to mix over their pure strategies, and utility is calculated by taking expectations of utility over pure strategies when the expectation is taken with respect to the lotteries that agents play – their mixed strategies. Since the mixed extension of a normal form game is itself a normal form game, our definition of NE applies directly to mixed extensions. Nash demonstrated that any finite strategy normal form game has a mixed strategy NE.

Theorem 16 *(Nash) Given a normal form game $\Gamma = \langle N, S, u\rangle$ in which S is finite, the mixed extension $\Gamma^m = \langle N, \Delta, u^m\rangle$ has at least one NE. In otherwords every finite game has a mixed (possibly degenerate) strategy NE.*

The proofs of the last two theorems hinge on fixed point theorems due to Brouwer and Kakutani and we devote a subsequent * ed subsection to proving these results.

8.3 Calculating NE

8.3.1 Pure strategy NE in finite games

Given a finite game, one way to characterize all of the pure strategy NE is to test whether each profile $s' \in S$ is a NE. To do this one starts with a profile $s' = (s'_1, ..., s'_n)$ and asks the following sequence of questions:

1. Holding $s'_2, ..., s'_n$ fixed is there a strategy s''_1 for which $u_1(s''_1, s'_{-1}) > u_1(s'_1, s'_{-1})$. If so then s' is not a NE. If not then continue

2. Holding $s'_1, s'_3..., s'_n$ fixed is there a strategy s''_2 for which $u_2(s''_2, s'_{-2}) > u_2(s'_2, s'_{-2})$. If so then s' is not a NE. If not then continue

.

.

.

i. Holding $s'_1, ...s'_{i-1}, s'_{i+1}.., s'_n$ fixed is there a strategy s''_i for which $u_i(s''_i, s'_{-i}) > u_i(s'_i, s'_{-i})$. If so then s' is not a NE. If not then continue

.

.

.

n. Holding $s'_1, ..., s'_{n-1}$ fixed is there a strategy s''_n for which $u_n(s''_n, s'_{-n}) > u_n(s'_n, s'_{-n})$. If so then s' is not a NE. If not then s' is a NE.

This algorithm is then repeated for each profile in S.

In two player finite games –which are representable by matrices– the algorithm is particularly straightforward. Start with a profile (matrix entry) and see whether there is an entry in the same column that makes the row player better off. If so then the original profile is not a NE. If not then repeat the exercise interchanging row and column.

Example 5.

Consider the following example

$$\begin{bmatrix} \text{player 1\textbackslash player 2} & l & c & r \\ t & 5,4 & 2,3 & 6,2 \\ m & 2,5 & 3,6 & 5,5 \\ b & 5,2 & 0,3 & 7,4 \end{bmatrix}$$

We begin by conjecturing that (t, l) is a pure strategy NE. Note that player 1 can only affect the row choice. Given that $s_2 = l$ player 1 chooses between $u_1(t, l) = 5$, $u_1(m, l) = 2$ and $u_1(b, l) = 5$. Accordingly $b_1(l) = \{t, b\}$. Given $s_1 = t$ player 2 chooses between utilities of 4, 3 and 2 so $b_2(t) = \{l\}$. Accordingly (t, l) is a NE. Since $b_2(l)$ is a singleton we know that this is the only pure strategy NE in which t is played. Recall that $b_1(l) = \{t, b\}$ so (m, l) is not a NE as player 1 would deviate to either t or b if she anticipated player 2 selecting l. Now we conjecture that (m, c) is a NE and note that $b_1(c) = \{m\}$ and $b_2(m) = \{c\}$ and thus our conjecture is correct. since $r \notin b_2(m)$ we note that the only pure strategy NE in which m is played is (m, c). Now if we conjecture that (b, l) is a NE we will note that $b_2(b) = \{r\}$ and thus we see that our conjecture is incorrect. If we conjecture that (b, c) is a NE we note that $b_1(c) = \{m\}$ and so our conjecture is incorrect. Finally conjecturing that (b, r) is a NE we observe that neither player has an incentive to deviate and that the conjecture is correct. We thus conclude that the set of pure strategy NE is $\{(t, l), (m, c), (b, r)\}$.

8.3.2 Mixed strategy NE in finite games

To characterize the mixed NE to a finite game the procedure is somewhat different. A key piece of intuition is that in a mixed strategy NE player $i \in N$ must be indifferent between playing any of the pure strategies in the support of her mixed strategy σ_i (recall that pure strategy s_i is in the support of mixed strategy σ_i if $\sigma_i(s_i) > 0$). To demonstrate this logic reconsider the "once, twice, three shoot" game of example 3.

$$\begin{bmatrix} \text{player 1}\backslash\text{ player 2} & \text{one} & \text{two} \\ \text{one} & 0,1 & 1,0 \\ \text{two} & 1,0 & 0,1 \end{bmatrix}.$$

Let σ_i denote the probability that player i plays pure strategy one and $1-\sigma_i$ denote the probability that player i plays two. If $\sigma_i \in (0, 1)$ in a pure strategy NE it must be the case that

$$\begin{aligned} \sigma_2 u_1(one, one) + (1 - \sigma_2)u_1(one, two) = \\ \sigma_2 u_1(two, one) + (1 - \sigma_2)u_1(two, two). \end{aligned} \tag{1}$$

If this condition were not true then player one would have an incentive to deviate from the strategy in which $\sigma_1 > 0$ and $1 - \sigma_1 > 0$ to one in which one of these probabilities were equal to 1. A subtle point to note is that equating player $1'$s expected utility over the support of her mixed strategy imposes a condition on the mixtures that player 2 is using. In other words in a mixed strategy NE each players mixtures are sufficient to make the other players indifferent between the pure strategies in the support of their own mixed strategy. This observation allows us to setup a system of equations whos solution is a mixed strategy NE. Specifically, lets conjecture that there is a mixed strategy NE to the above game in which all strategies are played with positive probability. Such a strategy profile is often termed **completely mixed**. In such an equilibrium the equation above, as well as the analogous equation for player 2 to be indifferent must be satisfied. Substituting in the values of the utility functions we attain the following system.

$$\begin{aligned} (1 - \sigma_2) &= \sigma_2 \\ \sigma_1 &= (1 - \sigma_1) \end{aligned}$$

It is clear that this system has a unique solution $\sigma_i = \frac{1}{2}$ for each $i \in N$.

We now consider a generalization of this game.

$$\begin{bmatrix} \text{player 1}\backslash\text{ player 2} & \text{one} & \text{two} \\ \text{one} & a, b & c, d \\ \text{two} & e, f & g, h \end{bmatrix}.$$

In fact varying the parameters $a, b, ..., h$ allows us to at one time characterize the mixed NE to all normal form games with 2 players and 2 strategies each. We will first focus on the completely mixed NE. The relevant system is

$$\sigma_2 a + (1 - \sigma_2)c = \sigma_2 e + (1 - \sigma_2)g$$
$$\sigma_1 b + (1 - \sigma_1)f = \sigma_1 d + (1 - \sigma_1)h.$$

Thus given the payoffs $a, b, ..., h$ if $(\sigma_1, \sigma_2) \in (0, 1)^2$ solve this system they are a mixed strategy NE to the game.

Exercise 22 *Using this system characterize the mixed strategy NE to the PD, and chicken (examples 1 and 2) above.*

Exercise 23 *Extend this logic to games with larger strategy spaces and characterize the mixed strategy NE to example 5.*

8.3.3 Pure strategy NE in non finite games

In games where the strategy space is not finite, the algorithm exhibited above will not work. In most applications one assumes that the functions $u_i(s)$ are twice differentiable. In this case the best response functions can be characterized by the following algorithm:

1. For each $i \in N$ take the derivative $\frac{\partial u_i(s_i, s_{-i})}{\partial s_i}$ and set this expression equal to 0. This equation is termed a first order condition (FOC). If S_i has more than one dimension then the term $\frac{\partial u_i(s_i, s_{-i})}{\partial s_i}$ is actually a vector (where each coordinate is the partial derivative with respect to one coordinate of s_i) and the quantity 0 denotes the vector of 0's.

2. Solve the system

$$\frac{\partial u_i(s_i, s_{-i})}{\partial s_i} = 0$$

for s_i as a function of s_{-i}. If second order conditions are satisfied (for a maximum) then $b_i(s_{-i})$ is equivalent to this function relating s_i to s_{-i}.

3. Repeat this procedure for each $i \in N$.

4. Find a simultaneous solution to the system

$$s_1^* = b_1(s_{-1}^*)$$
$$.$$
$$.$$
$$.$$
$$s_i^* = b_i(s_{-i}^*)$$
$$.$$
$$.$$
$$.$$
$$s_n^* = b_n(s_{-n}^*).$$

Example 6.

To demonstrate this logic we consider a simple interest group game. Two interest groups $N = \{1, 2\}$ want to influence a government policy. Both groups know that the final policy will be a function of how much support they give to the government. The first groups most preferred policy is 0 and the second groups most preferred policy is 1. The government favors the policy $\frac{1}{2}$ but is willing to sell out for contributions. Each group can contribute an amount $s_i \in [0, 1]$ and the final policy is given by $x(s_1, s_2) = \frac{1}{2} - s_1 + s_2$. So both groups contribute to the government simultaneously and then the government enacts the policy given by the function $x(s_1, s_2)$. The government of course keeps all of the contributions to buy advertisements for the next election. This example is a form of an all pay auction over a divisible good. We assume that the interest groups each have utility functions over their contribution and the final policy of the form:

$$u_1(s_1, s_2) = -(x(s_1, s_2))^2 - s_1$$
$$u_2(s_1, s_2) = -(1 - x(s_1, s_2))^2 - s_2$$

Substituting the policy function into the utility functions we attain:

$$u_1(s_1, s_2) = -(\frac{1}{2} - s_1 + s_2)^2 - s_1$$
$$u_2(s_1, s_2) = -(1 - (\frac{1}{2} - s_1 + s_2))^2 - s_2$$

The first order conditions are given by differentiation:

$$FOC_1 : 2(\frac{1}{2} - s_1 + s_2) - 1 = 0$$
$$FOC_2 : 2(1 - (\frac{1}{2} - s_1 + s_2)) - 1 = 0$$

Solving FOC_1 yields the best response correspondence

$$b_1(s_2) = s_2$$

Solving FOC_2 yields the best response correspondence

$$b_2(s_1) = s_1$$

Now the set of pure strategy NE to this game is uncountable and it is given by:

$$\{(s_1, s_2) \in [0,1]^2 : s_1 = s_2\}$$

This result has a very straightforward interpretation. Any pair of equivalent contributions is a NE. The resulting policy from such a profile of contributions is $\frac{1}{2}$. No contributor wants to unilaterally deviate because the marginal gain of an additional unit of contribution (in terms of pulling policy in the desirable direction) is exactly offset by the marginal cost of loosing another unit of resources. As in the PD NE of this game are inefficient. That is, the contributors would rather commit to not giving any money to the government. However, since no such commitment is possible in the game, the fact that each contributor has an incentive to deviate from such an agreement means that this outcome is not supportable.

8.4 Dominance

Nash equilibria require that agents play strategies from which no agent has an incentive to unilaterally deviate. Returning to voting theory let us consider the extent to which this requirement pins down behavior.

Example 7:

Suppose $N = \{1, 2, 3\}$, $S_i = \{u, d\}$ where the former denotes a vote to go up to heaven and the latter denotes a vote to go down to hell. Let preferences be represented by the utility functions

$$u_i(s) = \begin{cases} 1 \text{ if } s \in \{(u,u,u), (u,u,d), (d,u,u), (u,d,u) \\ 0 \text{ otherwise} \end{cases}$$

for $i \in N$. This game can be interpreted as a decision by three individuals about where to spend eternity in which the decision is made by majority rule. It will come as no surprise that a pure strategy NE in which $s_i = u$ for all $i \in N$ exists and that this equilibrium is **Pareto Efficient**.[4] It may also not be surprising that as an example the profile (u, u, d) is a NE. This is true because such a profile gives every one utility 1 and this is the highest possible utility an agent can receive. What is troubling is the fact that (d, d, d) is also a NE. In this game there is perfect alignment of preferences and we see that the least desirable of the two outcomes is possible. While each player would prefer to

[4]The normative condition Pareto Efficiency has the following definition. A strategy profile $s \in S$ is Pareto Efficient if there is no other strategy profile $s' \in S$ for which $u_i(s') \geq u_i(s)$ for every $i \in N$ and the inequality is strict for some $i \in N$.

go up, if player 1 correctly conjectures that players 2 and 3 are going to vote to go down, then player 1 will be indifferent between voting u or d. Accordingly voting d is a best response. While such a profile is a NE, it seems altogether unreasonable. In this game it seems entirely unreasonable for player 1 to resolve her indifference in this manner. Accordingly we will consider a refinement to NE which rules out this type of equilibria.

While the discussion of example 7 concluded that the notion of NE may be too weak, in other settings it may be to strong or unreasonable. Consider the following game.

Example 8:

player 1\player 2	l	c	r
t	$1,4$	$3,2$	$6,2$
m	$2,5$	$7,6$	$9,5$
b	$5,8$	$1,7$	$10,4$

In this example it is not difficult to see that the set of pure strategy NE is $\{(b,l),(m,c)\}$. But neither of these strategies is naturally more reasonable then the other. One may ask what happens if player 1 anticipates that the first equilibrium will be played and player 2 anticipates that the second equilibrium will be played. The result would then be (b,c) but this is not even a NE. Since a NE requires best responses, and a selection of mutual best response from a possible large set of such mutual best responses the requirement that players know which profile is being played (without communication) is demanding. Accordingly we may ask what predictions can be made if one does not assume that agents will coordinate on an equilibrium. One weak notion of rationality in normal form games is to require players to not use strategies which are dominated by some other mixed strategy.

Definition 40 *A pure strategy $s_i \in S$ is strictly dominated if there exists a $\sigma_i' \in \Delta(S_i)$ s.t.*

$$U_i(\sigma_i', s_{-i}) > u_i(s_i, s_{-i}) \text{ for every } s_{-i} \in S_{-i}.$$

The strategy s_i is weakly dominated if there is a σ_i' for which the inequality holds weakly for every $s_{-i} \in S_{-i}$ and strictly for some $s_{-i}' \in S_{-i}$.[5]

In example 8 t is strictly dominated by m and r is weakly dominated by c. Accordingly if we just require that players not use weakly dominated strategies we can rule out profiles with either t or m. In example 7 a vote for d is weakly dominated, and if we seek pure strategy NE involving strategies that are not weakly dominated (i.e. a pure strategy NE involving weakly dominant strategies) then the unique prediction is (u,u,u). In the PD of example 1, we do not need to seek NE to make a unique prediction. The only profile of strategies

[5]Recall that $U_i(\sigma_i', s_{-i}) = \sum_{s_i' \in S_i} u_i(s_i', s_{-i})\sigma_i(s_i)$.

that does not involve a strictly dominate strategy is $(confess, confess)$. In this particular game the weaker condition of excluding strictly dominated strategies, and the stronger condition of requiring a fixed point to the best response correspondence coincide. In example 2, however no strategy is strictly (or weakly) dominated and thus this requirement does not pin behavior down very much. A solution concept that is sometimes useful is attained by iteratively removing strictly dominated strategies. Such a procedure is defined as follows:

Definition 41 *Given a normal form game* $\Gamma^0 = \langle N, S^0, u^0 \rangle$ *the process of iteratively deleting strictly dominated strategies is attained through the following algorithm: for* $t = 1, 2, ...$

In period t arbitrarily select a player $i^t \in N \backslash i^{t-1}$ *and remove from* S_i^{t-1} *each strategy that is strictly dominated in the game* Γ^{t-1} *Call the set of strategies that survive* S_i^t. *Let* $S_j^t = S_j^{t-1}$ *for* $j \in N \backslash i^t$ *and let* u_z^0 *be the restriction of* u_z *to* S^t *for each* $z \in N$.

If at τ *there is no* $i^\tau \in N$ *having a strictly dominated strategy in the game* $\Gamma^{\tau-1}$ *then call the set* $S^{\tau-1}$ *the set of outcomes that survive iterative deletion of strictly dominated strategies.*

It can be shown that regardless of what sequence of players is chosen the same set of actions will be reached. A justification for this procedure is to consider agents who reason in the following manner.

> I know that my opponents will not use strictly dominated strategies, and I know that my opponents know that I will not use strictly dominated strategies. Given this we are all really choosing from the smaller strategy space that survives the first n iterations. But I know that my opponents will not use a strategy that is strictly dominated in this game, and I know that my opponents know that I wont play a strategy that is strictly dominated in this new game,.............*ad infinitum.*

We can relate this procedure to the set of mixed strategy NE for finite games.

Theorem 17 *Given a finite normal form game with a mixed strategy NE* σ^* *if the strategy* s_i *is in the support of* σ_i^* *then it survives itereted deletion of strictly dominated strategies.*

Example 6 *Prove the theorem.*

Theorem 18 *If the set of outctomes that survive iterative deletion of strictly dominated strategies is a singleton, then this profile is the unique pure strategy NE, and the unique mixed strategy NE places probability 1 on playing this profile.*

Exercise 24 *Prove the theorem.*

In applications the characterization of mixed strategy NE is usually simplified by applying iterative deletion of strictly dominated strategies and then using the theorem to characterize the mixed strategy NE of the smaller game. In principal one may consider applying an algorithm that deletes weakly dominated strategies. This procedure will not necessarily lead to the same irreducible set for all sequences of players. In other words the result may depend on which player's strategy set is reduced first. Accordingly this procedure is less defensible. In practice weak dominance is used to refine NE and strict dominance is used to either simplify the characterization of NE or make predictions without appeal to the stronger conditions of NE.

8.5 Proving the existence of NE*

Given a correspondence. $c : A \to A$ a fixed point $x^* \in A$ is a point s.t. $x^* \in c(x^*)$. If $c(\cdot)$ is a function then a fixed point is a point x^* s.t. $x^* = c(x^*)$. We will see that it is possible to reduce the question of whether a NE exists to the question of whether a specific correspondence has a fixed point. This should not be surprising as a NE involves a strategy profile s^* for which $s_i^* \in b_i(s_{-i}^*)$ for every $i \in N$. Thus if we consider

$$b(s) = (b_1(s_{-1}), ..., b_i(s_{-i}), ..., b_n(s_{-n}))$$

which is a correspondence $b : S \to\to S$ then a NE is a fixed point of the correspondence $b(\cdot)$. A fixed point theorem is a statement of the form: if correspondence $c : A \to\to A$ satisfies particular conditions then a fixed point exists. The following conditions will be of interest to us:

Definition 42 *A correspondence $c : A \to\to A$ is convex valued if for every $a \in A$ $c(A)$ is a convex subset of A.*

In other words a correspondence $c(\cdot)$ is convex valued if for every $x \in A$ if $y, z \in c(x)$ then for any $\lambda \in [0, 1]$ the point $\lambda y + (1 - \lambda)z \in c(x)$.

Definition 43 *For a correspondence $c : A \to\to A$ the upper inverse of a set $B \subseteq A$, is $c^+(B) = \{x \in A : c(x) \subset B\}$.*

So the upper inverse of a set B is the set of points in the domain which the correspondence maps into subsets of B.

Definition 44 *A correspondence $c : A \to\to A$ is upper hemi continuous if for every open set $O \subseteq A$ the set $c^+(O)$ is open.*

Definition 45 *A correspondence $c : A \to\to A$ has a closed graph if for any two sequences $x^n \to x \in A$ and $y^n \to y \in A$ with $x^n \in A$ and $y^n \in c(x^n)$ for every n we have $y \in c(x)$.*

Theorem 19 *If A is compact the a correspondence $c : A \to\to A$ is upper hemi continuos if it has a closed graph.*

Exercise 25 *Prove the theorem.*

So a correspondence is upper hemi continuous if the upper inverse of an open set is open. When the image of the correspondence is compact having a closed graph implies upper hemi continuity. The intuition behind the closed graph condition is not difficult to see. When a correspondence has a closed graph, if we have two sequences x^n and y^n of points each in A that converges to x and y both in A with $y^n \in c(x^n)$ it must be the case that $y \in c(x)$. In other words for any sequence in the domain converging to x and any selection of points y^n that are in the image of the first sequence which are converging to a point y it must be the case that the limit y is in the image of the correspondence evaluated at the limit of the first sequence, x. Graphically for a correspondence that has a closed graph the set $\{(x, y) \in A^2 : y \in c(x)\}$ is closed in the space A^2. For a more complete treatment of these concepts see Border (1985).

It is not difficult to see that if the correspondence $c : A \to\to A$ is actually single-valued for every $a \in A$ then $c(\cdot)$ is a function. If a single valued correspondence is upper hemi continuos then it is also a continuos function.

Exercise 26 *Show that an upper hemi continuos correspondence that is single valued is a continuous function.*

To establish theorem 15 we will use the following fixed point theorem

Theorem 20 *(Brouwer) Suppose $A \subset R^d$ is a compact and convex set. If $f : A \to A$ is a continuos function then $f(\cdot)$ has a fixed point in A.*

To establish theorem 16 we will use the following fixed point theorem.

Theorem 21 *(Kakutani) Suppose that $A \subset R^d$ is a compact and convex set with $c : A \to\to A$ a correspondence satisfying the conditions:*
(1) $c(x)$ is non-empty for every $x \in A$
(2) $c(\cdot)$ is convex valued
(3) $c(\cdot)$ is upper hemi continuos
then $c(\cdot)$ has a fixed point in A.

Several proofs of these results appear in Border. In order to establish the existence of Nash equilibria in either mixed strategies or pure strategies when the appropriate assumptions are satisfied, we will need to show that in the case of theorem 15 $b(s)$ is a continuos function and in the case of theorem 16 $b(s)$ is a correspondence that satisfies the conditions 1-3 in Kakutani's fixed point theorem. A result that is useful in its own right, as well as helpful in demonstrating that $b(s)$ is non-empty and upper hemi continuos is the Theorem of the Maximum.

Theorem 22 *(Theorem of the Maximum) Let $X \subset \mathbb{R}^d, M \subset \mathbb{R}^z$ be compact and convex sets. Let $f(x, m) : X \times M \to R^1$ be continuos in x and m then the correspondence $c : M \to\to X$ defined as*

$$c(m) = \arg\max_{x \in X}\{f(x, m)\}$$

is non-empty and upper hemi continuos.

The fact that the set of optimal choices is non-empty is interesting in its own right. This result was stated in a previous section. The fact that the correspondence $c(\cdot)$ defined in the theorem is upper hemi continuos has the following interpretation. Calling the vector m a parameter vector of the optimization problem, if we consider a sequence of parameter vectors m^n converging to m then for any selection of optimal policies $x^n \in c(m^n)$ that converge to x it will be the case that $x \in c(m)$.

We can now prove theorem 15.

> **Proof of theorem 15:** Assume that: S_i is a convex subset of $\mathbb{R}^d$ (for some d) for each $i \in N$ and for each $i \in N$, $u_i(s) : S \to \mathbb{R}^1$ is continuos and for each $s'_{-i} \in S_{-i}$, $u_i(s_i, s'_{-i})$ is strictly quasiconcave in s_i. By the Theorem of the Maximum the correspondence $b_i(s_{-i}) : S_{-i} \to\to S$ defined as
>
> $$b_i(s_{-i}) = \arg\max_{s_i \in S_i}\{u_i(s_i, s_{-i})\}$$
>
> is non-empty and upper hemi continuos. By theorem 4 and exercise 22 $b_i(s_{-i})$ is a singleton for every $s_{-i} \in S_{-i}$. Combining these facts yields the conclusion that $b_i(s_{-i})$ is a continuos function from S_{-i} into S_i for each $i \in N$. We now construct the function
>
> $$b(s) : S \to S$$
>
> by defining $b(s_1, ..., s_n) = (b_1(s_{-1}), ..., b_n(s_{-n}))$. Since $s_{-i}(s)$ is a projection it is continuos. Since $b_i(\cdot)$ is continuos and the composition of continuos functions is continuos, and the product of continuos functions is continuos in the product space the function $b(s)$ is continuos. By Brouwer's fixed point theorem this mapping has a fixed point, $s^* = b(s^*)$. But this means that for every $i \in N$, $b_i(s^*_{-i}) = s^*_i$ and thus s^* is a NE.■

The proof of theorem 16 is similar. We will show that for any finite game Γ in the mixed extension game Γ^m the best response correspondence satisfies the conditions of Kakutani's fixed point theorem.

> **Proof of theorem 16:** Assume that in the game Γ, S_i is finite for each $i \in N$. This implies that in the mixed extension Γ^m Δ_i is a compact and convex subset of a finite dimensional Euclidean space. By definition 37 we can see that $U(\sigma_i, \sigma_{-i})$ is linear and therefore continuos in σ. Letting $b_i(\sigma_{-i}) : \Delta_{-i} \to\to \Delta_i$ be defined as

$$b_i(\sigma_{-i}) = \arg\max_{\sigma_i \in \Delta_i} \{U(\sigma_i, \sigma_{-i})\}$$

the Theorem of the Maximum implies that this correspondence is non-empty for every $\sigma_{-i} \in \Delta_{-i}$ and upper hemi continuos. Since $U(\sigma_i, \sigma_{-i})$ is linear for any σ_{-i} and any two σ_i', σ_i'' if $U(\sigma_i', \sigma_{-i}) = U(\sigma_i'', \sigma_{-i})$ we have $U(\lambda\sigma_i' + (1-\lambda)\sigma_i'', \sigma_{-i}) = U(\sigma_i', \sigma_{-i})$ so $b_i(\sigma_{-i})$ is convex valued. Combining these facts we see that the correspondence

$$b(\sigma) : \Delta \to \Delta$$

defined as $b(\sigma_1, ..., \sigma_n) = (b_1(\sigma_{-1}), ..., b_n(\sigma_{-n}))$ satisfies the requirements of Kakutani's fixed point theorem. Thus there is a mixed strategy profile satisfying the condition $\sigma^* \in b(\sigma^*)$. Such a profile is a NE to Γ^m and thus a mixed strategy NE to Γ.■

8.6 Bayesian Games

In the previous section we developed the concepts of normal form games and equilibrium concepts. The games considered involved no uncertainty. More precisely the payoffs of the players were common knowledge. Suppose however, in the setting of example 6 above (the inteterst groups) that each group was unsure of the marginal value of policy to its opponent. How would each firm behave in the presence of such uncertainty? How would we model such a setting?

In this section we develop tools to analyze richer models involving agents that do not know with certainty the payoffs of the other players. This feature is termed **incomplete information**. The standard practice (originated by Harsanyi 1967-68) is to convert such a game into one where a fictional player (nature) moves first drawing the utility functions of the agents from a distribution that is known to the players. Following this draw agents knowing only their own draw simultaneously select their actions. This approach is called one of **imperfect information.**

We now modify our basic normal form structure Γ by adding to it lotteries over utility functions.

Definition 46 *A Bayesian Game with perfect private signals is a tuple* $\Gamma^B = \langle N, \Theta, F(\theta) S_i, u_i(s;\theta_i)\rangle$ *in which* $\Theta = \prod_{i\in N}\Theta_i$ *is the space of player type profiles,* $\Theta_i \subset \mathbb{R}^{d_i}$ *is the space of types for player* i*,* $F(\theta)$ *is the cumulative joint prior distribution on type profiles* θ*, and for each* θ_i*,* $u_i(s;\theta_i)$ *is a Bernoulli utility function for player* i *over strategy profiles* $s \in S$*.*

Example 9:

As a simple example consider a contributions model in which two firms $N = \{l, r\}$ simultaneously offer bribes to a president (in exchange for the provision of a non-divisible cabinet seat). Suppose the government is indifferent between selling out to either firm and thus will sell the cabinet seat to the highest bidder. Of course the president is quite moral and will only keep the bid of the winning firm. Let $\Theta_i = [0, 1]$ for $i \in N$ denote the set of types and suppose the prior probability that $\theta_i < c$ is c if $c \in [0, 1]$. Thus the prior is the uniform distribution on $[0, 1]$. The payoff to a firm that has type θ_i, offers bid s_i and is awarded the cabinet seat is $\theta_i - s_i$. The payoff to a firm that is not awarded the seat is 0. Each firm learns its own θ_i and the firms then simultaneously submit their bids to the government.

Given a Bayesian game with perfect private signals Γ^B we can extend the notion of a NE in the following manner.

Definition 47 *Given a Bayesian game with perfect private signals, Γ^B, a pure strategy Bayesian Nash equilibrium BNE is a n-tuple of mappings $z_i^* : \Theta_i \to S_i$ (1 for each $i \in N$) s.t. for each $i \in N$ and each $\theta_i \in \Theta_i$*

$$z_i^*(\theta_i) \in \arg\max_{s_i \in S_i} \left\{ \int_{\theta_{-i}} u_i(s_i, z_{-i}^*(\theta_{-i}); \theta_i) dF(\theta_{-i}). \right\}$$

Of course if Θ is a finite set then this definition can be rewritten with summations instead of integrals. When both Θ and S are finite sets the existence of mixed strategy BNE can be established by using the Nash result.

Exercise 27 *For a Bayesian game with perfect private signals, Γ^B with finite S and Θ : (1) Define an appropriate notion of a mixed strategy BNE. (2) Use Theorem 16 to establish the existence of mixed strategy BNE.*

Returning to example 9 we can characterize a symmetric pure strategy BNE. In this example a pure strategy is a mapping $s_i : [0.1] \to [0, 1]$. We start by conjecturing that firm $-i$ is using a mapping $s'_{-i}(\theta_{-i}) = y\theta_{-i}$. Then firm i faces the following objective function:

$$s_i^*(\theta_i) \in \arg\max_{s_i \in [0,1]} (\theta_i - s_i) \frac{s_i}{y}$$

The FOC is:

$$\frac{\theta_i}{y} = \frac{2s_i}{y}$$

Solving the FOC yields

$$s_i = \frac{1}{2}\theta_i$$

Thus $s_l(\theta_l) = \frac{1}{2}\theta_l$ and $s'_r(\theta_r) = \frac{1}{2}\theta_r$ is a BNE to the game.

Returning to example 6 above suppose now each interest groups utility is of the form

$$u_1(s_1, s_2) = -\theta_1(\frac{1}{2} - s_1 + s_2)^2 - s_1$$
$$u_2(s_1, s_2) = -\theta_2(1 - (\frac{1}{2} - s_1 + s_2))^2 - s_2$$

where $\theta_i \in \{1, 2\}$ and each possible combination of types occurs with equal probability. Now a strategy is a mapping from $\{1, 2\}$ into $[0, 1]$ assinging a bid to each possible type. Given a strategy for group 2, group 1 solves the optimization problem:

$$\max_{s_1 \in [0,1]} \left\{ \begin{array}{l} -\frac{1}{2}\left[\theta_1(\frac{1}{2} - s_1 + s_2(\theta_2 = 1))^2 - s_1\right] \\ -\frac{1}{2}\left[\theta_1(\frac{1}{2} - s_1 + s_2(\theta_2 = 2))^2 - s_1\right] \end{array} \right\}$$

Note that player 1 actually solves 2 optimization problems, one when her type is $\theta_1 = 1$ and one when her type is $\theta_1 = 2$. Letting s_i^j denot player i's bid when her type is j, and solving first order conditions for each player yields the system of equations:

$$b_1(s_2^1, s_2^1; \theta_1 = 1) = \frac{1}{2}s_2^1 + \frac{1}{2}s_2^2$$
$$b_1(s_2^1, s_2^1; \theta_1 = 2) = \frac{1}{3} + \frac{1}{2}s_2^1 + \frac{1}{2}s_2^2$$
$$b_2(s_1^1, s_1^1; \theta_2 = 1) = \frac{1}{2}s_1^1 + \frac{1}{2}s_1^2$$
$$b_2(s_1^1, s_1^1; \theta_2 = 2) = \frac{1}{4} + \frac{1}{2}s_1^1 + \frac{1}{2}s_1^2$$

Inspection of this system suggests that it has no solution. For any pair $s_1^1 < s_1^2$ player 2 will want submit bids that are higher than s_1^1, but then player 2 would respond to this by raising her bids. Note however that this does not mean there is not a pure strategy BNE. In differentiating and analyzing first order conditions, we neglected the fact that agents face the constraint $s_i \in [0, 1]$. Thus a corner solution ($s_i^j \in \{0, 1\}$) does not need to satisfy the system above. We may conjecture (based on the logic that players want to outbid each other) that $s_i^j = 1$ for all i,$j \in \{1, 2\}$ might be a BNE. The argument for this fact is left as an exercise.

Exercise 28 *Demonstrate that* $s_i(\theta_i) = 1$ *for all* i, j *is a BNE.*

In the last two examples conditional on knowing her type, and the actions taken by the other players, a agent faced no uncertainty. If we allow for the possibility that an agent's preferences may not be perfectly revealed by her type then we need a slightly richer class of games.

Definition 48 *A Bayesian Game with imperfect private signals is a tuple* $\Gamma^B = \langle N, \Omega, F(\omega), \Theta, F(\theta \mid \omega), S_i, u_i(s; \theta_i, \omega)\rangle$ *in which* $\omega \in \Omega$ *is an unknown state* *variable with prior distribution,* $F(\omega)$, $\Theta = \prod_{i \in N} \Theta_i$ *is the space of player type profiles,* $\Theta_i \subset \mathbb{R}^{d_i}$ *is the space of types for player* i, $F(\theta \mid \omega)$ *is the cumulative joint prior distribution on type profiles* θ *conditional on state* $\omega \in \Omega$, *and for each* (θ_i, ω), $u_i(s; \theta_i, \omega)$ *is a Bernoulli utility function for player* i *over strategy profiles* $s \in S$.

The difference between this class of games and the former, is the existence of a state variable ω which effects the players utility functions and the distribution of types.

Definition 49 *Given a Bayesian game with imperfect private signals* Γ^B, *a pure strategy Bayesian Nash equilibrium BNE is a n-tuple of mappings* $s_i^* : \Theta_i \to S_i$ *(1 for each* $i \in N$*) s.t. for each* $i \in N$ *and each* $\theta_i \in \Theta_i$

$$s_i^*(\theta_i) \in \arg\max_{s_i \in S_i} \left\{ \int_\omega \int_{\theta_{-i}} u_i(s_i, s_{-i}^*(\theta_{-i}); \theta_i, \omega) dF(\theta_{-i}) dF(\omega). \right\}$$

To demonstrate the concepts we consider an example of the model Austen-Smith and Banks (199x) analyze.

Example 2:

Consider three jurors $N = \{1, 2, 3\}$ responsible for deciding whether to convict or acquit a defendant. They must choose $x \in \{c, a\}$. The jurors simultaneously cast ballots $v_i \in S_i = \{c, a\}$ and the outcome is chosen by majority rule. Each player faces uncertainty about whether or not the defendant is guilty, G, or innocent, I. So $\Omega = \{G, I\}$. In the former state the jurors receive utility 1 from convicting and 0 from acquitting. In the latter state the jurors receive utility 1 from acquitting and 0 from convicting. Each player assigns prior probability $\pi > \frac{1}{2}$ to the defendant being guilty. Before voting each player also receives a private signal $\theta_i \in \{0, 1\}$. We assume that $prob(\theta_i = 1 \mid \omega = G) = prob(\theta_i = 0 \mid \omega = I) = p > \frac{1}{2}$. Of course the reciprocal conditional events occur with probability $1 - p$. Voter i selects her vote to maximize $prob(x = c \text{ iff } \omega = G)$. If weakly undominated voting strategies are used, this requires that voter i cast ballot $v_i = c$ if she believes that the defendant is guilty with probability greater than a half.

In analyzing games of this form **Bayes' rule** is quite useful.

Definition 50 *Given events* A *and* B *Bayes' rule is:*

$$prob(A \mid B) = \frac{prob(A) prob(B \mid A)}{prob(B)}$$

Applying this rule to example 2, we see that the probability that $\omega = G$ conditional on having type $\theta_i = 1$ is

$$prob(G \mid \theta_i = 1) = \frac{\pi p}{\pi p + (1 - \pi)(1 - p)}$$

Similar expressions yield the other relevant conditional probabilities. An insight that Austen-Smith and Banks demonstrate is that in deciding how to vote, if voters are voting according to their signals (c if $\theta_i = 1$ and a if $\theta_i = 0$), a voter should not just condition on her private information (θ_i). But rather a voter should condition on being pivotal. In the example that means that a voter should vote as if she knows that 1 of her colleagues received a signal of 0 and 1 received a signal of 1. Thus a voter receiving a private signal $\theta_i = 1$ should use the posterior probability

$$prob(G \mid \theta_i = 1) = \frac{\pi p^2(1-p)}{\pi p^2(1-p) + (1-\pi)(1-p)^2 p}.$$

A voter receiving a private signal of $\theta_i = 0$ should use the posterior probability

$$prob(G \mid \theta_i = 0) = \frac{\pi p(1-p)^2}{\pi p(1-p)^2 + (1-\pi)(1-p)p^2}.$$

The logic behind this conclusion is elegant. When voter i compares her expected utility of voting c to her expected utility of voting a she needs to sum over the possible ballots that her colleagues cast. Since i's ballot is irrelevant when both of her colleagues vote c or a, the only time there will be a difference in the expected utilities is when the colleagues split.

Exercise 29 *Assume that in example 2 $p = \frac{3}{4}$ and $\pi = \frac{2}{3}$, chacterize the set of BNE to the game. Now assume that instead of majority rule, a version of unanimity rule is used – if all agents vote to convict the defendent is convicted, if at least one agent votes to acquit the defendent is acquited. Characterize the BNE to this game (again assuming that $p = \frac{3}{4}$ and $\pi = \frac{2}{3}$).*

A classic example relating dominance and Bayesian games is the second price auction

Example 3. Suppose n players each assign utility θ_i to a painting. Each player knows her own type θ_i and believes that the type of each other player is uniformly distributed on $[0, 1]$. The auctioneer will accept secret bids from each player. The player that submits the highest bid will receive the painting and pay the auctioneer the amount of the second highest bidders bid. Assume that the payoff to a player with type θ_i that wins and pays c is $\theta_i - c$. The payoff to an agent that does not win is 0.

Exercise 30 *Characterize the set of strategies which are not weakly dominated in the game described in example 3. Characterize the BNE to the game described in example 3.*

It is important to note that the Bayesian games considered are really just special cases of Normal form games, in which each player simultaneously selects a strategy (where a strategy is a mapping from Θ_i into S_i) and the payoffs are defined as the agents expected utility over strategy profiles. We now consider games in which there is an interesting notion of timing.

9 Game Theory II: Extensive form games of complete information

9.1 Introduction

So far the we have only considered games that do not have a notion of timing. Normal form game involve each player simultaneously choosing her strategy. But in real social interactions there tends to be an important element of timing. We begin by considering a simple normal form game:

1/2	a	b	c	d
x	$4,3$	$4,3$	$0,2$	$0,2$
y	$2,2$	$1,3$	$2,2$	$1,3$

The pure strategy NE of this game are: $\{(x,a),(x,b),(y,d)\}$.Now suppose that player 1 gets to choose x or y. Then following this move player 2 gets to choose l or r. Player 2 chooses a mapping from S_1 into S_2. Of course there are 4 such mappings. $\{(x\to l, y\to l),(x\to l, y\to r),(x\to r, y\to l),(x\to r, y\to r)\}$

[insert figure 1 here]

We can label each of these mappings a, b, c, and d. Suppose that the payoffs following any pair of strategies in this new 2 period game are given by the normal form game above. In this case the tree depicted in figure 1 captures the sequence of moves and the payoffs associated with each pair of moves. We can now ask whether all of the pure strategy NE are equally reasonable when the game has this 2 period interpretation. Now the profile (y,d) is a NE, but in the game with 2 periods, this profile is unreasonable. Note that we identified the strategy d with the mapping $(x\to r, y\to r)$. Given that 2 is playing this strategy 1 chooses y because the outcome corresponding to (y,r) is more desirable than the profile corresponding to (x,r). The interesting question is whether it is reasonable for player 1 to expect that player 2 would respond to x by playing r. Note that if player 1 has already played x then 2 chooses between the payoffs 2 and 3, and thus will prefer to play l. If player 1 is convinced that this is how player 2 would respond to strategy x then player 1 would actually prefer playing x and getting 4 instead of playing y and getting 1. It turns out that the only NE of the game which survives this type of logic is (x,b). The other NE involve strategies by player 2 which call for her to take actions which are suboptimal at the time she moves. One interpretation of how these NE are inappropriate is that we can think of players announcing their strategies (as in the normal form game) but then having the opportunity to change their mind once they reach a period in which they are required to move. In this interpretation player 2 may announce that she will play d or $(x\to r, y\to r)$. But if player 1 were to actually play x the choice of r would not be optimal, and so the announcement of $(x\to r, y\to r)$ is not credible.

In this chapter we develop a theory of games in which timing is included in the datum of a game, and consider a solution concept which captures the intuition sketched out here. We begin with the simplest case of games that involve a finite number of periods, and in each period and agent has a finite number of feasible actions available to her. Our model of such a game is called a finite extensive form game.

9.2 Extensive form games

In this section we formalize the concepts involved in writing down a game involving a notion of timing. The simple example of figure 1 should be kept in mind in studying the formalism here.

Definition 51 *A* ***finite extensive from game*** Γ^E *is tuple* $\langle N, H, p(\cdot), U\rangle$ *where* N *is the set of players,* H *is a finite set of histories satisfying the following properties:*

(1) The empty set $\emptyset \in H$*. This is called the null history.*

(2) If the vector $(a_1, a_2, ..., a_n) \in H$ *then it is also the case that for any* $k < n$, $(a_1, a_2, ..., a_{n-k}) \in H$

(3) A history $(a_1, a_2, ..., a_n) \in H$ *for which there is not some history* $(a_1, a_2, ..., a_n, a_{n+1}) \in H$ *is termed a terminal history. The set of terminal histories is denoted* H^T*.*

(4) The mapping $p(h) : H \backslash H^T \to N$ *assigns to each non-terminal history* h *an agent who must make a decision at* h*.*

(5) The utility functions U *are a list of Bernoulli utility functions* $u_i(h) : H^T \to \mathbb{R}^1$ *for each* $i \in N$*.*

The histories H can be interpreted as a tree with histories being nodes. At each node (history h) agent $p(h)$ must choose from the set of branches which emanate from the node (actions in the set $A(h) := \{a : (h, a) \in H\}$) . The tree begins at the null history. Agent $p(\emptyset)$ must make a choice from the set $A(\emptyset) = \{a : (\emptyset, a) \in H\}$. Following the choice a' by $p(\emptyset)$ agent $p(\emptyset, a')$ must choose from the set $A(\emptyset, a') = \{a : (\emptyset, a', a) \in H\}$. The game ends when some agent makes a choice a'' for which the history $(\emptyset, a', ..., a'') \in H^T$. Player i's utility is given by $u_i((\emptyset, a', ..., a''))$. Such a tree will satisfy a few simple conditions: no branches double back, meaning for every node there is a unique chain of previous choices that reaches the node; no cycles exist in the tree, meaning one a node is passed it can never be reached again. In this section we focus on games in which agents always know what history they are at. In subsequent sections this strong informational assumption will be relaxed.

Given an extensive form game, a strategy profile is a schedule specifying what action agent i will select at every history h in the set $H_i = \{h \in H : p(h) = i\}$.

Definition 52 *Given an extensive form game* Γ^E*, a* ***strategy profile*** *for player* $i \in N$ *is a mapping* $s_i(h) : H_i \to A(h)$*. A strategy profile is a mapping* $s(h) : H \backslash H^T \to A(h)$*.*

Some contemplation will demonstrate that every extensive form game induces a normal from game.

Exercise 31 *Prove that every finite extensive form game has a Nash equilibrium in mixed strategies.*

9.3 Subgame perfect Nash equilibria

As suggested in the introduction to this chapter, the requirements of Nash equilibria are very weak in extensive form games. We desire an equilibrium concept that requires that at history h agent $p(h)$ anticipates that at subsequent histories agents will select actions which are reasonable (or optimal). One solution is to observe that at every history $h \in H \backslash H^T$ we can define a new extensive form game which has h as the null history. This new smaller extensive form game is termed a **subgame**. If we require that strategies constitute Nash equilibria to every subgame, then the strategies will involve optimal play both at histories that are reached (**on the equilibrium path**) and at histories which are not reached in equilibrium (**off the equilibrium path**).

Definition 53 *Given an extensive form game Γ^E, the set of* ***subgames*** *are all of the extensive form games constructed by selecting $h' \in H \backslash H^T$ and indentifying $\emptyset$ with h' and considering the restrictions of H, $p(\cdot)$ and $u_i(\cdot)$ to histories that include h'.*

Definition 54 *Given an extensive form game Γ^E, a strategy profile $s(\cdot)$ is a subgame perfect Nash equilibrium (SGPNE) if in every subgame to Γ^E the restriction of the strategy profile $s(\cdot)$ to the subgame is a NE of the subgame.*

The term restriction appearing in both of the previous definitions, is convenient and simple, but probably unfamiliar. An example will demonstrate the meaning. The restriction of H to histories that include h' is the subset $H' \subset H$ of histories that are of the form $(h', a_t, ..., a_k)$. When we identify h' with $\emptyset$ then this history looks like $(\emptyset, a_t, ..., a_k)$. A slightly less formal version of this definition is: *A SGPNE is a strategy profile that is a NE in every subgame.* In the finite extensive form games considered here, the characterization of SGPNE is straightforward, as a simple algorithm termed **backwards induction** may be used. The procedure is to start with the choice nodes that immediately precede terminal histories and determine which choices are optimal for the agents that choose at each of these histories. Once these optimal choices have been noted, we consider the choice nodes that immediately precede the nodes that were just analyzed. The optimal choice for agents at these histories is found by attributing to any choice, the terminal history that will be reached when choice in the final stage follows the characterization just given. Moving up the tree in this manner, we solve the game.

Definition 55 *To solve a finite extensive form game Γ^E by* ***backwards induction:***

-Step 1: Start with the set of histories that immediately proceed terminal histories and lead only to terminal histories H^{T-1} (we denote this set by H^{T-1}).

For each $h \in H^{T-1}$select the action $a \in A(h)$ which is optimal for $p(h)$. This action is termed $a^*(h) = \arg\max_{a \in A(h)} u_{p(h)}((h,a))$.

-Step 2: Consider the set of histories that immediately proceed the histories $H^{T-1}$and only lead to histories in H^{T-1}(we denote this set by H^{T-2}). For each $h \in H^{T-2}$ select the action $a \in A(h)$ which is optimal for $p(h)$ given that subsequent play is as in step 1. This action is termed $a^*(h) = \arg\max_{a \in A(h)} u_{p(h)}((h,a,a^*(h,a)))$.

........Step k: For each $h \in H^{T-k}$ solve for $a^*(h) = \arg\max_{a \in A(h)} u_{p(h)}((h,a,a^*(h,a)))$.

Continue until $H^{T-k} = \{\emptyset\}$ (the initial node).

Note that we add the clause "and lead only to ..." because some branches may be longer than others. Accordingly, at each step only nodes for which all subsequent play has already been resolved are analyzed. To see that any strategy profile $a^*(h)$ that is attained by this procedure is a SGPNE note that by starting at the bottom of the tree and working upwards, the procedure insures that in every subgame $a^*(h)$ gives NE play. We now turn to the question of existence of SGPNE.

Theorem 23 *Every finite extensive form game has a SGPNE. Moreover, if no player is indifferent between any two terminal histories then the SGPNE is unique.*

Exercise 32 *Prove the proposition (Hint use induction).*

Example 1: As an example we will consider a problem of sequential voting by 3 players $N = \{1,2,3\}$. Suppose that the choices x, y, z are to be voted on with the agenda: chose between x and y first, and then compare the winner with z, enacting either the winner from the first vote or z depending on which proposal gets the most votes. We assume that at each stage of voting ballots are cast simultaneously. Figure x depicts the game tree

[figure x about here]

Assume that the player have the following preferences over the enacted policy xP_1yP_1z; yP_2zP_2x; zP_3xP_3y. Applying subgame perfection and requiring that strategies are not weakly dominated (so voting is sincere) we see that if the final vote is between x and z then players 2 and 3 will vote for z. In contrast if the final vote is between y and z then players 1 and 2 will vote for y Accordingly, in voting over x and y in the first period, strategic agents will anticipate that the real choice is between the **sophisticated equivalents**, z and y. Accordingly players 1 and 2 will vote for y over x. Note that player 1 prefers x to y, but in a SGPNE she casts a strategic vote for y over x because she realizes that a vote for x is really a vote for z which she finds very unappealing.

9.4 Imperfectly observed actions

The basic extensive form games that satisfy the above definition are ones in which each agent knows exactly what actions have already occured. Accordinly, games of this form cannot capture situations inwhich players move simultaneously. The problem is that if only 1 player can move at a time, and evey previous action is publicly available, then one agent must know the other agents action already. This problem is easily resolved by introducting the notion of information sets. The idea is that players can be forced to choose without knowing what actions have proceeded this point in the game. An information set is a collection of histories which all give a specific player the same feasible choices, and which the player cannot discearn between the histories at the time the decision must be made.

Definition 56 *An* ***extensive form game with imperfectly observed actions,*** Γ^{EI}*, is a tuple* $\langle N, H, I, p(\cdot), U\rangle$ *satisfying (1)-(5) of definition xx where* $I = \{I_1, ..., I_z\}$ *is a partition of* H *into information sets (ie* I *is a list of disjoint subsets of* H *for which the union of all subsets in* I *is equivalent to* H*), satsifying the condition:*

(6) If $h, h' \in I_j$ *then* $p(h) = p(h')$ *and for ever action* a *the number of histories in* H *of the form* (h, a) *and* (h', a) *coincide.*

In drawing game trees representing games of this type it is customary to connect nodes in the same information set by a dashed line. Note that any extensive form game with perfectly observed actions is also an extensive form game with imperfectly observed actions involving only singleton information sets, (ie $I = H$). Given a game Γ^{EI} we need to ensure that strategy profiles preserve the information sets.

Definition 57 *Given an extensive form game* Γ^{EI}*, a* ***strategy profile*** *for player* $i \in N$ *is a mapping* $s_i(h) : H_i \rightarrow A(h)$*. s.t.* $s_i(h) = s_i(h')$ *if* h, h' *are in the same informaion set* I_j*.*

No modifications to the equilibrium concepts discussed are needed. Note, however, that in utilizing backwards induction the information sets need to be preserved.

Example 2: We can now represent the situation in wich three voters vote overy polices a, b, c in two periods, as a an extensive form game with imperfectly observed actions. Assume that preferences are aP_1bP_1c; bP_2cP_2a; cP_3aP_3b. First the voters simultaneously cast ballots between a and b. They then simultaneously cast ballots between the winner (by majority rule) of this comparison and c. The winner of this comparison (by majority rule) is enacted.

Exercise 33 *Represent example 2 as an extensive form game with imperfectly observed actions, and characterize the SGPNE.*

Exercise 34 *Characterize the SGPNE in weakly undominated strategies.*

9.5 Non finite extensive form games

Thus far, we have considered only extensive form games with a finite number of histories, (H is a finite set). The convenience of this assumption is that (1) we can actually represent the game by a tree, and more importantly (2) backwards induction is possible. In considering extensive form games in which there are an infinite number of non-terminal histories, backwards induction is not always possible. Nevertheless, the concept of subgame perfection is still well defined in these games. Moreover, in many applications it is possible to use backwards induction, even though the extensive form game is non finite.

We begin by extending the definition of an extensive form game, and then consider a few examples. There are no difficulties in considering extensive forms with H countably or even non countably infinite, as long as there is some finite k s.t. every terminal history consists of atmost k nodes. Such games are non-finite because agents have a large number of choices, but the game can be thought of as occuring over a finite number of time periods. In contrast games in which the length of histories is unbounded require somewhat different techniques of analysis. We will consider these games in a subsequent section. The extension we consider here is.

Definition 58 *An* ***extensive form game with with imperfectly observed actions and finite number of periods****,* Γ^{EIF} *is tuple* $\langle N, H, I, p(\cdot), U\rangle$ *with* H *a possibly non finite set, (1)-(6) of definition xx satisfied and satisfying the condition:*

(7) There exists some $k < \infty$ *s.t. for every* $h \in H^T$ *it is the case that* h *is a list of at most* k *actions.*

In games of this form the procedure of backwards induction may be used to characterize SGPNE if they exist, but no general existence results can be attained. In some games the preferences are sufficiently continous and convex and the choice sets at every information set are strictly convex and compact, so that iterative use of theorems xx and xx can be applied to establish existence. A a caution, note that, very simple games without SGPNE can be constructed.

Example 3: As an example we consider first a simple bargaining problem, in which player 1 must choose an offer $b_1 \in [0, 1]$ and then player 2 chooses to accept or reject the offer. If player 2 accepts the offer the payoffs are $b_1, 1 - b_1$ respectively. If player 2 rejects the offer she makes an alternative offer $b_2 \in [0, \delta]$ with $\delta \in (0, 1)$. In this case the payoffs are $\delta - b_2, b_2$ respectively. It is not difficult to see that if player 2 is ever given the opportunity to propose she will select $b_2 = \delta$. Accordingly, player 2 will only accept an offer tha gives her at least this much utility ($b_1 \leq 1 - \delta$). Given this the unique optimal proposal for player 1 is $b_1 = 1 - \delta$. The unique SGPNE is thus: $b_1 = 1 - \delta$, 2 accepts any offer in which $b_1 \leq 1 - \delta$ and rejects any other offer, If given the chance to propose $b_2 = \delta$.

Example 4: We now consider a simple example in the spirit of the model in Baron and Ferejohn (19xx). There are 3 legislators, $N = \{1, 2, 3\}$ selecting the allocation of a dollar. Thus the feasible allocations are $\Delta = \{(x_1, x_2, x_3) \in [0,1]^3 : \sum_{i=1}^{3} x_i = 1\}$. The sequence of moves is as follows: Player 1 is chosen to announce an allocation $x^1 \in \Delta$. The three players then vote for or agains the proposal. If a majority support the proposal the game is over and the payoffs are given by $u_i = x_i^1$. If the propsal does not pass, then player 2 is chosen to make an offer $x^2 \in \Delta$. Again there is a vote, and if the proposal passes the payoffs are $u_i = \delta x_i^2$. If the proposal fails then player 3 makes an offer $x^3 \in \Delta$. There is then a vote on this proposal. If the proposal passes, the payoffs are $u_i = \delta^2 x_i^3$. If the proposal fails the payoffs are $u_i = 0$. We will let $v_i^t = 1(0)$ denote a vote by player i in favor (against) a proposal made by player t. Clearly an optimal voting strategy for each player in the last round is:

$$v_i^3 = \begin{cases} 1 \text{ if } x_i^3 \geq 0 \\ 0 \text{ otherwise.} \end{cases}$$

When $x_i^3 = 0$ it is also optimal to vote 0 but it is not possible to construct an equilibrium when more then 1 player uses this strategy. Given this the optimal proposal for player 2 is $x^2 = (0, 0, 1)$. This proposal would pass and it gives 3 the greatest feasible utility if passed ($u_3 = \delta^2$). Given this when deciding whether to vote for or against a proposal by player 2 the choice for player 3 is between $u_i = \delta^2$ if the proposal fails and $u_3 = \delta x_3^2$ if it passes. Simmilarly player 1 chooses between $u_1 = 0$ if the proposal fails, and $u_1 = x_1^2$ if it passes. Accordingly optimal strategies are:

$$v_1^2 = \begin{cases} 1 \text{ if } x_i^2 \geq 0 \\ 0 \text{ otherwise.} \end{cases}$$

$$v_3^2 = \begin{cases} 1 \text{ if } x_i^2 \geq \delta \\ 0 \text{ otherwise.} \end{cases}$$

Given this the optimal proposal for 2 is $x^2 = (0, 1, 0)$. Simmilar logic indicates that player 3 will support anything that player 1 proposes, and thus her optimal proposal is $x^1 = (1, 0, 0)$.

Exercise 35 *Write out the complete SGPBE strategy profile to example 4 and verify that no other allocation is supportable as a SGPNE in weakly undominated strategies.*

In many cases it is very convenient to introduce moves by nature throughout the extensive form game. Baron and Ferejohn present a model very simmiliar to example 5, but involving uncertainty about which player will propose in periods 2 and 3 (if these periods are reached).

Example 5: Consider example 4 with the modification that the second and third proposers are randomly selected (by nature). Assume that if x^1 is rejected then each player has a $\frac{1}{3}$ chance of being recognized to propose in period 2. Simmiliarly, if x^2 is rejected then each player has a $\frac{1}{3}$ chance of being recognized to propose in period 3. In this game the voting strategy

$$v_i^3 = \begin{cases} 1 \text{ if } x_i^3 \geq 0 \\ 0 \text{ otherwise.} \end{cases}$$

is still a best response. Given this if i is recognized to propose in period 3 she will give 0 to the other players. Accordingly if x^2 is rejected, then players can expect the expected utility of $\frac{1}{3}$. This is the expected utility of bing the proposer with proability $\frac{1}{3}$ which results in the payoff 1. However because of discounting, in comparing the payoff to voting for x^2 or agaist it, the choice is between the share from x^2 and the reservation value $\frac{\delta}{3}$. Given this the optimal offer for the proposer of the second proposal is to give $\frac{\delta}{3}$ to one player and keep $1 - \frac{\delta}{3}$ for herself. We can assume that players that are recognized to make the second period proposal, that have not made the first period proposal, choose to give $\frac{\delta}{3}$ to the player that has not yet proposed and that if the same player is recognized in both of the first periods she randomizes (with equal probability) over which agent to give $\frac{\delta}{3}$ to. This proposal will pass with a bare majority. Given that this is the propsal to be made in period 2, the expected utility to not approving x^1 is $\frac{1}{3}\left(1 - \frac{\delta}{3}\right) + \frac{1}{3}\left(\frac{\delta}{3}\right) + \frac{1}{3}\left(\frac{1}{2}\right)\left(\frac{\delta}{3}\right)$. This value simplifies to $\frac{1}{3} + \frac{1}{18}\delta$. Accordingly, in the first period a voter will vote infavor of the offer x^1 if it gives her more than her reservation value $\delta\left(\frac{1}{3} + \frac{1}{18}\delta\right)$. Accordigly, the optimal first period offer is to keep $1 - \delta\left(\frac{1}{3} + \frac{1}{18}\delta\right)$ and give $\delta\left(\frac{1}{3} + \frac{1}{18}\delta\right)$ to one of the other players. This offer will be supported by a bare majority.

Exercise 36 *Characterize all SGPNE in weakly undominated strategies to the game described in example 5.*

Example 6: The setter model:

Example 7: The 2 period spatial bargaining model:

10 Game Theory III:Extensive form games of incomplete information

10.1 Introduction

Thus far, we have considered games with varying degree of uncertainty. In the Bayesian games of section 7.6, players faced uncertainty about the preferences (or types) of the other players. In the extensive form games with imperfectly observed actions, players were not certain what histories they had reached. While subgame perfection can rule out some unreasonable Nash equilibrium, in many extensive form games with imperfectly observed actions a stronger equilibrium concept is needed. Consider the extensive form game depicted in Figure 2. Player 1 chooses whether to secretly accumulate a military capability to attack an island. She can either not accumulate and thus not send any ships, NA, or she can accumulate by sending a small line of ships (b) or a big line of ships (B). Player 2 can only observe whether there was an accumulation, as she sees ships coming, but can't determine how many ships are coming. If no accumulation occurs then the payoffs are (0,5) as player 2 keeps the island. If there is an accumulation, then player 2 must decide whether to respond to the attack (R). If there is no response (NR) then player 1 wins the island. If there is a response, then player 2 wins the island but the casualties for player 2 are much higher under b then under B. The casualties for player 1 are higher under B then under b.

[Figure 2 about here]

There are three Nash Equilibria to this game. The first is (NA,R). This means that player 1 does not accumulate, but if she did player 2 would respond. The second Nash equilibrium is (B,NR). Player 1 accumulates a big line of ships, and player 2 does not respond. Inspection of the first NE indicates that it is not very reasonable. Regardless of whether B or b is played, player 2 is better off playing NR. But since this game has no proper subgames, the profile (NA,R) is a SGPNE. Simmiliary (b,NR) is a SGPNE.

The argument that this profile is not reasonable, hinges on the claim that player 1 should anticipate a rational response from player 2 at player 2's information set. We incorporate this type of **sequential rationality** into an equilibrium concept, by requiring that agents form **beliefs** about which history they are at and that they select best responses given these beliefs. Equilibria of this form are termed Perfect Bayesian Equilibria (PBE). In the example, no belief about which history player 2 is at will justify the selection of R as a best response. We now consider a slight modification of the game in figure 2. In this game player 1 can only win the island if she selects B. Moreover, player 2 would rather defend the island if player 1 has selected b. Figure 3 depicts the relevant payoffs.

[Figure 3 about here]

In this game whether playing R or NR is sequentially rational, hinges on what beliefs player 2 assigns to the 2 possible histories in the information set. If she believes that b was played then R is sequentially rational. Conversely if she believes that B was played then NR is sequentially rational. What strategy profile is reasonable in this game? In this section we will present the techniques needed to analyze games of this form.

10.2 Perfect Bayesian Equilibria

We now define the concepts needed to characterize PBE. We start with beliefs over histories.

Definition 59 *Given an extensive form game with imperfectly observed actions, Γ^{EI} a* ***belief on information set*** *$I_j \in I$ is a probability distribution on I_j. A* ***belief*** *is a mapping $b : H \to [0,1]$ s.t. for every $I_j \in I$ $b(\cdot)$ is a belief on I_j (that is for every $I_j \in I$ $\sum_{h \in I_j} b(h) = 1$ if I_j is finite and $\int_{h \in I_j} db(h) = 1$ if I_j is not finite).*

So in the examples above, a belief on player 2's information set, is a probability measure on the space $\{b, B\}$. Given a belief for each information set in I we can define sequential rationality, a condition on strategies. Throughout we will use the notation $p(I_j)$ and $s(I_j)$ to denote the player that moves at information set I_j and the action called for (by a strategy profile) at information set I_j. These terms are equivalent to $p(h)$ and $s(h)$ when $h \in I_j$. For a fixed strategy profile $s(\cdot)$ we denote the expected utility to player $p(I_j)$ associated with the choice a at history h by $Eu_{p(h)}(a, h, s(\cdot))$. This is an expected utility (as opposed to a utility) because players other than $p(h)$ may mix. When player $p(I_j)$ assigns probability $b(h)$ to being at history $h \in I_j$ conditional upon being at the information set I_j, the expected utility to taking action a at information set I_j is

$$Eu_{p(I_j)}(a, I_j, s(\cdot), b(\cdot)) = \sum_{h \in I_j} b(h) Eu_{p(h)}(a, h, s(\cdot)).$$

A strategy profile is sequentially rational relative to a belief if it involves optimal actions at each information set, when players evaluate the desirability of action a using $Eu_{p(I_j)}(a, I_j, s(\cdot), b(\cdot))$.

Definition 60 *Given an extensive form game with imperfectly observed actions, Γ^{EI} and a belief $b(h)$ on each information set, the strategy profile $s(\cdot)$* ***is sequentially rational (relative to the beliefs) at information set*** *I_j if given available action b we have*

$$Eu_{p(I_j)}(s(I_j), I_j, s(\cdot), b(\cdot)) \geq Eu_{p(I_j)}(b, I_j, s(\cdot), b(\cdot)).$$

If the strategy profile is sequentially rational (relative to the beliefs) at every information set, then it is ***sequentially rational (relative to the beliefs)****.*

Returning to the example in figure 3 above, if the beliefs assign probability close to 1 on b (B) then R (NR) is sequentially rational at the information set. We now consider a condition on beliefs. Recall that Bayes' rule is given by

$$prob(A \mid B) = \frac{prob(A\&B)}{prob(B)}.$$

Applying this formula to the probability that a particular history $h \in I_j$ is reached conditional on reaching the information set I_j and players using the strategy profile $s(\cdot)$ we find that

$$prob(h \mid I_j, s(\cdot)) = \frac{prob(h \text{ is played under } s(\cdot))}{prob(I_j \text{ is played under } s(\cdot))}$$

The term in the denominator is the probability that information set I_j is reached under strategy profile $s(\cdot)$. The term in the numerator is the probability that history h is reached under strategy profile $s(\cdot)$. Not that since $h \in I_j$ $prob(h$ is played under $s(\cdot)) = prob(h$ and I_j are played under $s(\cdot))$.

Definition 61 *Given an extensive form game with imperfectly observed actions, Γ^{EI} and a strategy profile $s(\cdot)$ we say that the beliefs $b(\cdot)$ are consistent relative to strategy $s(\cdot)$ if $b(h) = prob(h \mid I_j, s(\cdot))$ whenever*
$prob(I_j$ is played under $s(\cdot)) > 0$.

Consistency does not impose any restrictions on beliefs on information sets that do not occur under a particular strategy profile. This weakness is sometimes problematic, as we will see in a subsequent section. Combining consistency of beliefs and sequential rationality of strategies yields the equilibrium concept PBE.

Definition 62 *Given an extensive form game with imperfectly observed actions, Γ^{EI} a **perfect Bayesian equilibrium** (PBE) is a pair $(s(\cdot), b(\cdot))$ s.t.: (1) the strategy profile $s(\cdot)$ is sequentially rational relative to the belief $b(\cdot)$, and (2) the belief $b(\cdot)$ is consistent relative to the strategy profile $s(\cdot)$.*

Thus a PBE requires the construction of beliefs. The existence of beliefs allows us to define a notion of sequential rationality (optimality of choices at histories). Moreover, the beliefs that players entertain are related to the equilibrium strategies, in that histories which are relatively more likely to be reached under a strategy profile, are believed to occur with a higher probability by a player at a non-trivial information set.

Returning to the game in figure 2, we can now consider what strategy profiles occur in a PBE. Clearly the NE of (NA,R) is not supportable in a PBE, because for any beliefs about which history b or B player 2 is at when her information set is reached, NR is the unique response that is sequentially rational for 2 at the information set. Now given that player 2 is choosing NR, player 1's optimal choice is to play either b or B. Now if player 1 chooses B then consistent beliefs must assign probability 1 to player 2 being at history B. Thus one PBE is

(B,NR), $prob(B) = 1$, where $prob(B)$ is the posterior probability of B given that player 2's information set is reached under player 2's beliefs. Similarly there is a PBE of the form (b,NR), $prob(B) = 0$.

Now consider the game in figure 3. If player 2 believes that $prob(B) = 1$ then NR is the best response. On the other hand if player 2 believes that $prob(B) = 0$ then R is the best response. One candidate for a PBE is (NA,R), $prob(B) = 0$. Note that since no constraint is imposed on beliefs over the histories B and b when player 1 plays NA, the belief $prob(B) = 0$ is consistent relative to the strategy NA. But, the strategy profile (NA,R) is not sequentially rational as player 1 would prefer to play B than NA when she conjectures that player 2 is playing R. It is also clear that NA cannot be a best response to NR. Alternatively we can try to characterize a pure strategy PBE in which NA is not played. If B is played and beliefs are consistent than the only sequentially rational strategy by 2 will involve NR. But if player 1 conjectures that player 2 is playing NR she will want to play b. So we cannot have B played in a pure strategy PBE. On the other hand if b is played then consistent beliefs must assign probability 1 to player 2 being at this history. Thus, the only sequentially rational action will involve playing R. But if player 1 conjectures that player 2 is playing R then she will want to play B. Thus we cannot have a pure strategy PBE in which b is played. We have thus shown that there is no pure strategy PBE to the game.

It is not difficult to characterize the mixed strategy PBE to the game. Suppose that player 1 plays B with probability q and b with probability $(1 - q)$.Further suppose that player 2 plays R with probability z and NR with probability $(1-z)$. Consistency of beliefs requires that $prob(B) = q$. Now for player 2 to be indifferent between R and NR it must be the case that

$$q(-5) + (1-q)2 = q(0) + (1-q)0$$

This requires that $q = \frac{2}{7}$. Now in order for player 1 to be indifferent between playing B and b it must be the case that

$$(1-z)5 + z(-2) = (1-z)4 + z3$$

This requires that $z = \frac{1}{6}$. Accordingly the strategy profile, b with probability $\frac{5}{7}$, B with probability $\frac{2}{7}$, R with probability $\frac{1}{6}$ and NR with probability $\frac{5}{6}$ is supportable as a PBE, with the beliefs $prob(B) = \frac{2}{7}$.

Exercise 37 *Consider the game of figure 3, with the payoff to the path B,NR being (5,0) instead of (4,0). Characterize all of the PBE (mixed and pure strategy) to the game.*

Exercise 38 *Consider the game of figure 3, with the payoff to the path NA being (w,5) instead of (0,5). Here w is an exogenous parameter known to the agents that is ranging from* $[-2, 5]$. *For what regions of this range are there PBE in which NA occurs with positive probability. In other words for what subset of* $[-2, 5]$ *are there PBE in which NA is played.*

We now consider an extensive from game, in which nature draws a type $\theta \in \{A, B\}$ for player 1. Player 1 observes her type and chooses $s_1 \in \{a, b\}$. Player two observes the action chosen by player 1 but does not observe her type, and chooses $s_2 \in \{l, h\}$.Figure 4 depicts the game form. Note that under thes payoffs regardless of the strategy that player 2 is using, a player 1 of type A prefers action a and a player 1 of type B prefers action b. Upon observing action a player 2 prefers to play l (h) if she thinks that type A (B) has attained. Conversely upon observing action b player 2 prefers to play h (l) if she thinks that type B (A) has attained.

Exercise 39 *Find all of the PBE of the game depicted in figure 4.*

[Figure 4 about here]

11 Game Theory IV: Repeated Games

11.1 Introduction

Consider the normal form game

$$\begin{bmatrix} \text{player 1}\backslash\ \text{player 2} & \tilde{}\ \text{confess} & \text{confess} \\ \tilde{}\ \text{confess} & r, r & s, v \\ \text{confess} & v, s & p, p \end{bmatrix}$$

With $v > r > p > s$ the game is our old friend the PD. Suppose two enterprising but non-stealthy prisoners play the same game every day. While we know the unique NE of the one shot PD is (confess,confess) we may speculate that the repetition allows for the creation of a reputation, whereby player 1 plays ˜ confess because she anticipates player 2 will base future actions on player 1's behavior in the current period. As a stark demonstration of this logic suppose that each player adopts the following simple strategy: as long as the other player has always played ˜confess, I will play ˜confess. If the other player has ever played confess, I will always play confess. Given this strategy (by say player 2) we can investigate the calculus by player one at any period in which no one has previously played confess. Relative to playing ˜confess, the action confess yields the utility gain of $v - r$ in the current period. In all subsequent periods player 2 will play confess, and thus the best that player 1 can hope to get is p in all subsequent periods. Accordingly the desirability of the short term gain hinges on whether a one time gain of $v - r$ offsets the the long term cost of $p - r$ in all subsequent periods. In this chapter we consider a common models of intertemporal preferences and analyze the set of NE and SGPNE to repeated games.

11.2 Preferences over an infinite horizon

To motivate this section consider an agent choosing between two infinite sequences of per period monetary payoffs. The first sequence is $\{\$1, \$1, \$1, ...\}$, and the second sequence is $\{\$2, \$2, \$2, ...\}$. Which sequence is more desirable? The latter sequence giving \$2 in every period seems more attractive than the sequence giving \$1 in every period. But suppose that the agents utility function over such sequences is given by $\$1 + \$1 + \$1 + ...$, and $\$2 + \$2 + \$2 + ...$ Both of these sums are equal to ∞ (more precisely neither sum is convergent). Accordingly, if one just sums over per period payoffs then many sequences must be in the same indifference set. In many applications a reasonable solution is to introduce a notion of discounting. Letting $\{x_t\}_{t=1}^{\infty}$ be a sequence of per period outcomes with each $x_t \in X$, and $u_i : X \to R^1$ be a per period utility function, we will represent preferences over sequences by the intertemporal utility function $U_i(\{x_t\}_{t=1}^{\infty}) = \sum_{t=1}^{\infty} \delta_i^{t-1} u_i(x_t)$, where $\delta_i \in [0, 1]$ is interpreted as agent i's discount rate. When $\delta_i = 1$ the agent values each period the same. When $\delta < 1$

the agent values earlier periods more than later periods. It can be shown that $\sum_{t=1}^{\infty} \delta_i^{t-1} = \frac{1}{1-\delta}$ meaning that as long as $\delta < 1$ and there is some $k < \infty$ s.t. $\max_x u_i(x) < k$, the sum $\sum_{t=1}^{\infty} \delta_i^{t-1} u_i(x_t)$ is finite. In fact this sum is bounded by $\frac{k}{1-\delta}$. The discount rate thus serves two roles, it captures an often plausible notion of impatience, and it allows us to represent by a finite number the utility of some infinite streams of per period payoffs. Other methods of representing preferences over infinite streams have been explored, but throughout we focus on discounting.

11.3 Folk theorems

Our theory of repeated games will take as primitives a normal form stage game $\Gamma = \langle N, S, u \rangle$ and pair of agent discount rates ($\delta = (\delta_i, ..., \delta_n)$). The interpretation is that in each period $t \in \{1, 2, 3, ...\}$ the normal form game Γ is played. In this context the game Γ is often called the stage game to distinguish it from the repeated game. Before agent i selects $s_i^t \in S_i$, her strategy in period t she observes the strategy profile s^{t-1} played in period $t-1$. Moreover, we maintain the assumption of perfect recall, meaning that s_i^t can be conditioned on the history $h^{t-1} = (s^1, ..., s^{t-1}) \in S^{t-1} := \prod_{j=1}^{t-1} S$. The null history is $h^0 = \emptyset$. A pure strategy for player i is then a sequence of mappings $\left\{s_i^t(h^{t-1}) : S^{t-1} \to S_i\right\}_{t=1}^{\infty}$. A mixed strategy is a sequence of mappings $\left\{\sigma_i^t(h^{t-1}) : S^{t-1} \to \Delta(S_i)\right\}_{t=1}^{\infty}$. Given a sequence of lotteries over stage game profiles $\{\sigma^t\}_{t=1}^{\infty}$ agent i's expected utility is given by $\mathbb{E}U_i(\{\sigma^t\}_{t=1}^{\infty}) = (1-\delta_i) \sum_{t=1}^{\infty} \delta_i^{t-1} \mathbb{E}_{\sigma^t} u_i(s^t)$ where $\mathbb{E}_{\sigma^t} u_i(s^t)$ takes the expectation of $u_i(s^t)$ over the mixture σ^t. The multiplier $(1-\delta_i)$ is included so that for a constant sequence σ^t, $\mathbb{E}U_i(\{\sigma^t\}_{t=1}^{\infty}) = u_i(\sigma^t)$. We denote the repeated game induced by a stage game, by $\Gamma^{\infty} = \langle N, S, u, \delta \rangle$. Of course a repeated game is also an extensive form game and our notions of NE and SGPNE are well defined in the repeated game.

We now focus on repeated games generated by finite normal form stage games. Given proposition (xx) we know that every such stage game has at least one mixed strategy NE. It is not surprising then that every such repeated game has as a mixed strategy SGPNE the infinite repetition of the stage game mixed strategy NE.

Proposition 1 *If σ^* is a SGPNE of the stage game then the repeated game profile $\sigma_i^t(h^{t-1}) = \sigma_i^*$ for every (h^{t-1}) for every t for every i is a SGPNE of the repeated game.*

Exercise 40 *Prove the proposition.*

An interesting feature of repeated games is that the set of SGPNE is usually very large. The class of results termed "Folk theorems" serve to quantify the set of equilibrium payoffs that are supportable in an equilibrium. We prove a particularly useful and simple Folk theorem. We first need several definitions.

Definition 63 *The payoff vector* $v \in \mathbb{R}^n$ *is termed individually rational if* $v_i \geq \min_{s_{-i} \in S_{-i}} \{\max_{s_i \in S_i} u_i(s_i, s_{-i})\}$.

The value $\min_{s_{-i} \in S_{-i}} \{\max_{s_i \in S_i} u_i(s_i, s_{-i})\}$ is the minimum stage game utility that player i can attain when she plays a best response. This value is identified by letting the players $-i$ select s_{-i} so as to minimize the utility to i of playing a best response to s_{-i}.

Definition 64 *The payoff vector* $v \in \mathbb{R}^n$ *is termed feasible if there is some sequence of pure strategy stage game profiles* $\{s^t\}_{t=1}^{\infty}$ *s.t. for each* $i \in N$, $\mathbb{E}U_i(\{s^t\}_{t=1}^{\infty}) = v_i$.

Proposition 2 *For every feasible and individually rational payoff vector* $v \in \mathbb{R}^n$ *there is an* n*-tuple of discount rates* δ *s.t. the payoff vector* v *occurs in a NE of the repeated game with the discount rates* δ.

Proof: Assume that v is feasible and individually rational. Let $\{s^{vt}\}$ be a strategy profile that calls for playing the strategy that attains the payoff vector v as long as no one has previously deviated from this strategy or more than two players have deviated, and plays the strategy $\{s^{pt} = \arg\min_{s_{-i} \in S_{-i}} \{\max_{s_i \in S_i} u_i(s_i, s_{-i})\}\}$ which punishes the unique player that deviated in all subsequent periods. At any period t the payoff to i of playing $\{s^{vt}\}$ is v_i and the payoff to deviating is bounded by

$$(1 - \delta_i^t)v_i + \delta_i^t(1 - \delta_i)\max_{s \in S} u_i(s) + \delta_i^{t+1} \min_{s_{-i} \in S_{-i}} \left\{ \max_{s_i \in S_i} u_i(s_i, s_{-i}) \right\}$$

This value is less than v_i if

$$\delta_i \geq \frac{\max_{s \in S} u_i(s) - v_i}{\max_{s \in S} u_i(s) - \min_{s_{-i} \in S_{-i}} \{\max_{s_i \in S_i} u_i(s_i, s_{-i})\}}$$

Since $\max_{s \in S} u_i(s) \geq v_i \geq \min_{s_{-i} \in S_{-i}} \{\max_{s_i \in S_i} u_i(s_i, s_{-i})\}$ the right hand side is strictly less than 1. Thus as long as this condition is satisfied for each $i \in N$ the conjectured strategy profile is a NE to the repeated game.■

The equilibria used in the proof need not be SGPNE as the punishment might be very costly to impose. We can quantify a set of payoff vectors supportable in SGPNE to the repeated game using reversion to stage game NE strategies as the punishment.

Proposition 3 *If* $v \in \mathbb{R}^n$ *is a feasible payoff vector for which there is some mixed strategy stage game NE which yields the payoff vector* v' *s.t.* $v'_i < v_i$ *for every* $i \in N$ *then there is a SGPNE in the repeated game which yields the payoff vector* v*.*

Exercise 41 *Prove the proposition.*