POLS 601: Definitions and Theorems Regarding Game Theory

Ahmer Tarar
Department of Political Science
Texas A&M University
email: ahmertarar@tamu.edu

December 1, 2016

In these notes, I will mainly just be providing formal definitions and theorems, so that I don't have to write them in their entirety on the blackboard in class. Most of these are in the Tadelis textbook (some of which I have modified slightly), but some are not. I have indicated where in the text (or alternative texts, e.g. Osborne) these can be found. In some places I have provided some explanation and details, but this is much less than what I go over in class—so these should supplement your lecture notes, not replace them. But all formal definitions and theorems will be provided in these notes, so in class you don't need to worry about writing these down fully and accurately.

1 Rational Choice Theory

Rational Choice Theory can be thought of as the rigorous study of rational decision-making. It can be said to have 3 main branches: (a) decision theory (the study of single-player decision-making), (b) game theory (the study of interdependent or "strategic" decision-making, i.e., when at least 2 decision-makers are involved), and (3) social choice theory (how individual preferences are aggregated into group preferences or group choice, also called collective choice).

2 Decision Theory

2.1 Complete Information (No Uncertainty)

(Tadelis, p.4)

Definition 1 *A* ***decision theory model under complete information*** *consists of:*

- *A decision maker (also called a player or actor)*
- *A set A (finite or infinite) of actions available to the decision-maker*
- *A set X (finite or infinite) of outcomes that can result from those actions*
- *An outcome function $g : A \to X$ mapping actions to outcomes, where for any $a \in A$, $g(a) \in X$ is the outcome resulting from a*[1]
- *The decision maker's preference ordering of the elements of X*

Important: preferences are over outcomes, not actions. Actions are just means to outcomes, which are what the decision-maker ultimately cares about.

Definition 2 *The* ***weak preference relation*** *is denoted by $\succeq$. For any $x, y \in X$, $x \succeq y$ means "x is at least as good as y". (Sometimes this is instead denoted by xRy.)*

Definition 3 *The* ***strict preference relation*** *is denoted by $\succ$. For any $x, y \in X$, $x \succ y$ means "x is strictly better than y". (Sometimes this is instead denoted by xPy.)*

Definition 4 *The* ***indifference preference relation*** *is denoted by $\sim$. For any $x, y \in X$, $x \sim y$ means "x and y are equally good". (Sometimes this is instead denoted by xIy.)*

[1] Tadelis does not include this in his formal definition. But he refers to this function on p.10, using the notation $x(a)$. But because he also uses x to refer to a generic element of X, this is confusing. Hence we use $g(a)$ instead of $x(a)$.

Usually $\succeq$ is considered the fundamental or primitive of the decision-maker's preference ordering, because $\succ$ and $\sim$ can be defined in terms of it as follows.

Definition 5 *For any* $x, y \in X$, $x \succ y \iff x \succeq y$ *and* $\neg(y \succeq x)$.

Definition 6 *For any* $x, y \in X$, $x \sim y \iff x \succeq y$ *and* $y \succeq x$.

Thus, if we have the entire $\succeq$ relation, this allows us to deduce the entire $\succ$ and $\sim$ relations. (Preference relations are more rigorously defined in terms of binary relations on X, which are covered in the optional appendix.)

(Tadelis, p.5)

Definition 7 *The preference relation* $\succeq$ *is* ***complete*** *if for any* $x, y \in X$, *either* $x \succeq y$ *or* $y \succeq x$ *(or both).*

(Tadelis, p.6)

Definition 8 *The preference relation* $\succeq$ *is* ***transitive*** *if for any* $x, y, z \in X$, *if* $x \succeq y$ *and* $y \succeq z$, *then* $x \succeq z$.

(Tadelis, p.6)

Definition 9 *The preference relation* $\succeq$ *is* ***rational*** *if it is complete and transitive.*

Henceforth, we will assume that each decision-maker has a rational preference relation. For a decision-maker that does not, we cannot systematically study which action it will choose, which is our goal in rational choice theory.

The Condorcet Paradox: transitive individual preferences need not imply a transitive "group preference", where by group preference we mean the preference ordering given by majority rule voting over pairs of outcomes.

(Tadelis, p.7)

Definition 10 *A* ***payoff function*** *or* ***utility function*** *representing the preference relation* $\succeq$ *is a function* $u : X \to \mathbb{R}$ *with the property that for any* $x, y \in X$, $u(x) \geq u(y)$ *if and only if* $x \succeq y$.

Utilities and utility functions are convenient ways of representing preferences, rather than dealing with the cumbersome $\succeq$. But the preference ordering is what is fundamental, not the utility function used to conveniently represent it.

This is because with decision-making under certainty, preferences are just ordinal (i.e., just the order matters for what the decision-maker ends up choosing, not the intensity of preferences), and hence many different (in fact, an infinite number of) utility functions represent a given preference relation. No specific utility function (i.e., any specific assignment of utilities to the outcomes) is special.

(Osborne, p.6)

Proposition 1 *If* $u : X \to \mathbb{R}$ *is a utility function representing* $\succeq$, *then so is any function* $f : X \to \mathbb{R}$ *with the property that for any* $x, y \in X$, $f(x) > f(y) \iff u(x) > u(y)$.

[Sidenote to ambitious students: For any given proposition in these notes, think about how you would go about proving it. Thinking about this will guide you to seeing whether this is a "do-able" proof for students at this level, or if it is instead extremely challenging. If you think it is "do-able", try proving it. Also think about whether some statements in the proposition are crucial, or whether the proposition can be stated more generally. For example in the above proposition, is $>$ crucial, or is it also true for $\geq$? This is all optional, however, and the required homework will be enough for most students.]

For example, if $u(\cdot)$ reflects the decision-maker's preferences, then so does $f(\cdot) = u(\cdot) + 1$. So do many other functions.

Proposition 2 *If $u : X \to \mathbb{R}$ is a utility function representing $\succeq$, then so is the composite function $f \circ u : X \to \mathbb{R}$, for any $f : \mathbb{R} \to \mathbb{R}$ that is a strictly increasing function.*

(Tadelis, p.8)

Proposition 3 *If the set of outcomes X is finite, then any rational preference relation over X can be represented by a payoff function.*

A decision theory model is often conveniently represented by a **decision tree**, with a **decision node**, **terminal nodes** (outcomes), and **payoffs** or **utilities** assigned to those terminal nodes. This is much more convenient than specifying all of the components of the formal definition given above.

Although preferences and hence payoffs are defined over outcomes, it is sometimes more convenient to talk about the payoff for an action, which is of course the payoff for the outcome that that action leads to. Formally:

(Tadelis, p.10)

Definition 11 *Let $g : A \to X$ be the outcome function (see Definition 1), and let $u : X \to \mathbb{R}$ be a payoff function that represents $\succeq$. Then the corresponding **payoff function over actions** is the composite function $v = u \circ g : A \to \mathbb{R}$. That is, for any $a \in A$, $v(a) = u(g(a))$ is the payoff (utility) for action a.*

The following definition embodies the **principle of rational choice** for decision-making under certainty.

(Tadelis, p.10)

Definition 12 *A player facing a decision problem with a payoff function $v(\cdot)$ over actions is **rational** if he chooses an action $a \in A$ that maximizes his payoff. That is, he chooses an $a^* \in A$ with the property that $v(a^*) \geq v(a)$ for all $a \in A$.*

Equivalently:

Definition 13 *A player facing a decision problem with a payoff function $u(\cdot)$ is **rational** if he chooses an $a^* \in A$ with the property that $u(g(a^*)) \geq u(g(a))$ for all $a \in A$.*

Equivalently:

Definition 14 *A player facing a decision problem with a payoff function $u(\cdot)$ is **rational** if he chooses an $a \in A$ that solves the problem $\max_{a \in A} u(g(a))$.*

Equivalently:

Definition 15 *A player facing a decision problem with a payoff function $u(\cdot)$ is **rational** if he chooses an element of the set $\arg\max_{a \in A} u(g(a))$.*[2]

Equivalently (and getting rid of utility functions altogether):

Definition 16 *A player facing a decision problem is **rational** if he chooses an $a^* \in A$ with the property that $g(a^*) \succeq g(a)$ for all $a \in A$.*

[2]For any real-valued function $f : X \to \mathbb{R}$, $\arg\max_{x \in X} f(x)$ is the set of **maximizers** of f (on X). Formally, $\arg\max_{x \in X} f(x) = \{x^* \in X \mid f(x^*) \geq f(x) \forall x \in X\}$. The notation $\max_{x \in X} f(x)$ refers to the maximum value of f (on X), also called simply the **maximum** of f (on X). Formally: $\max_{x \in X} f(x) = f(x^*)$ such that there exists an $x^* \in X$ such that $f(x^*) \geq f(x)$ for all $x \in X$. Alternatively, $\max_{x \in X} f(x)$ is the unique element, if any, of the set $\{f(x^*) \mid \exists x^* \in X \text{ s.t. } f(x^*) \geq f(x) \forall x \in X\}$.

Going back to the original definition, how do we know that such an $a^* \in A$ exists? That is, how do we know that an optimal action exists? The following two results give sufficient conditions.

Proposition 4 *If A is finite, then $v : A \to \mathbb{R}$ has a maximum (and, incidentally, a minimum) on A. That is, there exists an $a^* \in A$ with the property that $v(a^*) \geq v(a)$ for all $a \in A$. (Such an a^* is called a* ***maximizer*** *of v.)*

Proposition 5 *(A is infinite) If A is a closed interval and $v : A \to \mathbb{R}$ is continuous on A, then $v(\cdot)$ has a maximum (and, incidentally, a minimum) on A.*

The latter is just the Extreme Value Theorem (Proposition 23 in the summer math camp notes that I provided you) applied to A and $v(\cdot)$.

In rational choice theory, we assume that each decision maker has a rational preference relation and is rational. Both of these notions have now been formally defined. If we do not assume both of these, we cannot systematically study the action that a decision-maker will choose, which is our goal.

2.1.1 Appendix (Optional): More Advanced Topics and Exercises

These are things that I would not typically include in a first graduate course on game theory and hence are optional, but that you might want to think about and solve if you are so inclined.

Preference relations are more rigorously defined using binary relations. To define a binary relation, first note that for any two sets A_1 and A_2, the **Cartesian product** of A_1 and A_2 is denoted by $A_1 \times A_2$ and is defined as the set of all ordered pairs (a_1, a_2) of elements of A_1 and A_2, where the first element in the pair is a member of A_1 and the second element is a member of A_2.[3] Formally, $A_1 \times A_2 = \{(a_1, a_2) \mid a_1 \in A_1 \text{ and } a_2 \in A_2\}$. More generally, for any finite collection of sets $A_1, A_2, \ldots, A_n$, we define $A_1 \times A_2 \times \ldots \times A_n$ to be the set of all ordered "n-tuples" of elements from these sets. Formally, $A_1 \times A_2 \times \ldots \times A_n = \{(a_1, a_2, \ldots, a_n) \mid a_1 \in A_1, a_2 \in A_2, \ldots, a_n \in A_n\}$.

Definition 17 *For any finite collection of sets $A_1, A_2, \ldots, A_n$, the **Cartesian product** $A_1 \times A_2 \times \ldots \times A_n$ is the set of all ordered "n-tuples" of elements from these sets. Formally, $A_1 \times A_2 \times \ldots \times A_n = \{(a_1, a_2, \ldots, a_n) \mid a_1 \in A_1, a_2 \in A_2, \ldots, a_n \in A_n\}$. (Sometimes we use the notation $\prod_{i=1}^{n} A_i$ for the Cartesian product, i.e., $\prod_{i=1}^{n} A_i = A_1 \times A_2 \times \ldots \times A_n$. The notation $\times_{i=1}^{n} A_i$ is also used.)*

Definition 18 *Let X be a set. A **binary relation** B on X is a subset of $X \times X$, i.e., $B \subseteq X \times X$.*

Let X be the set of outcomes. Then $X \times X = \{(x, y) \mid x \in X \text{ and } y \in X\}$. We say that the binary relation B_R on X represents the weak preference relation $\succeq$ if for any $x, y \in X$, $(x, y) \in B_R \iff x \succeq y$.

Notice then that if we have the entire set B_R for a decision-maker, this gives us the entire weak preference relation for that decision-maker. It also gives us the entire strong and indifference preference relations, as follows. Suppose $(x, y) \in B_R$, meaning that $x \succeq y$. If $(y, x) \in B_R$ as well, then obviously $x \sim y$ because $y \succeq x$ as well. On the other hand, if $(y, x) \notin B_R$, then $x \succ y$. So B_R gives us the decision-maker's entire preference ordering over the set of outcomes.

More formally, we have the following definitions.

[3] Unfortunately, we use the same notation for an ordered pair as we do for an open interval on the real line, even though they are completely different things. In any given context, you should be able to recognize which one we mean when we use that notation.

Definition 19 *A **weak preference relation** is a binary relation on X, denoted by B_R, defined as follows: $B_R = \{(x, y) \in X \times X \mid x \text{ is at least as good as } y\}$. (If $(x, y) \in B_R$, we interpret that to mean that "x is at least as good as y".)*

Definition 20 *A **strict preference relation** is a binary relation on X, denoted by B_P, defined as follows: $B_P = \{(x, y) \in X \times X \mid x \text{ is strictly better than } y\}$. (If $(x, y) \in B_P$, we interpret that to mean that "x is strictly better than y".)*

Definition 21 *An **indifference preference relation** is a binary relation on X, denoted by B_I, defined as follows: $B_I = \{(x, y) \in X \times X \mid x \text{ and } y \text{ are equally good}\}$. (If $(x, y) \in B_I$, we interpret that to mean that "x and y are equally good".)*

Optional Exercises:

(A1) Define what it means for $\succ$ to be complete, and to be transitive.

(A2) Define what it means for $\sim$ to be complete, and to be transitive.

(A3) Is it reasonable to impose the assumption that a rational preference ordering should require $\succ$ to be complete? Explain why or why not. What about transitive?

(A4) Answer the same questions for $\sim$.

(A5) Are any of the following statements true? If so, prove it. If not, but the statement is at least coherent, provide a counterexample (a specific preference ordering) to show that it is not true in general. If it is not even coherent, explain why.

(a) "If $\succeq$ is transitive, then $\succ$ is transitive."

(b) "If $\succeq$ is complete, then $\succ$ is complete."

(c) "If $\succeq$ is transitive, then $\sim$ is transitive."

(d) "If $\succeq$ is complete, then $\sim$ is complete."

(A6) Let $A_1 = \{a, b, c\}$ and $A_2 = \{3, 10\}$. List $A_1 \times A_2$. List $A_2 \times A_1$.

(A7) If $\succeq$ is complete, what does that imply for B_R?

(A8) If $\succeq$ is transitive, what does that imply for B_R?

(A9) If $B_R = X \times X$, what does that imply for the decision-maker's preference ordering?

(A10) Suppose that the set of actions is infinite, and is the union of two disjoint closed intervals A_1 and A_2. That is, $A = A_1 \cup A_2$ and $A_1 \cap A_2 = \emptyset$. Suppose that $v(\cdot)$ is continuous on A_1 and A_2. Does $v(\cdot)$ have a maximum on A? If so, prove it. If not, provide a counterexample to show that it is not true in general.

2.2 Incomplete Information (Uncertainty)

Note in Definition 1 that the outcome function $g : A \to X$ is a function from actions to outcomes, and hence assigns exactly one outcome to each action. In other words, there is a deterministic and known mapping from actions to outcomes, and hence a rational decision-maker simply chooses an action that leads to a most-preferred outcome. Defining rational choice with decision-making under certainty is easy.

Oftentimes, however, the decision-maker is not certain what outcome will result from one or more of its actions, but can at least assign probabilities to the different outcomes. To capture this, we now replace g with h, where $h(a)$ is a (possibility **degenerate**, i.e., assigning probability 1 to a certain outcome) probability distribution, also called a simple lottery, over X, where $h(a)$ is the probability distribution over outcomes induced by the action $a \in A$.

(Tadelis, p.16)

Definition 22 *A* ***simple lottery*** *over a finite set of outcomes* $X = \{x_1, x_2, \ldots, x_n\}$ *is defined as a* ***probability distribution*** $p = (p(x_1), p(x_2), \ldots, p(x_n))$*, where* $0 \leq p(x_k) \leq 1$ *is the probability that* x_k *occurs and* $\sum_{k=1}^{n} p(x_k) = 1$*.*

(Tadelis, p.18)

Definition 23 *A* ***simple lottery*** *over an interval set of outcomes* $X = [\underline{x}, \overline{x}]$ *is given by a* ***cumulative distribution function*** $F : X \to [0, 1]$*, where* $F(\hat{x}) = \Pr(x \leq \hat{x})$ *is the probability that the outcome is less than or equal to* $\hat{x}$*.*

If F is differentiable, then the function $f = F'$ is the **density function** for the simple lottery F. We usually visualize continuous probability distributions by their density functions rather than their CDF's. For example, we are familiar with the density function of the uniform distribution and the normal distribution.

(Tadelis, p.102)

Definition 24 *Let X be a set. We use the notation ΔX to denote the set of all simple lotteries (probability distributions) over X. (When X is finite, we sometimes call ΔX the* ***simplex*** *of X.) Formally, we have the following formidable definitions involving sets of functions. If X is finite, then $\Delta X = \{p : X \to [0,1] \mid \sum_{k=1}^{n} p(x_k) = 1\}$. If X is an interval, then $\Delta X = \{F : X \to [0,1] \mid F \text{ is a cumulative distribution function}\}$.*

Definition 25 *A* ***decision theory model under uncertainty*** *consists of:*

- *A decision maker (also called a player or actor)*
- *A set A (finite or infinite) of actions available to the decision-maker*
- *A set X (finite or infinite) of outcomes that can result from those actions*
- *A function $h : A \to \Delta X$ mapping actions to simple lotteries over X, where for any $a \in A$, $h(a) \in \Delta X$ is the simple lottery over outcomes induced by a*
- *The decision maker's preference ordering of the elements of X*

With decision-making under uncertainty, the problem is that the decision-maker's preferences are over outcomes, but actions are associated with lotteries over outcomes rather than with outcomes. The natural extension of the principle of rational choice to decision-making under uncertainty is that the decision-maker should choose an action that maximizes its payoff function over actions (i.e., choose a maximizer of this function), but now it is no longer clear what this payoff function should be, unlike in the previous section where it was simply $v(a) = u(g(a))$, i.e., simply determined by a payoff function over outcomes (and of course by the outcome function).

It is clear that the decision-maker's preference over actions should be determined by its **preference over lotteries** (over outcomes), because there is a one-to-one link between actions and lotteries. So to predict what a rational decision-maker will do, what we really need to know is its preferences over lotteries (over outcomes). But this is not in general indicated by its preferences over outcomes (which are the fundamentals or primitives of the decision theory model). von Neumann and Morgenstern (1944) showed that if the decision-maker's preferences over lotteries satisfy certain conditions (henceforth called the **vN-M conditions**), then the decision-maker's preferences over lotteries are given by the **expected value** of a (certain) payoff function $u(x)$ (over outcomes) under those lotteries, and hence a rational decision-maker should be defined as one who chooses an action associated with a highest expected value of its payoff function over outcomes.

(Tadelis, p.19)

Definition 26 *Let $u(x)$ be the player's payoff function over outcomes in the finite set $X = \{x_1, x_2, \ldots, x_n\}$, and let $p = (p_1, p_2, \ldots, p_n)$ be a lottery over X such that $p_k = \Pr(x = x_k)$. Then we define the player's **expected payoff from the lottery** p as*

$$E[u(x) \mid p] = \sum_{k=1}^{n} p_k u(x_k) = p_1 u(x_1) + p_2 u(x_2) + \ldots + p_n u(x_n).$$

The expected payoff is a weighted average of payoffs, where the weight given to each payoff is the probability that the player receives that payoff.

(Tadelis, p.20)

Definition 27 *Let $u(x)$ be the player's payoff function over outcomes in the interval $X = [\underline{x}, \overline{x}]$, and consider a lottery given by the cumulative distribution function $F(x)$, with density function $f(x) = F'(x)$. Then we define the player's* ***expected payoff from the lottery*** *F as*

$E[u(x) \mid F] = \int_{\underline{x}}^{\overline{x}} u(x)f(x)dx.$

(If F is not differentiable, then $E[u(x) \mid F] = \int_{x \in X} u(x)dF(x)$. The intuition behind this is that $f(x) = \frac{dF(x)}{dx}$, and hence $f(x)dx = dF(x)$.)

von Neumann and Morgenstern showed the following.

(McCarty and Meirowitz, p.36)

Proposition 6 *If the player's preferences over lotteries satisfy the vN-M conditions (we call such preferences* ***vN-M preferences****), then there exists a utility function $u : X \to \mathbb{R}$ representing $\succeq$ such that for any two lotteries $p, p' \in \Delta X$ (alternatively, $F, F' \in \Delta X$ if X is an interval), the player weakly prefers p to p' (alternatively, F to F') if and only if $E[u(x) \mid p] \geq E[u(x) \mid p']$ (alternatively, $E[u(x) \mid F] \geq E[u(x) \mid F']$). The utility function u is called a* ***Bernoulli utility function****.*

In other words, to compare two lotteries, we simply compare the expected value of $u(x)$ under those two lotteries.

We can now give the principle of rational choice for decision-making under uncertainty as follows.

Definition 28 *Suppose a player facing a decision problem under uncertainty has preferences over lotteries that satisfy the vN-M conditions, and that these preferences are therefore represented by the expected value of a Bernoulli payoff function $u : X \to \mathbb{R}$ that represents*

$\succeq$. Let $p_a \in \Delta X$ denote the simple lottery over X induced by action $a \in A$ if X is finite, and denote this by $F_a \in \Delta X$ if X is an interval. Let $v(a) = E[u(x) \mid p_a]$ if X is finite, and $v(a) = E[u(x) \mid F_a]$ if X is an interval, be the corresponding ***payoff function over actions*** *(clearly, $v : A \to \mathbb{R}$). Then the player is* ***rational*** *if he chooses an $a^* \in A$ such that $v(a^*) \geq v(a)$ for all $a \in A$, i.e., if he chooses an action that maximizes $v(a)$. Sometimes we call $v(a)$ the* ***vN-M expected utility function****.*

Henceforth, whenever we are dealing with decision-making (whether single-player or multi-player) where uncertainty (and hence probabilities and lotteries) is involved, we will assumed that the decision-maker's preferences over lotteries satisfy the vN-M conditions, and hence will simply use expected utilities to evaluate the player's options. Clearly, vN-M's **expected utility theory** makes decision-making under uncertainty much more tractable than it would otherwise be. We don't have to list the decision-maker's entire preference ordering over the (always infinite) set of lotteries ΔX, but can instead just work with the expected value of a utility function over outcomes.

With decision-making under certainty, there were an infinity of utility functions that represent any given preference ordering $\succeq$, and any one of them was acceptable to use, because the predicted action that the decision-maker will choose is the same regardless. Is the same true with decision-making under uncertainty? Note that Proposition 6 simply states that there is *a* utility function whose expected value represents the decision-maker's preferences over lotteries. Is this u unique? The answer is given in the following result.

(Osborne, p.148)

Proposition 7 *Suppose that there are at least three possible outcomes. The expected values of the Bernoulli payoff functions $u : X \to \mathbb{R}$ and $w : X \to \mathbb{R}$ represent the same preferences* *over lotteries if and only if there exist numbers a and b with $b > 0$ such that $w(x) = a + bu(x)$ for all $x \in X$.*

That is, the u specified in Proposition 6 is *not* unique; for any given preference ordering over lotteries, there are an infinite number of Bernoulli payoff functions whose expected value represents that preference ordering, but they are related to each other in an exact way. Namely, for any two, each can be written as a strictly increasing *linear* function of the other. The same is not true of the ordinal utility functions that are sufficient for decision-making under certainty (e.g., Proposition 2 above). Let $X = [0, 1]$. Then $u(x) = x$ and $w(x) = \sqrt{x}$ represent the same ordinal preferences over outcomes, but different preferences over lotteries, because one cannot be written as a strictly increasing linear function of the other. On the other hand, $u(x) = x$ and $z(x) = 5x - 2$ not only represent the same ordinal preferences over outcomes, but the same preferences over lotteries as well. (This does *not* mean that Bernoulli payoff functions have to be linear functions of x. $w(x) = \sqrt{x}$ and $k(x) = 10\sqrt{x} - 3$ represent the same preferences over lotteries.)

With decision-making under certainty, we said that we are assuming no more than **ordinal** preferences, because that is all we need to postulate what a rational decision-maker will choose. If I strictly prefer outcome x_1 over all others, I will choose an action that leads to x_1 no matter *how much* I prefer x_1 to other outcomes: it doesn't matter whether we specify the utility function $b(x_1) = 100$ or $c(x_1) = 1000$, as long as these are the highest numbers. Thus, we say that there is no meaningful difference between $b(x)$ and $c(x)$, as they represent the same ordinal preferences.

But with decision-making under uncertainty, intensity of preferences matter, because this influences the preference between lotteries: expected utilities, and hence preferences between lotteries, depend on the specific Bernoulli payoff function being used, and hence so too does the specific action chosen (up to the non-uniqueness allowed by Proposition 7). Thus, Bernoulli payoff functions also incorporate the **intensity** of preferences over outcomes, and these are referred to as **cardinal** utility functions. With decision-making under uncertainty, saying that the player's Bernoulli payoff function is $c(x_1) = 1000$ rather than $b(x_1) = 100$ is saying something meaningful, as this implies a different preference ordering over lotteries (and hence potentially the action that the player chooses).[4]

Assuming that the vN-M conditions hold does not impose restrictions on the player's attitudes to risk: depending on the shape of his Bernoulli payoff function, he may be **risk-averse** (concave payoff function), **risk-neutral** (linear utility function), or **risk-loving** (convex utility function). Regardless of the shape of the utility function, he chooses an action that maximizes the expected value of the utility function.

2.2.1 Appendix (Optional)

<u>Optional Exercises:</u>

(B1) Osborne 149.2

(B2) Let $w(x) = \sqrt{x}$ and $k(x) = 10\sqrt{x} - 3$ be Bernoulli payoff functions. Applying Proposition 7, determine the value of a and b when each is written as a strictly increasing linear function of the other.

[4]This is of course assuming that $c(x)$ cannot be written as a strictly increasing linear function of $b(x)$.

2.3 Homework Assignment #1, Due to TA Thiago Silva on Wednesday October 5

(1) Tadelis 1.4

(2) Tadelis 2.1

(3) Tadelis 2.2

(4) Tadelis 2.3

(5) Tadelis 2.4

3 Static Games of Complete Information

3.1 Defining Normal-Form Games

Static games are also called normal-form games, normal games, strategic-form games, strategic games, and simultaneous-move games. I prefer the latter term because it is the most informative.

The Cartesian product of a collection of sets is an important concept in giving a formal definition of a normal-form game.

Definition 29 *For any finite collection of sets A_1, A_2, $\ldots$, A_n, the* ***Cartesian product*** *$A_1 \times A_2 \times \ldots \times A_n$ is the set of all ordered "n-tuples" of elements from these sets. Formally, $A_1 \times A_2 \times \ldots \times A_n = \{(a_1, a_2, \ldots, a_n) \mid a_1 \in A_1, a_2 \in A_2, \ldots, a_n \in A_n\}$. (Sometimes we use the notation $\prod_{i=1}^n A_i$ for the Cartesian product, i.e., $\prod_{i=1}^n A_i = A_1 \times A_2 \times \ldots \times A_n$. The notation $\times_{i=1}^n A_i$ is also used.)*

(Tadelis p.47; the following definition is more general, because it allows an infinite number of players; notice how)

Definition 30 *A* ***normal-form game*** *consists of the following:*

- *A set of* ***players****, N, where $\mid N \mid \geq 2$*
- *For each player $i \in N$, a set of* ***actions*** *(also called* ***pure strategies****) S_i*
- *For each player $i \in N$, a preference relation $\succeq_i$ over the set of* ***action (or strategy) profiles*** *$S = \times_{i \in N} S_i$. Assuming that each preference relation is complete and transitive, it is represented by a payoff function $v_i : S \to \mathbb{R}$. For any action profile $s \in S$, $v_i(s)$ is player i's payoff for s. Action profiles are also sometimes called* ***outcomes****.*

(Tadelis p.51)

Definition 31 *If N is finite and each S_i is finite, then we say that the normal-form game is* ***finite****.*

Sometimes we use the notation $< N, (S_i), (u_i) >$ to denote a game (or a minor variant, as in Tadelis p.48): the set of players, the collection of the sets of actions, and the collection of the utility functions.

When we are talking about some player $i \in N$, it is common to use $-i$ to refer to the set of players *other* than i, i.e., $-i = N \setminus \{i\}$.[5] That is, $-i$ is the complement of $\{i\}$. Similarly, $S_{-i} = \times_{j \in N \setminus \{i\}} S_j$ is the set of action profiles of all players other than i. If $s_{-i} \in S_{-i}$ and $s_i \in S_i$, then (s_i, s_{-i}) or (s_{-i}, s_i) is the action profile in which player i chooses action s_i and the players in $-i$ choose their actions in s_{-i}.

Finally, if s is an action profile, then s_i is the action that player i chooses in the profile s, and s_{-i} is the action profile of the players other than i in the profile s. The usefulness of this notation will become clear in the definitions of solution concepts given later.

Some common, simple normal-form games are:

(1) Prisoner's Dilemma (PD). $N = \{1, 2\}$. $S_i = \{C, D\}$ for all $i \in N$, where C is interpreted as "cooperate" with the other player, and D is interpreted as "defect" against the other player. Note that $S = S_1 \times S_2 = \{CC, CD, DC, DD\}$. Finally, any utility functions v_1 and v_2 such that $v_1(DC) > v_1(CC) > v_1(DD) > v_1(CD)$ and $v_2(CD) > v_2(CC) > v_2(DD) > v_2(DC)$ represent the players' preferences.

(2) Stag Hunt (SH) (also called Assurance). Same as PD, but $v_1(CC) > v_1(DC) > v_1(DD) > v_1(CD)$ and $v_2(CC) > v_2(CD) > v_2(DD) > v_2(DC)$. For each player's preference

[5]Recall that $\setminus$ denotes "set subtraction": $A \setminus B = \{x \in A \mid x \notin B\}$.

ordering, the top two outcomes are flipped relative to PD.

(3) Battle of the Sexes (BoS) (also called a mixed coordination game). $N = \{1, 2\}$. $S_i = \{B, S\}$ for all $i \in N$, where B is interpreted as "Bach", and S is interpreted as "Stravinsky" (Osborne's notation). Note that $S = S_1 \times S_2 = \{BB, BS, SB, SS\}$. Finally, $v_1(BB) > v_1(SS) > v_1(BS) = v_1(SB)$ and $v_2(SS) > v_2(BB) > v_2(BS) = v_2(SB)$.

Simple games like these in which there are no more than 3 players and each player has a finite number of actions can be conveniently represented using a **payoff matrix**, which fully represents the game (N, each S_i, and each v_i).

As in single player decision theory, our goal is to predict what a rational (not yet defined in a game-theoretic context) player will do in a normal-form game. Of course, there are multiple players in such games, so our goal is to predict what each player will do, assuming that each player is rational. That is, our goal is to predict which action profile (outcome) will occur.

There are a number of different ways of thinking about (or predicting) what rational players will do in normal-form games, each requiring different assumptions about knowledge of each other's rationality and of each other's chosen action. Each way of predicting what rational players will do is called a **solution concept** (a formal definition is given in the appendix). Before getting into these, we need to define rationality in a game-theoretic setting.

3.2 Decision-Theoretic Rationality in a Game-Theoretic Setting

(Tadelis, p.72)

Definition 32 *(pure beliefs over pure strategies) A* ***belief*** *of player i is a possible profile of* *his opponents' strategies, $s_{-i} \in S_{-i}$.*

(Tadelis, p.70)

Definition 33 *(Best response in pure strategies) For player i, strategy $s_i \in S_i$ is a* ***best response*** *to his opponents' strategy profile $s_{-i} \in S_{-i}$ if $v_i(s_i, s_{-i}) \geq v_i(s_i', s_{-i})$ $\forall s_i' \in S_i$.*

(Tadelis, p.70, p.54)

Definition 34 *(Decision-theoretic rationality in a game-theoretic setting, in pure strategies and pure beliefs) Player i is* ***rational*** *if, given belief $s_{-i} \in S_{-i}$, he chooses a best response to s_{-i}.*

This is very similar to the definition of rationality in decision theory, but here given some belief about how the opponents will behave. But the basic idea is, given its beliefs, the decision-maker will choose an action that leads to a preferred outcome (action profile).

3.3 Solution Concept #1: Strict Dominance

(Tadelis p.60)

Definition 35 *(Strict dominance in pure strategies) Player i's action s_i' is* ***strictly dominated*** *by its action s_i if $v_i(s_i, s_{-i}) > v_i(s_i', s_{-i})$ for all $s_{-i} \in S_{-i}$. We also say that s_i* ***strictly dominates*** *s_i'.*

That is, s_i strictly dominates s_i' if player i is strictly better off choosing s_i than s_i' *for every combination of actions of the other players.*

(Tadelis, p.60)

Proposition 8 *A rational player never chooses a strictly dominated action.*

(Tadelis p.61)

Definition 36 *(Strict dominance in pure strategies) Player i's action s_i is **strictly dominant** if it strictly dominates every $s_i' \in S_i$ such that $s_i' \neq s_i$. That is, if for every $s_i' \in S_i$ such that $s_i' \neq s_i$, $v_i(s_i, s_{-i}) > v_i(s_i', s_{-i})$ for all $s_{-i} \in S_{-i}$.*

Proposition 9 *A player can have at most one strictly dominant action.*

Proposition 10 *If player $i \in N$ has a strictly dominant action s_i, he will choose it.*

Note that "strictly dominates" and "strictly dominated" are pairwise comparisons, i.e., are concepts that apply to any pair of actions of a player. "Strictly dominant" is a player-specific concept, in that we ask the question of whether a player has a strictly dominant action. Of course, in a game in which each player has just 2 actions, these concepts are identical: if s_i strictly dominates s_i', s_i is strictly dominant. (Also note that we can ask the question of whether a player has a strictly dominated action without being explicit about which action strictly dominates it, and in this sense can be thought of as a player-specific concept.)

In PD, each player has a strictly dominant action, namely D. Thus, if both players are rational, both will choose D. The dilemma is that each player strictly prefers the action profile CC to DD, i.e., *CC strictly Pareto dominates DD.*

(Tadelis p.57)

Definition 37 *An action profile $s \in S$ **Pareto dominates** action profile $s' \in S$ if $v_i(s) \geq v_i(s')$ $\forall i \in N$, and $v_i(s) > v_i(s')$ for at least one $i \in N$. We also say that s' is **Pareto***

dominated *by s. (We sometimes use the terms* ***strictly Pareto dominates*** *and* ***strictly Pareto dominated*** *if all of the inequalities hold strictly, i.e., $v_i(s) > v_i(s')$ $\forall i \in N$.) An action profile is* ***Pareto optimal*** *if it is not Pareto dominated by any other action profile.*

Pareto-optimality is an efficiency concept, not a solution concept (i.e., a prediction concept). If a predicted outcome is not Pareto optimal, we lament this because there are unrealized gains that wouldn't make anyone worse off, but we don't therefore just naively change our prediction.

If a player has a strictly dominant action, then from that player's perspective the decision setting is virtually a single player decision theory setting rather than a game-theoretic setting, because the player doesn't really have to think about how the other players are likely to behave. That player will choose its strictly dominant action regardless of how it expects the other players to behave. Indeed, it doesn't even need to have an expectation of how the other players will behave.

If *all* players have strictly dominant actions, as in PD, then we have a unique prediction of what will happen in the game, i.e., the action profile that will occur if all players are rational.

(Tadelis p.61)

Definition 38 *The strategy profile $s^D \in S$ is a* ***strictly dominant strategy equilibrium (SDSE)*** *if, for each player $i \in N$, $s_i^D \in S_i$ is a strictly dominant strategy for player i.*

Proposition 11 *A game can have at most one SDSE.*

(Tadelis, p.70)

Proposition 12 *If the strategy profile $s^* \in S$ is an SDSE, then for all $i \in N$, s_i^* is a best response to s_{-i}^*.*

In most interesting games, no player has a strictly dominant strategy, and hence there is no SDSE. PD is a rare exception. It is the most basic solution concept, because all it requires is that all players are rational, but is rarely applicable.

3.4 Solution Concept #2: Iterated Elimination of Strictly Dominated Strategies (IESDS)

If each player has a strictly dominant strategy, wonderful. We have a clear prediction that just assumes that each player is rational. But this is rare. In fact, other than PD, any such game will be a contrived game.

Consider a more frequent situation, where at least 1 player has at least 1 strictly dominated action. We know that if each player is rational, none will play a strictly dominated action, and hence we can eliminate these from consideration. This might eliminate a significant number of actions, and hence allow us to focus on a reduced set of strategy profiles as being possible outcomes of the game among rational players.

Can we go any further? Suppose we assume not only that each player is rational, but that each player knows that everyone is rational. Then each player considers a "reduced game" in which all strictly dominated strategies, for all players, are eliminated. If I know that all of my opponents are rational, I know that they won't consider playing their strictly dominated actions. But so far this doesn't get us any further; it just gets us to the point where all players are considering the "reduced game" where not only their own strictly dominated actions are eliminated, but those of every other player as well.

But suppose we assume not only that every player knows that every player is rational, but also that every player knows that every player knows that every player is rational. Then I know that all of my opponents are also considering this "reduced game." And because we are all rational and know that we are all rational, I know that we will all eliminate any strictly dominated actions within this "reduced game." That is, it is possible that within the reduced game some actions are strictly dominated, even though they were not in the original game. This may result in a "second reduced game" that is smaller than the "reduced game."

If we assume **common knowledge of rationality (CKR)**, then the players will conduct this process until they arrive at a game that cannot be reduced any further, i.e., in which no one has a strictly dominated action. This process is known as **iterated elimination of strictly dominated strategies (IESDS)**, and is well-described by the algorithm in Tadelis, p.65.

(Tadelis p.45)

Definition 39 *An event E is **common knowledge** if (1) everyone knows E, (2) everyone knows that everyone knows E, and so on ad infinitum.*

Definition 40 ***CKR** means that (1) every player knows that every player is rational, (2) every player knows that every player knows that every player is rational, and so on ad infinitum.*

(Tadelis, p.65)

Definition 41 *The strategy profile $s^{ES} \in S$ is an **iterated-elimination equilibrium (IEE)** if, for each player $i \in N$, $s_i^{ES} \in S_i$ survives the process of IESDS.*

At least one IEE exists for every game, unlike with SDSE. However, in games with no strictly dominated actions, every strategy profile is an IEE, and hence it doesn't predict much, because it doesn't rule out anything from happening. And unfortunately, in most interesting games there are no strictly dominated actions, and hence IEE doesn't have any bite in most games.

Also note that the "epistemic conditions" for IEE are much higher than for SDSE: not only that every player is rational, but also CKR.[6]

(Tadelis, p.68)

Proposition 13 *If s^* is an SDSE, then s^* is the unique IEE.*

(Tadelis, p.70)

Proposition 14 *If the strategy profile $s^* \in S$ is the unique IEE, then for all $i \in N$, s_i^* is a best response to s_{-i}^*.*[7]

3.5 Solution Concept #3: Rationalizability

The notion of a correspondence generalizes the notion of a function.

Definition 42 *A **correspondence** C from a set X into a set Y (i.e., $C : X \rightarrow\rightarrow Y$) is a rule that assigns to each $x \in X$ a set $C(x)$, where $C(x) \subseteq Y$. We call X the **domain set** (or **preimage set**) and Y the **range set** (or **image set**).*

A function is a special type of correspondence in which, for every $x \in X$, $C(x)$ has exactly one element.

(Tadelis, p.72)

[6] The branch of game theory that deals with the rationality and knowledge requirements for each solution concept is called *epistemic game theory*, and an important paper here is Aumann and Brandenburger (1995).

[7] Later on, we will have the following equivalent statement: "If the strategy profile $s^* \in S$ is the unique IEE, then s^* is a Nash equilibrium."

Definition 43 *(Best response correspondence in pure strategies) For player $i \in N$, the* ***best response correspondence*** $BR_i : S_{-i} \to\to S_i$ *is defined as follows: for any $s_{-i} \in S_{-i}$, $BR_i(s_{-i}) = \{s_i \in S_i \mid v_i(s_i, s_{-i}) \geq v_i(s'_i, s_{-i})\ \forall s'_i \in S_i\}$. That is, $BR_i(s_{-i})$ is the set of all best responses of player i to s_{-i}.*

(Tadelis, p.70)

Definition 44 *(Alternative definition of decision-theoretic rationality in a game-theoretic setting, in pure strategies and pure beliefs) Player i is* ***rational*** *if, given belief $s_{-i} \in S_{-i}$, he chooses some $s_i \in BR_i(s_{-i})$.*

(Tadelis, p.73)

Definition 45 *(Never-best response in pure strategies) For player i, strategy $s_i \in S_i$ is a* ***never-best response*** *if, for every belief $s_{-i} \in S_{-i}$, $s_i \notin BR_i(s_{-i})$.*

Proposition 15 *A rational player never chooses a never-best response.*

Rationalizability (Tadelis, p.73): We can engage in the process of **iterated elimination of never-best responses**. (Question to ponder: is the following statement true? "A strategy $s_i \in S_i$ is a never-best response if and only if it is strictly dominated.") Whatever strategies remain can be called **rationalizable**. Like IESDS, this relies on assuming not only that every player is rational, but also that this is common knowledge. A formal definition requires mixed strategies, which we will encounter later.

Another way of thinking about rationalizability in 2-player games: player 1's strategy $s_1 \in S_1$ is rationalizable if and only if it is a best response to some $s_2 \in S_2$, which is a best response to some $s_1' \in S_1$ (it is possible that $s_1' = s_1$), which is a best response to some $s_2' \in S_2$ (it is possible that $s_2' = s_2$), and so on *ad infinitum.*

Player 1 can rationalize playing such a strategy, when it is common knowledge that both players are rational. Player 1 might not be choosing a best response to what player 2 is *actually* choosing, but without knowing for sure what player 2 is choosing, player 1 can *rationalize* choosing s_1 under CKR. It is a best response to something that player 2 might be doing, that is a best response to something that player 1 might be doing, that is a best response to something that player 2 might be doing, and so on *ad infinitum.*

(Tadelis, p.70)

Proposition 16 *If $s_i \in S_i$ is a strictly dominated strategy for player i, then it is a never-best response.*

3.6 Solution Concept #4: Weak Dominance and Iterated Elimination of Weakly Dominated Strategies (IEWDS)

We have the following weaker notion of dominance.

(Tadelis p.63; this is slightly different from Tadelis's definition; think about why it is more commonly defined this way rather than the way Tadelis defines it)

Definition 46 *(Weak dominance in pure strategies) Player i's action s_i' is* ***weakly dominated*** *by its action s_i if $v_i(s_i, s_{-i}) \geq v_i(s_i', s_{-i})$ for all $s_{-i} \in S_{-i}$, and $v_i(s_i, s_{-i}) > v_i(s_i', s_{-i})$ for at least one specific $s_{-i} \in S_{-i}$. We also say that s_i* ***weakly dominates*** *s_i'.*

That is, s_i does at least as well as s_i' no matter what the other players do, and for at least one specific action profile of the other players, s_i does strictly better than s_i'.

Unlike with strictly dominated actions, we can't conclusively say that a rational player i will not choose a weakly dominated action, because it can be a best response to some belief $s_{-i} \in S_{-i}$. However, we have the following.

Proposition 17 *If $s'_i \in S_i$ is weakly dominated, then there exists no $s_{-i} \in S_{-i}$ for which $BR_i(s_{-i}) = \{s'_i\}$.*

That is, a weakly dominated action can never be the *unique* best response to some belief $s_{-i} \in S_{-i}$.

Although we can't conclusively rule out a rational player choosing a weakly dominated action s'_i, it is not clear why a rational player *would* choose s'_i, unless he is absolutely certain of the belief (if any) to which s'_i is a best response. Choosing the action that weakly dominates it is a safer bet, especially when there is uncertainty as to what the other players are going to do.

Definition 47 *(Weak dominance in pure strategies) Player i's action $s_i \in S_i$ is **weakly dominant** if it weakly dominates every $s'_i \in S_i$ such that $s'_i \neq s_i$.*

Proposition 18 *A player can have at most one weakly dominant action. (Note that this is not true under Tadelis's definition of "weakly dominates".)*

If a rational player has a weakly dominant action, we can't conclusively say that it will choose it, although it is not clear why it would *not* choose it, as it will be a safer bet than any other action, especially when there is uncertainty as to what the other players are going to do.

(Tadelis, p.76)

Definition 48 *The strategy profile* $s^W \in S$ *is a* ***weakly dominant strategy equilibrium (WDSE)*** *if, for each player* $i \in N$, $s_i^W \in S_i$ *is a weakly dominant strategy for player* i.

Proposition 19 *A game can have at most one WDSE. (Again, this is not true under Tadelis's definition of "weakly dominates".)*

Just like with SDSE, a game can have no WDSE. Even if a game has a WDSE, we cannot conclusively say that it will be the outcome if every player is rational. But again, it is hard to see why it would *not* be the outcome, as each player's best bet is its weakly dominant action.

We can conduct the process of **iterated elimination of weakly dominated strategies (IEWDS)** (and define a corresponding equilibrium, analogous to IEE). However, this is less compelling than IESDS because we cannot conclusively say that a rational player will eliminate its weakly dominated actions from consideration. Therefore, CKR does *not* imply that a strategy profile consisting of strategies that survive IEWDS will occur; players can't assume that their rational opponents will eliminate weakly dominated actions from consideration.

3.7 Assessing the Solution Concepts So Far

SDSE is a very attractive solution concept, because it just relies on each player being rational (i.e., has minimal "epistemic conditions"). However, it is rare to find an interesting game in which even a single player has a strictly dominant action, much less every player. PD is really the only interesting game for which an SDSE exists.

IESDS (and IEE) has stronger epistemic conditions than SDSE, namely CKR. On the plus side, at least one IEE exists for every game. On the downside, for most interesting games, the set of IEE consists of the set of *all* strategy profiles, because no player has any strictly dominated action. So this type of equilibrium always exists, but usually does not predict anything because it does not rule out anything.

Rationalizability, or iterated elimination of never-best responses, suffers the same problems as IESDS, because it also usually does not rule out much. In most interesting games, players have no never-best responses.

We will see that many interesting games do have weakly dominated actions, but not weakly dominant actions. Hence, WDSE rarely exist. IEWDS can sometimes eliminate some strategies from consideration, but usually not many. Moreover, IEWDS does not have firm epistemic justification: we cannot conclusively say that rational players will eliminate weakly dominated actions from consideration, and CKR does not imply IEWDS.

The limitations of all of these solution concepts suggest a need for an alternative solution concept that (a) always exists (i.e., at least one equilibrium always exists), (b) usually eliminates at least a significant number of strategy profiles from consideration as predictions of what will happen, and (c) has epistemic conditions that clearly justify it. These considerations will lead us to Nash equilibrium as the preferred solution concept for normal-form games.

3.8 Appendix (Optional)

Definition 49 *Let G be the set of all normal-form games, and for any normal-form game $g \in G$, let S_g be the set of strategy profiles of g. Let $S = \bigcup_{g \in G} 2^{S_g}$ be the set containing every subset of S_g, for every $g \in G$.*[8] *A* ***solution concept*** *is a function $E : G \to S$, where*

[8] Let A be a set. The **power set** of A, denoted by 2^A, is the set of all subsets of A. That is, $2^A = \{X \mid X \subseteq A\}$. For example, if $A = \{1, 2\}$, then $2^A = \{\emptyset, \{1\}, \{2\}, \{1, 2\}\}$.

*for any normal-form game $g \in G$, $E(g) \subseteq S_g$ is the set of all strategy profiles of g that are **equilibria** of the game (under that solution concept).*[9]

That is, a solution concept is a function E that assigns to any normal-form game g a subset of the set of strategy profiles of the game, where the strategy profiles in this subset are the equilibria of the game (under that solution concept).

Desirable qualities of solution concepts are that: (a) for each $g \in G$, $E(g) \neq \emptyset$ (i.e., an equilibrium always exists), (b) for most $g \in G$, $E(g) \subset S_g$ (i.e., the set of equilibria is usually smaller than the set of strategy profiles, i.e., the solution concept usually rules out at least some strategy profiles as predictions), and (c) there are clear epistemic conditions that justify the solution concept, i.e., if these conditions hold, we can reasonably expect the players to play according to one of the equilibria (although we may not be able to clearly determine which one, if there are multiple equilibria).

Optional Exercises:

(C1) For each of the following, is the statement true? If so, prove it. If not, provide a counter-example to show that it is not true in general.

(a) If the strategy profile s^* is an SDSE, then it is a Pareto-optimal strategy profile.

(b) If the strategy profile s^* is the unique IEE, then s^* is an SDSE. (This is the converse of Tadelis's Proposition 4.2 on p.68)

(c) Consider a finite game in which there is a strategy profile s^* that *strictly* Pareto-dominates every other strategy profile. Then s^* is the unique IEE.

[9] The reason this is a function instead of a correspondence is that the range is the set of all *subsets* of the set of strategy profiles for all normal-form games, and E assigns exactly one such subset (the set of equilibria) to every game $g \in G$. That is, $E(g)$ is a set, but is a specific set from a range of possible sets, and hence E is a function rather than a correspondence. Admittedly, this can be confusing.

(d) Consider a finite game in which there is a strategy profile s^* that *strictly* Pareto-dominates every other strategy profile. Then s^* is an IEE.

(e) If strategy $s_i \in S_i$ is a never-best response, then s_i is strictly dominated.

(f) If s_i strictly dominates s_i', then it weakly dominates it.

(g) If s_i weakly dominates s_i', then it strictly dominates it.

(h) If s_i is strictly dominant, then it is weakly dominant.

(i) If s_i is weakly dominant, then it is strictly dominant.

3.9 Homework Assignment #2, Due to TA Thiago Silva on Wednesday October 12

(1) Tadelis 3.3

(2) Tadelis 3.7

(3) Tadelis 4.1

(4) Tadelis 4.3

(5) Tadelis 4.5

3.10 The Main Solution Concept: Nash Equilibrium (NE)

(Tadelis, p.80)

Definition 50 *(Nash equilibrium in pure strategies) A strategy profile* $s^* \in S$ *is a* ***Nash equilibrium (NE)*** *if, for each player* $i \in N$, $v_i(s_i^*, s_{-i}^*) \geq v_i(s_i, s_{-i}^*)$ *for all* $s_i \in S_i$.

That is, a NE is strategy profile in which each player is choosing a best response to what the other players are doing. In fact:

(Osborne, p.36)

Proposition 20 *A strategy profile* $s^* \in S$ *is a NE if and only if for each* $i \in N$, $s_i^* \in BR_i(s_{-i}^*)$.

[Note: "if and only if" results give an alternative way of defining a concept. That is, we could have defined a NE using the above proposition, and then turned the above definition into an "if and only if" proposition.]

Informal definition: "A NE is a strategy profile in which no player can increase its payoff by unilaterally changing its action (strategy)."

NE is intuitively thought of as a "stability concept." That is, if the players get to a NE, no one has an incentive to change their behavior, and hence they will remain there. This is not true of strategy profiles that are not NE.

Similarly, if I as an outside analyst predict to the players that a strategy profile will occur that happens to be a NE, I have good reason to expect that this strategy profile *will* in fact occur, because no one has an incentive to choose differently. Again, this is not true of strategy profiles that are not NE; at least one player can benefit by choosing differently, and hence I won't expect that strategy profile to actually occur.

The "epistemic conditions" for NE are: (1) each player is rational, i.e., chooses a best response to its belief about the opponents' strategies, and (2) each player's belief about its opponents' strategies is correct. Note that common knowledge of rationality (CKR), which is a strong condition, is *not* needed. However, condition (2) is fairly strong. Experimental evidence tends to suggest that experienced players play according to Nash equilibria but inexperienced players often do not, and one explanation of this data is that experienced players know how their opponents will choose and hence (2) is satisfied for them, but not necessarily for inexperienced players.

Finding NE in payoff matrices: (i) directly, and (ii) more efficiently, marking best responses using dots. [best-response dots also allow us to efficiently determine issues like strict dominance; see the optional exercises in the appendix]

[Do some examples]

[Pareto-dominance as an **equilibrium selection criterion** when multiple NE exist]

[Schelling's notion of **focal points** as another selection criterion in **coordination games**]

[Matching pennies as a game that has no NE in pure strategies]

Two types of NE:

(Osborne, p.33)

Definition 51 *(Strict and non-strict NE in pure strategies) A strategy profile $s^* \in S$ is a **strict NE** if, for each player $i \in N$, $v_i(s_i^*, s_{-i}^*) > v_i(s_i, s_{-i}^*)$ for all $s_i \in S_i$ such that $s_i \neq s_i^*$. A NE that is not a strict NE is called a **non-strict NE**.*

[Tips for finding the NE of games that can't be shown using payoff matrices and hence the cook-book, algorithmic best-response-dots method can't be applied. (1) Begin by considering *simple* strategy profiles and check whether they are NE, and this will guide you towards other, more complicated strategy profiles that you need to check. (2) Often useful to introduce a new variable to help you conveniently list strategy profiles. For example, in an n-player game in which each player has 2 actions, C ("contribute") or NC ("not contribute"), "Let x be the number of players choosing C. Consider the simple strategy profile $x = 0 \ldots$ Consider the simple strategy profile $x = n \ldots$"]

[Examples: in the optional exercises in the appendix]

[Finding NE of games with a continuum of actions: Cournot Duopoly, auctions, Hotelling-Downs Model of Electoral Competition] [some require calculus (FOC and SOC) to find the NE, some don't]

The following are some results linking some earlier solution concepts with NE. [Optional exercise: Each of the following is a $P \Rightarrow Q$ result. Think about whether the converse holds, i.e., $P \Leftarrow Q$.]

(McCarty and Meirowitz, p.99)

Proposition 21 *If a strategy profile $s^* \in S$ is a NE, then it is also an IEE (i.e., no s_i^* is eliminated through IESDS).*

(Tadelis, p.80)

Proposition 22 *If a strategy profile $s^* \in S$ is either (a) a strictly dominant strategy equilibrium (SDSE), (b) the unique IEE (i.e., the unique survivor of IESDS), or (c) the unique rationalizable strategy profile, then s^* is the unique NE.*

(Osborne, p.45)

Proposition 23 *If $s_i' \in S_i$ is a strictly dominated action of player i, then there exists no NE containing s_i'.*

(Osborne, p.47)

Proposition 24 *If $s_i' \in S_i$ is a weakly dominated action of player i, then there exists no strict NE containing s_i'.*

3.11 Appendix (Optional)

Optional Exercises:

(D1) Give an elegant, self-contained formal definition of a non-strict NE, i.e., a definition that makes no reference to strict NE.

(D2) Give a simple payoff matrix in which a weakly dominated action is part of a NE.

(D3) Give a simple payoff matrix of a finite 2-player game that has a unique NE, in which each player's action is weakly dominated.

(D4) Is the following result true? If so, prove it. If not, provide a counter-example to show that it is not true in general.

(a) "In a finite 2-player game in which some player i has a weakly dominant action $s_i \in S_i$, there exists at least 1 NE containing s_i."

(b) "An action $s_i \in S_i$ is a strictly dominant action of player i if and only if for every $s_{-i} \in S_{-i}$, $BR_i(s_{-i}) = \{s_i\}$."

(c) "An action $s_i \in S_i$ is a weakly dominant action of player i if and only if for every $s_{-i} \in S_{-i}$, $s_i \in BR_i(s_{-i})$." [Using the definition of weak dominance in these notes rather than Tadelis's definition.]

(d) The converse of Proposition 21

(e) The converse of Proposition 22

(f) The converse of Proposition 23

(g) The converse of Proposition 24

(D5) Two candidates, A and B, compete in an election. There are n citizens, where $n \geq 3$ is an odd integer. Each citizen can vote for A or vote for B (no abstentions are allowed). The candidate who obtains the most votes wins. Each citizen either strictly prefers A winning to B winning, or strictly prefers B to A. A majority of citizens strictly prefer A to B.

(a) Formulate this as a normal-form game.

(b) Are there are any strictly or weakly dominated actions?

(c) Find all of the NE in pure strategies. For each one, determine whether it is a strict or non-strict NE.

(d) If there are multiple NE, is there a reasonable basis for selecting one of the them as the most likely to occur? Explain.

(D6) Two candidates, A and B, compete in an election. Of the $n \geq 2$ citizens, k support candidate A and m $(= n - k)$ support candidate B. Each citizen decides whether to vote, at a cost, for the candidate she supports, or to abstain. A citizen's payoff for her preferred candidate winning is 2, is 1 for a tie, and is 0 if her preferred candidate loses. A citizen who votes also pays a cost c, where $0 < c < 1$.

(a) For $k = m = 1$, draw the payoff matrix and find all of the NE in pure strategies. For each one, determine whether it is a strict or non-strict NE. Are there any strictly or weakly dominated actions?

(b) For $k = m$, find all of NE in pure strategies. For each one, determine whether it is a strict or non-strict NE. Are there any strictly or weakly dominated actions?

(c) For $k < m$, find all of the NE in pure strategies. For each one, determine whether it is a strict or non-strict NE. Are there any strictly or weakly dominated actions?

(D7) Each of $n \geq 2$ people chooses whether to contribute a fixed amount toward the provision of a public good. The good is provided if and only if at least k people contribute, where $2 \leq k \leq n$; if it is not provided, contributions are not refunded. Each person assigns payoffs as follows: 4 if the good is provided and I don't contribute, 3 if the good is provided and I contribute, 2 if the good is not provided and I do not contribute, and 1 if the good is not provided and I contribute.

(a) Formulate this as a normal-form game.

(b) Are there are any strictly or weakly dominated actions?

(c) Find all of the NE in pure strategies. For each one, determine whether it is a strict or non-strict NE.

(d) If there are multiple NE, is there a reasonable basis for selecting one of the them as the most likely to occur? Explain.

(D8) Consider an $n-$player Stag Hunt game, where each player can choose C (cooperate; pursue the stag) or D (defect; catch a rabbit instead), and $n \geq 3$. The stag is only caught if at least m hunters, where $2 \leq m < n$, pursue it; if fewer pursue it, the pursuers don't catch anything and go hungry. A captured stag is only shared by the hunters who catch it. Find all of the pure-strategy NE of the game in the following two scenarios:

(a) Each hunter prefers the fraction $\frac{1}{n}$ of the stag to a rabbit.

(b) Each hunter prefers the fraction $\frac{1}{k}$ of the stag to a rabbit, but prefers a rabbit to any smaller fraction of the stag, where k is an integer with $m \leq k \leq n$.

(D9) Find the pure-strategy NE of the following game: $N = \{1, 2\}$, $S_1 = S_2 = [0, \infty)$,

$v_1(s_1, s_2) = s_1(s_2 - s_1)$, and $v_2(s_1, s_2) = s_2(1 - s_1 - s_2)$.

(D10) Two people are engaged in a joint project. If person i (for $i = 1, 2$) puts in the effort $s_i \in [0, 1]$, which costs her $c(s_i)$, the outcome of the project is worth $f(s_1, s_2)$. The worth of the project is split equally between the two people, regardless of their effort levels. In each of the following scenarios, (i) formulate this as a normal-form game; (ii) find all of the pure-strategy NE; (iii) is there a pair of effort levels that yields higher payoffs for both players than do the NE effort levels?

(a) $f(s_1, s_2) = 3s_1 s_2$ and $c(s_i) = s_i^2$ for $i = 1, 2$.

(b) $f(s_1, s_2) = 4s_1 s_2$ and $c(s_i) = s_i$ for $i = 1, 2$.

3.12 Homework Assignment #3, Due to TA Thiago Silva on Wednesday October 19

(1) Tadelis 5.5 (note that this is the same game as an earlier HW problem, Tadelis 3.7). Also:

(c) For each NE, state whether it is a strict or non-strict NE.

(d) Does any player have a strictly dominated or weakly dominated action? If so, what?

(e) Which, if any, strategy profiles are Pareto optimal?

(2) (a) Find the pure-strategy NE for the game in Tadelis 4.3 (which was a previous HW problem). For each NE, state whether it is a strict or non-strict NE.

(b) Discuss how your answer to part (a), and your answer to Tadelis 4.3, illustrates Propositions 21 and 22 in the notes.

(c) Which, if any, strategy profiles are Pareto optimal?

(3) (a) Find the pure-strategy NE for the game in Tadelis 4.5 (which was a previous HW problem). For each NE, state whether it is a strict or non-strict NE.

(b) Discuss how your answer to part (a), and your answer to Tadelis 4.5, illustrates Propositions 21 and 22 in the notes.

(c) Which, if any, strategy profiles are Pareto optimal?

(4) Optional exercise (D6) above—this is no longer optional.

(5) Optional exercise (D9) above—this is no longer optional.

3.13 Mixed Strategies and Mixed Strategy Nash Equilibrium (MSNE)

Now we want to allow players to probabilistically choose among their actions.

(Tadelis, p.102)[10]

Definition 52 *(Mixed strategies with a finite set of actions) Suppose player i has a finite set of actions $S_i = \{s_{i1}, s_{i2}, \dots, s_{im}\}$. The set $\Delta S_i = \{\sigma_i : S_i \to [0,1] \mid \sum_{s_i \in S_i} \sigma_i(s_i) = 1\}$ is the set of probability distributions over S_i, and is called player i's **set of mixed strategies**. Any $\sigma_i \in \Delta S_i$ is a **mixed strategy** of player i. We can write $\sigma_i = (\sigma_i(s_{i1}), \sigma_i(s_{i2}), \dots, \sigma_i(s_{im}))$, where $\sigma_i(s_i)$ is the probability that player i chooses action $s_i \in S_i$.*

Even with just 2 actions, an infinite set of mixed strategies.

Now distinction between actions and strategies becomes important.

Now we have in mind that the utility function $v_i : S \to \mathbb{R}$ over action profiles is a Bernoulli utility function whose expected value captures the actor's preferences over lotteries over action profiles (since probabilities are now involved).

Definition 53 *A mixed strategy σ_i is called a **pure strategy** if it assigns probability 1 to some action $s_i \in S_i$, and probability 0 to all other elements of S_i. This pure strategy is usually just denoted s_i.*

Therefore, pure strategies are special cases of mixed strategies.

(Tadelis, p.104)

Definition 54 *Let $\sigma_i \in \Delta S_i$ be a mixed strategy of player i. We say that action $s_i \in S_i$ is **in the support of** σ_i if σ_i assigns positive probability to s_i, i.e., if $\sigma_i(s_i) > 0$.*

[10] In the following definition, it may be useful to refer to Definition 24 above.

Obviously, if only a single action is in the support of σ_i, then σ_i is a pure strategy.

Definition 55 *Let $\times_{i\in N}\Delta S_i$ be the Cartesian product of each player's set of mixed strategies. Then any $\sigma \in \times_{i\in N}\Delta S_i$ is a* ***profile of mixed strategies****. If $N = \{1, 2, \ldots, n\}$, then we typically denote a profile of mixed strategies by $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_n)$, where σ_i is the mixed strategy of player i.*

We use notation that is analogous to the notation introduced on p.21 for pure strategies (actions). Note that $\times_{j\in N\setminus\{i\}}\Delta S_j$ is the set of mixed strategy profiles of all players *other* than i (with pure strategies, we denoted this by S_{-i}).[11] If $\sigma_{-i} \in \times_{j\in N\setminus\{i\}}\Delta S_j$ and $s_i \in S_i$, then (s_i, σ_{-i}) or (σ_{-i}, s_i) is the mixed strategy profile in which player i chooses action s_i and the players in $-i$ choose their mixed strategies in σ_{-i}.

If $\sigma_{-i} \in \times_{j\in N\setminus\{i\}}\Delta S_j$ and $\sigma_i \in \Delta S_i$, then (σ_i, σ_{-i}) or (σ_{-i}, σ_i) is the mixed strategy profile in which player i chooses mixed strategy σ_i and the players in $-i$ choose their mixed strategies in σ_{-i}.

Finally, if σ is a mixed strategy profile, then σ_i is the mixed strategy that player i chooses in the profile σ, and σ_{-i} is the mixed strategy profile of the players other than i in the profile σ.

Now we can talk about a player's expected payoff for a profile of mixed strategies.

(Tadelis, p.105; note that Tadelis confusingly uses v_i for both the utility function as well as the expected utility function, and we will instead use EU_i for the latter.)

[11] Note that Tadelis sometimes uses the notation ΔS_{-i} for $\times_{j\in N\setminus\{i\}}\Delta S_j$, e.g., in his Definition 6.5 on p.105. But this is confusing, since S_{-i} is the set of action profiles of the other players, and hence ΔS_{-i} is the set of probability distributions over the action profiles of the other players, which is not the same as the set of mixed strategy profiles of the other players. Any specific mixed strategy profile of the other players implies a specific probability distribution over the action profiles of the other players, but these are different objects. So we will use the more tedious $\times_{j\in N\setminus\{i\}}\Delta S_j$ rather than the more succinct but misleading ΔS_{-i}.

Definition 56 *Let* $v_i : S \to \mathbb{R}$ *be player* i*'s Bernoulli payoff function over action profiles* $s \in S$*. Then the* ***expected payoff*** *of player* i *when he chooses the pure strategy* $s_i \in S_i$ *and the other players choose the mixed strategy profile* $\sigma_{-i} \in \times_{j \in N \setminus \{i\}} \Delta S_j$ *is*

$EU_i(s_i, \sigma_{-i}) = \sum_{s_{-i} \in S_{-i}} \sigma_{-i}(s_{-i}) v_i(s_i, s_{-i})$*, where* $\sigma_{-i}(s_{-i})$ *is the probability of* s_{-i} *occurring given* σ_{-i}*.*

Similarly, player i*'s expected payoff when he chooses* $\sigma_i \in \Delta S_i$ *and the others choose* $\sigma_{-i} \in \times_{j \in N \setminus \{i\}} \Delta S_j$ *is*

$EU_i(\sigma_i, \sigma_{-i}) = \sum_{s_i \in S_i} \sigma_i(s_i) EU_i(s_i, \sigma_{-i})$*.*

This looks complicated, but is just the regular notion of expected utility: given a mixed strategy profile σ, player i's expected payoff is calculated by multiplying the probability of each action profile with the payoff for that action profile, doing this for each action profile, and then adding them all up.

We are now in a position to define a mixed strategy Nash equilibrium.

(Tadelis, p.107)

Definition 57 *(Nash equilibrium in mixed strategies) A profile of mixed strategies* $\sigma^* \in \times_{i \in N} \Delta S_i$ *is a* ***mixed strategy Nash equilibrium*** *(MSNE) if, for each player* $i \in N$*,* $EU_i(\sigma_i^*, \sigma_{-i}^*) \geq EU_i(\sigma_i', \sigma_{-i}^*)$ *for all* $\sigma_i' \in \Delta S_i$*, where* $EU_i(\sigma)$ *is the expected value of* v_i *under* σ*.*

Definition 58 *(Best response correspondence in mixed strategies) For player* $i \in N$*, the* ***best response correspondence*** $BR_i : \times_{j \in N \setminus \{i\}} \Delta S_j \to\to \Delta S_i$ *is defined as* $BR_i(\sigma_{-i}) = \{\sigma_i \in \Delta S_i \mid EU_i(\sigma_i, \sigma_{-i}) \geq EU_i(\sigma_i', \sigma_{-i})\ \forall \sigma_i' \in \Delta S_i\}$ *for any* $\sigma_{-i} \in \times_{j \in N \setminus \{i\}} \Delta S_j$*. Any* $\sigma_i \in BR_i(\sigma_{-i})$ *is a* ***best response*** *of player* i *to* σ_{-i}*.*

Proposition 25 *A mixed strategy profile σ is a MSNE if and only if, for each $i \in N$, $\sigma_i \in BR_i(\sigma_{-i})$.*

Plot the best response correspondences to find all of the MSNE of 2 by 2 games.

(Osborne p.116; this subsumes Tadelis's Proposition 6.1 on p.108 and also includes more)

Proposition 26 *Suppose S_i is finite for each $i \in N$. Then a mixed strategy profile σ is a MSNE if and only if, for each $i \in N$, the following 2 conditions hold: (1) the expected payoff, given σ_{-i}, to every action to which σ_i assigns positive probability is the same, and (2) the expected payoff, given σ_{-i}, to every action to which σ_i assigns probability 0 is at most the expected payoff to any action to which σ_i assigns positive probability. (Also, each player's expected payoff in a MSNE σ is her expected payoff, given σ_{-i}, to any of her actions that σ_i assigns positive probability to.)*

This result gives a method for finding *all* of the MSNE of games in which each player has a finite number of actions.

Note that this result implies that a MSNE in which at least one player is genuinely mixing (i.e., at least 2 actions lie in the support of her mixed strategy) is never a *strict* NE; she has an infinity of alternative best responses to what the others are doing. In particular, any probability distribution over the actions in the support of her mixed strategy is a best response to what the others are doing.

(Osborne p.122)

Proposition 27 *Suppose that the action profile $s \in S$ is a NE of a normal-form game when mixing is not allowed. Then the mixed strategy profile in which each player $i \in N$ chooses s_i with probability 1 is a MSNE of the same normal-form game when mixing is allowed.*

This result tells us that finding pure-strategy NE is not a useless exercise now that we recognize that players may mix; these are still NE even when mixing is allowed. The change is simply that there may be *additional* NE that we previously didn't recognize.

(Tadelis, p.117)

Proposition 28 *Every normal-form game in which each player has a finite number of actions has at least one Nash equilibrium (in pure or mixed strategies).*

This result, known as an *existence* result, doesn't provide any guidance as to how to find the NE of a normal form game in which each player has a finite number of actions, but does tell us that at least 1 NE exists.

3.13.1 Optional Exercise

(E1) Recall optional exercise (D6) on p.41. (Each citizen votes for her preferred candidate or abstains, and voting is costly.) Suppose that $k \leq m$. Show that there is a value of p between 0 and 1 such that there is a MSNE in which every supporter of candidate A votes with probability p, k supporters of candidate B vote with certainty, and the remaining $m-k$ supporters of candidate B abstain. How do the probability p that a supporter of candidate A votes and the expected number of voters ("turnout") depend upon c? (HINT: if every supporter of candidate A votes with probability p, then the probability that exactly $k-1$ of them vote is $kp^{k-1}(1-p)$.)

3.14 Homework Assignment #4, Due to TA Thiago Silva on Wednesday October 26

(1) Optional exercise (D9) above—this is no longer optional.

(2) Tadelis 6.5

(3) Tadelis 6.7

(4) Tadelis 6.9

3.15 UPDATED Homework Assignment #4, Due to TA Thiago Silva on Wednesday October 26

[Note: If you already started the previous version of HW#4, your work isn't wasted as those problems will be included on HW#5.]

(1) Optional exercise (D9) above.

(2) Optional exercise (D8).

(3) Optional exercise (D10), just part (b). Note that in answering part (iii), it may help to find the (s_1, s_2) that maximizes the sum of their payoff functions, just like we did in class with the Tragedy of the Commons Game.

3.16 Additional Topics Regarding Mixing and MSNE

3.16.1 A Generalized Notion of Beliefs

Recall that when we wanted to define what it means for an actor to be rational in a game-theoretic setting, we said that the actor is rational if he chooses a best response to his belief about what the other actors are going to do. We defined a belief to be one specific strategy profile of the other players, i.e., one specific $s_{-i} \in S_{-i}$.

Now that we have expanded our thinking to allow the players to mix, it seems natural that we should define a belief of a player in a more general way, namely a belief should be a *probability distribution* over the strategy profiles of the other players. This is because a player's belief need not assign probability 1 to one specific strategy profile of the other players; the player may be uncertain about how the others are going to choose, but can at least assign probabilities to the different possibilities. This uncertainty may be due to the player believing that the others are going to choose mixed strategies, or it may just be uncertainty in general.

(Tadelis, p.105; Osborne, p.379)

Definition 59 *A **belief** for player i is given by a probability distribution $\pi_i \in \Delta S_{-i}$ over the strategy profiles of the other players. For any $s_{-i} \in S_{-i}$, $\pi_i(s_{-i})$ is the probability player i assigns to strategy profile s_{-i} occurring.*

Note something interesting here. Any specific $\sigma_{-i} \in \times_{j \in N \setminus \{i\}} \Delta S_j$ leads to a specific $\pi_i \in \Delta S_{-i}$, i.e., any specific mixed strategy profile of the other players leads to a specific probability distribution over the strategy profiles of the other players. However, the opposite is not true for games with 3 or more players, i.e., for such games there exist $\pi_i \in \Delta S_{-i}$ that cannot be generated by any $\sigma_{-i} \in \times_{j \in N \setminus \{i\}} \Delta S_j$.

For example, suppose $N = \{1, 2, 3\}$ and $S_i = \{C, NC\}$ $\forall i \in N$ (think of C as "contribute", and NC as "not contribute"). Consider π_1 where $\pi_1(C, C) = \frac{1}{2}$ and $\pi_1(NC, NC) = \frac{1}{2}$. That is, player 1 believes that the other players will both contribute with probability $\frac{1}{2}$, and will both not contribute with probability $\frac{1}{2}$. A moment's reflection should convince you that no possible mixed strategy profile of the other players can generate this probability distribution over strategy profiles of the other players (try to come up with one; you can also demonstrate this formally). But player 1 may believe that players 2 and 3 are *correlating* their strategies; they toss a coin, and if "heads" comes up they both choose C, and if "tails" comes up they both choose NC. Or, player 1 is simply uncertain about what the others are going to do, and for whatever reason assigns probability $\frac{1}{2}$ to (C, C) occurring and probability $\frac{1}{2}$ to (NC, NC) occurring. (Another solution concept, **correlated equilibrium**, allows for correlated strategies, while still maintaining the "each player is choosing a best response to what the others are doing" feature at the heart of noncooperative game theory.) Thus, beliefs are more encompassing than mixed strategy profiles; every mixed strategy profile of the other players generates a belief, but there exist (in games with 3 or more players) beliefs that cannot be generated by any mixed strategy profile of the other players.

(Osborne, p.379)

Definition 60 *Player i is **rational** if he chooses a best response (pure or mixed) to his belief $\pi_i \in \Delta S_{-i}$.*

3.16.2 Strict and Weak Dominance

(Osborne, p.120; Tadelis, p.114)

Definition 61 *(Strict dominance in mixed strategies) Player i's action $s_i' \in S_i$ is **strictly***

dominated *by its mixed strategy* $\sigma_i \in \Delta S_i$ *if* $EU_i(\sigma_i, s_{-i}) > v_i(s_i', s_{-i})$ *for all* $s_{-i} \in S_{-i}$. *We also say that* σ_i ***strictly dominates*** s_i'.

An action that is not strictly dominated by any pure strategy may be strictly dominated by a mixed strategy, and hence this is a more general notion of strict dominance. (This means that when we allow for mixing, we may be able to rule out actions from being possible choices by a rational player that we were not able to rule out when restricting ourselves to pure strategies. It also means that IESDS may get us further than when we were restricting ourselves to pure strategies, as in Tadelis's example on p.115.)

Think about why, in the above definition, we say for every $s_{-i} \in S_{-i}$ rather than for every $\sigma_{-i} \in \times_{j \in N \setminus \{i\}} \Delta S_j$.

We can modify the definition to allow for the notion of a mixed strategy to be strictly dominated:

Definition 62 *(Strict dominance in mixed strategies) Player i's mixed strategy* σ_i' *is* ***strictly dominated*** *by its mixed strategy* σ_i *if* $EU_i(\sigma_i, s_{-i}) > EU_i(\sigma_i', s_{-i})$ *for all* $s_{-i} \in S_{-i}$. *We also say that* σ_i ***strictly dominates*** σ_i'.

(Osborne, p.121)

Proposition 29 *A strictly dominated action is not used with positive probability in any mixed strategy Nash equilibrium.*

This means that when looking for MSNE, we can eliminate possibilities where a strictly dominated action is being chosen with positive probability.

(Osborne, p.121)

Definition 63 *(Weak dominance in mixed strategies) Player i's action s_i' is* ***weakly dominated*** *by its mixed strategy σ_i if $EU_i(\sigma_i, s_{-i}) \geq v_i(s_i', s_{-i})$ for all $s_{-i} \in S_{-i}$, and $EU_i(\sigma_i, s_{-i}) > v_i(s_i', s_{-i})$ for at least one specific $s_{-i} \in S_{-i}$. We also say that σ_i* ***weakly dominates*** *s_i'.*

Since weakly dominated actions (using our old definition, in pure strategies; but if an action is weakly dominated by another action, it is weakly dominated by the mixed strategy assigning probability 1 to that action) can be part of a pure-strategy NE (which is a special case of a MSNE), this means that weakly-dominated actions may be used with positive probability in a MSNE. Hence, when finding MSNE, we cannot summarily eliminate possibilities where a weakly dominated action is being chosen with positive probability.

(Osborne, p.122)

Proposition 30 *Every normal-form game in which each player has a finite number of actions has at least one Nash equilibrium (in pure or mixed strategies) in which no player's strategy is weakly dominated.*

3.16.3 Rationalizability, IESDS, and IEWDS Revisited

(Osborne, p.385; Tadelis, p.114)

Definition 64 *Player i's action $s_i' \in S_i$ is a* ***never-best response*** *if for every belief $\pi_i \in \Delta S_{-i}$ of player i, there exists a mixed strategy σ_i of player i that provides player i with a strictly higher expected utility (under π_i) than s_i'. That is, $\sum_{s_{-i} \in S_{-i}} \pi_i(s_{-i}) EU_i(\sigma_i, s_{-i}) > \sum_{s_{-i} \in S_{-i}} \pi_i(s_{-i}) v_i(s_i', s_{-i})$.*

This is a richer definition than under pure strategies, both because we are using a richer notion of beliefs, and are allowing the better response to be a mixed strategy. Note that the

σ_i in the definition can be different for each belief π_i. If there is a specific σ_i that satisfies the condition for *every* belief π_i, then obviously σ_i strictly dominates s_i'.

Recall that when we restrict ourselves to pure strategies, then a strictly dominated action is a never-best response, but a never-best response need not be strictly dominated. When we allow for mixing, the two notions are completely equivalent.

(Osborne, p.385; Tadelis, p.116)

Proposition 31 *If each player has a finite number of actions, then an action is a never-best response if and only if it is strictly dominated.*

Mixed strategies allow us to give a precise definition of rationalizability.

(Osborne, p.383)

Definition 65 *Player i's action $s_i^* \in S_i$ is **rationalizable** if for each player $j \in N$ there exists a subset $Z_j \subseteq S_j$ such that*

(i) Z_i contains s_i^*, and

(ii) for each player $j \in N$, every action $a_j \in Z_j$ is a best response to a belief $\pi_j \in \Delta Z_{-j}$ of player j (i.e., a belief that assigns positive probability only to action profiles in Z_{-j}).

This is a formidable definition, but the intuition was provided when we discussed rationalizability in the context of pure strategies. In 2-player games, player 1's strategy $s_1 \in S_1$ is rationalizable if it is a best response to some $s_2 \in S_2$, which is a best response to some $s_1' \in S_1$ (it is possible that $s_1' = s_1$), which is a best response to some $s_2' \in S_2$ (it is possible that $s_2' = s_2$), and so on *ad infinitum.* All of these elements lie in the Z_j sets mentioned in the definition, and no other elements do.

(Osborne, p.383)

Proposition 32 *Every action used with positive probability in some MSNE is rationalizable.*

We have been specifying the process of IESDS to mean that, at any given stage of the process, we eliminate *all* of the strictly dominated actions in the game that remains at that stage. But sometimes the process is instead specified as eliminating just 1 strictly dominated action at each stage (if there are any at that stage), and we can even imagine for some reason eliminating 2 but not all of them (if more) at that stage. Or 3. And so on. This raises the question of whether the actions that remain at the end of the process depend on these differences in how the process is conducted. The following result reassuringly says no. It also tells us that, because never-best responses and strictly dominated actions are equivalent, we can find the set of rationalizable actions by the relatively simple process of IESDS.

(Osborne, p.386)

Proposition 33 *If each player has a finite number of actions, then a unique set of action profiles survives IESDS, and this set is equal to the set of profiles of rationalizable actions.*

With IEWDS, the set of action profiles that survives the process *can* depend on the order in which the process is conducted (e.g., see the game on p.389 of Osborne), which is another reason why it is less compelling than IESDS.

Also note the following. With pure strategies, we said that we can't rule out a rational player choosing a weakly dominated action (that is not also strictly dominated), because it *may* be a best response to some strategy profile of the other players (but never the unique best response), i.e., to some belief. With mixing, we can conclude from Proposition 31 that a weakly dominated action that is not also strictly dominated *is* a best response to some belief; if it was a never-best response, it would be strictly dominated. This further supports the notion that we cannot rule out a rational player choosing a weakly dominated action, although again it is not clear *why* he would do so.

Osborne p.389-391 discusses additional things along these lines, as well as the notion of dominance solvability. (A game is **dominance solvable** if, when conducting IEWDS by eliminating *all* of the weakly dominated actions at each stage of the process, each player is indifferent among all action profiles that survive the process. Which means that the process "effectively" produces a unique action profile.)

3.16.4 Mixed Strategies With An Interval Set of Actions

(Tadelis, p.104)

Definition 66 *(Mixed strategies with an interval set of actions) Suppose player i has an interval set of actions $S_i = [\underline{s_i}, \overline{s_i}]$ with $\underline{s_i} < \overline{s_i}$. The set $\Delta S_i = \{F_i : S_i \to [0,1] \mid F_i$ is a cumulative distribution function$\}$[12] is the set of probability distributions over S_i, and is called player i's **set of mixed strategies**. Any $F_i \in \Delta S_i$ is a **mixed strategy** of player i, where $F_i(x) = Pr(s_i \leq x)$. If F_i is differentiable, then $f_i = F_i'$ is the density function, and we say that $s_i \in S_i$ is **in the support of** F_i if $f_i(s_i) > 0$.*

For example, in the Tragedy of the Commons Game on p.84 of Tadelis that we went over in class, a player could adopt a mixed strategy given by, say, a uniform distribution over $[0, \frac{K}{2}]$.

(Tadelis, p.106)

Definition 67 *(Expected payoff with interval action sets) Suppose $N = \{1,2\}$, $S_i = [\underline{s_i}, \overline{s_i}]$ $\forall i \in N$, and let $v_i : S \to \mathbb{R}$ be player i's Bernoulli payoff function over action profiles $s \in S$.*

[12] F_i is a cumulative distribution function if (i) $F_i(\underline{s_i}) = 0$, (ii) $F_i(\overline{s_i}) = 1$, and (iii) F_i is a weakly increasing function.

*Then the **expected payoff** of player i when he chooses the action $s_i \in S_i$ and player j chooses a mixed strategy given by the density functions f_j is*

$$EU_i(s_i, f_j) = \int_{\underline{s_j}}^{\overline{s_j}} v_i(s_i, s_j) f_j(s_j) ds_j$$

The following is the interval action set equivalent of Proposition 26.

(Osborne, p.142)

Proposition 34 *A mixed strategy profile σ^* is a MSNE if and only if, for each player $i \in N$, (i) σ_i^* assigns probability 0 to the set of actions s_i for which the strategy profile (s_i, σ_{-i}^*) yields player i an expected payoff less than her expected payoff to σ^*, and (ii) for no action s_i does the strategy profile (s_i, σ_{-i}^*) yield player i an expected payoff greater than her expected payoff to σ^*.*

3.16.5 Optional Exercises

(F1) Are the following statements true? If so, prove it. If not, provide a counter-example.

(a) A mixed strategy that assigns positive probability to a strictly dominated action is strictly dominated.

(b) A mixed strategy that assigns positive probability only to actions that are not strictly dominated is not strictly dominated.

3.17 Homework Assignment #5, Due to TA Thiago Silva on Wednesday November 2

[NOTE: Henceforth, when a problem says to find all of the NE, that means in pure as well as mixed strategies. When a problem says to find the MSNE, that also includes the pure-strategy NE, as these are special cases of MSNE. Hence, unless it specifically says to find just the pure-strategy NE, assume that it means you are to find *all* of the NE.]

(1) Tadelis 6.5

(2) Tadelis 6.7

(3) Tadelis 6.9

(4) Consider a 2-player public goods game in which each player can choose C ("contribute") or NC ("not contribute"). The public good is only produced if both choose C, and has a value of 1 for both sides. The cost of contributing for each player is $0 < c < 1$.

(a) Draw the payoff matrix. Be sure to use C as action 1 for both players, and NC as action 2.

(b) Find all of the NE in pure and mixed strategies.

(c) How do the equilibria change as c increases?

4 Dynamic Games of Complete Information

4.1 Perfect Information

Common terminology: dynamic games, sequential-move games, extensive-form games, extensive games

Game trees, **decision nodes**

(Osborne's and Osborne/Rubinstein's terms: **terminal histories**, non-terminal histories (essentially, decision nodes), **empty history**)

(Tadelis's and Fudenberg/Tirole's terms: **nodes**, **terminal nodes**, non-terminal nodes (essentially, decision nodes), precedence relation over the set of nodes, **root** (same as empty history))

Backwards induction as an intuitive way of predicting the players' choices (common knowledge of rationality implies BI)

Backwards induction leads to **subgame-perfect (Nash) equilibrium** (**SPE**)

A (pure-strategy) SPE is a profile s^* of **pure strategies**, can also talk about the **SPE outcome** and **SPE payoffs**

Player i has a set of pure strategies S_i, each $s_i \in S_i$ specifies an action for each decision node at which player i moves, can be thought of as a "complete plan of action" for the game (note that in extensive-form games, the distinction between actions and strategies is significant even without allowing for mixing, unlike in normal-form games where actions and pure strategies are equivalent)

BI when there are some indifferences in the payoffs

A **finite horizon** game is one in which the length of the longest terminal history is finite

(i.e., the game can't go on forever)

A **finite** game is one that is finite horizon *and* has finitely many terminal histories (i.e., each player has a finite number of actions at each decision node)

BI can be applied to every finite game, and hence every finite game has at least 1 SPE

NE in extensive-form games: the same definition as for normal-form games (a strategy profile with the property that no player can increase its payoff by adopting a different strategy, given the strategies of the other players)

A convenient method for finding the NE of extensive-form games: constructing the **strategic form** or **normal form** of the extensive-form

Problem with NE in dynamic games: each player is choosing a best response to what the others are doing, but may not be choosing optimally at decision nodes that are **off the equilibrium path** (i.e., not reached when the players follow their equilibrium strategies)

So NE behavior can be sustained by "threats" that are not credible under common knowledge of rationality, i.e. "non-credible threats"

BI rules out the possibility of such non-credible threats, so SPE never rely on non-credible threats, but NE may

You'll notice that every SPE is also a NE, but there may exist NE that are not SPE

So SPE is a "refinement" of NE

To formally define SPE ("a strategy profile that results from backwards induction" is not a formal definition), we need the concept of a **subgame** of a dynamic game

Let Γ be an extensive-form game. A strategy profile s^* is a SPE if for every subgame G of Γ, the restriction of s^* to G is a NE of G.

From this definition it clearly follows that an SPE is also a NE.

This is the formal definition of an SPE, and it is something to be proven that a strategy profile that results from BI is an SPE (intuitively it is fairly clear)

We want to allow the possibility of mixing, and there are two ways of thinking about this in extensive-form games

A **mixed strategy** for player i is a probability distribution over player i's set of pure strategies S_i

A **behavioral strategy** for player i specifies, for each of player i's decision nodes, an independent probability distribution over player i's set of actions at that decision node

Behavioral strategies better capture our intuitive notion of mixing in extensive-form games (mixed strategies in extensive-form games, as defined above, are kind of weird when you think about them), and it turns out that in games of **perfect recall**, which includes every game we will ever consider, they are essentially equivalent anyway, so we'll work with behavioral strategies

BI also allows us to determine SPE in mixed strategies (the usual principles with mixing apply, namely at a decision node, a player can only rationally be mixing between two or more actions if (i) it is indifferent among them, and (ii) can't strictly prefer over them an action that it is choosing with probability 0)

SPE of tic-tac-toe, chess, removing rocks game

Games with first-mover advantage (BoS or Chicken), second-mover advantage (MP), no advantage either way (PD), both players prefer a certain player to move first (1 PD, 2 SH), sequential play in SH guarantees they both go after Stag

Games involving sequential and simultaneous moves (Osborne, p.208), SPE and NE, BI

Extensive-form games with chance moves (exogenous uncertainty), Osborne p.226, SPE and NE, BI

Extensive-form games with continuum action sets, simple bargaining games

4.2 Homework Assignment #6, Due to TA Thiago Silva on Wednesday November 9

(1) Two people select a policy that affects them both by alternately vetoing policies until only one remains. First person 1 vetoes a policy. If more than one policy remains, person 2 then vetoes a policy. If more than one policy still remains, person 1 then vetoes another policy. The process continues until a single policy remains un-vetoed. Suppose there are 3 possible policies, X, Y, and Z, person 1 prefers X to Y to Z, and person 2 prefers Z to Y to X. Use the payoffs 2, 1, and 0 for each player.

(a) Draw the game-tree, and find all of the SPE in pure strategies.

(b) Draw the strategic form, and use it to find all of the NE in pure strategies.

(c) Are there any NE that are not SPE? For each one, indicate why it is not an SPE.

(2) The political figures Rosa and Ernesto have to choose either Berlin (B) or Havana (H) as the location for a party congress. They choose sequentially. A third person, Karl, determines who chooses first. Both Rosa and Ernesto care only about the actions they choose, not about who chooses first. Rosa prefers the outcome in which both she and Ernesto choose B to that in which they both choose H, and prefers this outcome to either of the ones in which she and Ernesto choose different actions; she is indifferent between these last two outcomes. Ernesto's preferences differ from Rosa's in that the roles of B and H are reversed. Karl's preferences are the same as Ernesto's. Use the payoffs 2, 1, and 0 for each player.

(a) Draw the game-tree, and find all of the SPE in pure strategies.

(b) Draw the strategic form, and use it to find all of the NE in pure strategies.

(c) Are there any NE that are not SPE? For each one, indicate why it is not an SPE.

4.3 Allowing for Imperfect Information, and Formal Definitions and Results Regarding Extensive-Form Games

(Tadelis, p.133; Fudenberg and Tirole, p.78)

Definition 68 *A* ***game tree*** *is a finite set of nodes X, with a precedence relation on X denoted by $>$: for any $x, x' \in X$, $x > x'$ means that "x precedes x'." The precedence relation is* ***transitive*** *($x > x', x' > x'' \Rightarrow x > x''$),* ***asymmetric*** *($x > x' \Rightarrow$ not $x' > x$), and* ***incomplete*** *(not every pair of nodes x, x' can be ordered). There is a special node called the* ***root*** *of the tree, denoted by x_0, that precedes every other node. Every node x (other than x_0) has exactly one immediate predecessor (i.e., one node $x' > x$ such that $x'' > x, x'' \neq x' \Rightarrow x'' > x'$). Nodes that do not precede other nodes are called* ***terminal nodes****, denoted by the set $Z \subset X$. Terminal nodes denote the final outcomes of the game with which payoffs are associated. Every node x that is not a terminal node is assigned either to Nature, or to a player using the* ***player function*** *$i : X \setminus Z \to N$. At non-terminal node x, player $i(x)$ has the action set $A_{i(x)}(x)$.*

The game-tree captures the physical aspects of a sequential interaction, but not what the players know when they move.

(Tadelis, p.135)

Definition 69 *Every player i has a set of* ***information sets*** *H_i. Each $h_i \in H_i$ is an information set of player i. H_i partitions the nodes of the game at which player i moves with the following properties:*

(1) If h_i is a singleton that includes only x, then player i who moves at x knows that he is at x.

(2) If $x \neq x'$ and if both $x \in h_i$ and $x' \in h_i$, then player i who moves at x (and x') does not know whether he is at x or x'.

(3) If $x \neq x'$ and if both $x \in h_i$ and $x' \in h_i$, then $A_i(x) = A_i(x')$. Therefore, sometimes we refer to the action set as $A_i(h_i)$.

(Tadelis, p.136)

Definition 70 *A game of complete information in which every information set is a singleton and there are no moves by Nature is called a* ***game of perfect information****. A game in which some information sets contain several nodes or in which there are moves by Nature is called a* ***game of imperfect information****.*

(Tadelis, p.139)

Definition 71 *A* ***pure strategy*** *for player i is a function $s_i : H_i \to A_i$ that assigns an action $s_i(h_i) \in A_i(h_i)$ for every information set $h_i \in H_i$. We denoted by S_i the set of all pure strategies for player i. (Note that $A_i = \bigcup_{h_i \in H_i} A_i(h_i)$ is the set of all actions for player i, i.e., the set of all actions at all of that player's information sets.)*

Proposition 35 *The number of pure strategies for player i is given by $|S_i| = \prod_{h_i \in H_i} |A_i(h_i)|$.*

Definition 72 *A* ***mixed strategy*** *for player i is a probability distribution over his set of pure strategies S_i.*

(Tadelis, p.140)

Definition 73 *A* ***behavioral strategy*** *for player i specifies for each information set $h_i \in H_i$ an independent probability distribution over $A_i(h_i)$. Formally, it is a function $\sigma_i : H_i \to \Delta A_i(h_i)$,[13] where $\sigma_i(a_i(h_i))$ is the probability that player i chooses action $a_i(h_i) \in A_i(h_i)$ at*

[13]Technically, this should actually be $\sigma_i : H_i \to \bigcup_{h_i \in H_i} \Delta A_i(h_i)$.

information set h_i.

(Tadelis, p.142)

Definition 74 *A game of **perfect recall** is one in which no player ever forgets information that he previously knew.*

(Tadelis, p.146)

Definition 75 *Let $\sigma^* = \{\sigma_1^*, \ldots, \sigma_n^*\}$ be a Nash equilibrium profile of behavioral strategies in an extensive-form game. We say that an information set is **on the equilibrium path** if given σ^* it is reached with positive probability. We say that an information set is **off the equilibrium path** if given σ^* it is never reached.*

(Tadelis, p.152)

Definition 76 *Given a profile of behavioral strategies σ_{-i} of the other players, we say that player i's behavioral strategy σ_i is **sequentially rational** if it has player i choosing a best response to σ_{-i} at each of his information sets.*

(Tadelis, p.153)

Proposition 36 *Any finite game of perfect information has a backwards induction solution that is sequentially rational for each player. If no two terminal nodes prescribe the same payoffs to any player then the backward induction solution is unique.*

Proposition 37 *Any finite game of perfect information has at least one sequentially rational Nash equilibrium in pure strategies. If no two terminal nodes prescribe the same payoffs to any player then the game has a unique sequentially rational Nash equilibrium.*

(Tadelis, p.154)

Definition 77 *A* ***subgame*** *G of an extensive-form game Γ consists of a single node and all its successors in Γ with the property that if $x \in G$ and $x' \in h(x)$, then $x' \in G$. The subgame G is itself a game tree with its information sets and payoffs inherited from Γ. A subgame G that does not contain the root x_0 of Γ is called a* ***proper subgame*** *of Γ.*

(Tadelis, p.157)

Definition 78 *Let Γ be an extensive-form game. A behavioral strategy profile $\sigma^* = \{\sigma_1^*, \ldots, \sigma_n^*\}$ is a* ***subgame-perfect (Nash) equilibrium (SPE)*** *if for every subgame G of Γ the restriction of σ^* to G is a Nash equilibrium in G.*

Although backwards induction, when applicable, leads to an SPE (this can formally be stated as a proposition), SPE is more general because it can even be applied to games of imperfect information, which BI cannot.

OSBORNE'S ALTERNATIVE DEFINITIONS

(Osborne, p.155)

Definition 79 *An* ***extensive-form game with perfect information*** *consists of:*

(1) a set of ***players***

*(2) a set of sequences (****terminal histories****) with the property that no sequence is a proper subhistory of any other sequence (by sequence, we mean sequence of actions)*

(3) a function (the ***player function****) that assigns a player to every sequence that is a proper subhistory of some terminal history*

*(4) for each player, **preferences** over the set of terminal histories.*

(Osborne, p.157)

Definition 80 *If the length of the longest terminal history is finite, then we say that the game has a **finite horizon**. If the game has a finite horizon and a finite number of terminal histories, then we say that the game is **finite**.*

(Osborne, p.159)

Definition 81 *A **strategy** of player i in an extensive-form game with perfect information is a function that assigns to each history h after which it is player i's turn to move an action in $A(h)$ (the set of actions available after h).*

4.4 Homework Assignment #7, Due to TA Thiago Silva on Wednesday November 16

(1) Consider the game-tree in Figure 7.14 on p.150 of Tadelis. Change the labels of the players so that it is player 1 moving first (choosing between a and b), then player 2 moves, and then player 1. (That is, just reverse all player labels; everything else remains the same.)

(a) Draw the normal form of this game, and find all of the NE in pure strategies.

(b) Find all of the SPE in pure strategies. (Hint: Tadelis p.158 provides an example.)

(2) There is a kid K who has to decide how large a cake to bake. The largest possible cake is of size A, where $A > 0$. Let the size of the cake that K chooses to bake be x, where $0 \leq x \leq A$. Suppose that the cost to K of baking a cake of size x is ax, where $a > 0$ is a positive constant that is the marginal cost of baking more cake. K's payoff for building a cake of size x is the size of the cake minus the cost.

(a) How large a cake will K bake, i.e., what x will it choose? (Hint: There will be a critical threshold for a that will be relevant for your answer.)

Now suppose that there is an additional actor, the parent P. There is an additional kid K_2, who is not an actor in this strategic interaction. Once K has chosen x, P chooses some y where $0 \leq y \leq x$, which is the share of the cake that goes to K_2. K's payoff is the remaining share minus the cost of producing the cake. Because a parent is only as happy as its least happy child, P's payoff is the smaller of the two shares that go to the kids, y and $x - y$. That is, P's payoff is $\min\{y, x - y\}$. (The parent does not worry about the cost to K of baking the cake.)

(b) Draw the game-tree, and find the subgame-perfect equilibria (SPE). Be sure to clearly indicate the size of the cake that will be baked in the SPE.

5 Dynamic Games of Incomplete Information

Dynamic games of imperfect information are those in which there is at least one information set that contains more than 1 decision node

That is, at least 1 player does not perfectly observe at least one move in the game

We can construct the strategic form (or normal form) and use it to find all of the NE

We can also find all of the SPE, either using backwards induction or the direct method (find all of the NE, and then for each of these, determine whether its restriction to each subgame is a NE in that subgame)

SPE is a very nice solution concept for sequential-move games of perfect information

Give an example of a sequential-move game of imperfect information in which SPE leads to non-reasonable equilibria (Gibbons, p.176)

So we will develop a new solution concept, called perfect Bayesian equilbrium (PBE), that is well-suited for sequential-move games of imperfect information, and that has 3 requirements

(Gibbons, p.177)

Requirement 1 *At each information set, the player who moves has a* ***belief*** *about which node in the information set has been reached by the play of the game. For a non-singleton (non-degenerate) information set, a belief is a probability distribution over the nodes in the information set; for a singleton (degenerate) information set, the player's belief puts probability 1 on the single decision node.*

Requirement 2 *Given their beliefs, the players' strategies must be* ***sequentially rational****. That is, at each information set the action taken by the player with the move (and the player's* *subsequent strategy) must be optimal given the player's belief at that information set and the other players' subsequent strategies (where a "subsequent strategy" is a complete plan of action covering every contingency that might arise after the given information set has been reached).*

Definition 82 *For a given equilibrium in a given extensive-form game, an information set is* ***on the equilibrium path*** *if it will be reached with positive probability if the game is played according to the equilibrium strategies, and is* ***off the equilibrium path*** *if it is certain not to be reached if the game is played according to the equilibrium strategies (where "equilibrium" can mean NE, SPE, BNE, or PBE).*

Requirement 3 *At information sets on the equilibrium path, beliefs are determined by* ***Bayes' rule*** *and the players' equilibrium strategies.*

Definition 83 *A* ***perfect Bayesian equilibrium (PBE)*** *consists of strategies and beliefs satisfying Requirements 1 through 3.*

Methods for finding PBE.

My backwards induction-type method, that can be applied to any finite extensive-form game. Label each information set, assign belief labels to each non-degenerate information set, then start doing BI, at non-degenerate information sets this will involve expected utility calculations using the belief labels

Give examples, along with how to apply Bayes' rule.

For information sets that are on the equilibrium path, we refer to the beliefs as **on-the-equilibrium-path beliefs**, and these are the ones to which Requirement 3 applies.

For information sets that are off the equilibrium path, we refer to the beliefs as **off-the-equilibrium-path beliefs**, and Requirement 3 does not apply to these beliefs. PBE does not impose restrictions on what these beliefs must be. Refinements of PBE do impose restrictions on these beliefs.

An alternative, less-preferred method is based on the following result:

(Osborne, p.329; Osborne calls PBE **weak sequential equilibrium**)

Proposition 38 *The strategy profile in any PBE is a Nash equilibrium.*

(It need not be an SPE; for an example, see the game in Gibbons, p.181.)

Thus, an alternative method is to find all of the NE, and then for each one determine what the beliefs must be at information sets that are on the equilibrium path (by Requirement 3). Then determine whether sequential rationality is satisfied at each information set. This is a tedious method (especially when moves by "nature" are involved) that is less preferred than the BI-type method, which is very systematic.

Games of incomplete information: at least one player is uncertain about the other player's payoffs, and the game begins with **nature** or **chance** probabilistically choosing the latter player's (the **informed player's**) **type** (each type is associated with one set of possible payoffs for that player), a move that the former player (the **uninformed player**) does not observe (but knows the probabilities), and hence a game of incomplete information is analyzed as a game of imperfect information

Signaling game: a game of incomplete information transformed into a game of imperfect information as above, and in which the informed player moves before the uninformed player and the informed player's actions are interpreted as **messages**, which may or may not affect payoffs (costly messages versus costless messages)

Sometimes the informed player is called the **sender** (of the message), and the uninformed player is called the **receiver**, and sometimes these are called **sender-receiver games**

Pooling equilibria: all types of the informed actor choose the same message (hence, no information is conveyed by the message)

Separating equilibria: all types of the informed actor choose different messages (hence, full information is conveyed by the message)

Semi-separating equilibria: the types are not fully pooling, but are also not fully separating, so that partial information is conveyed by the message (some updating occurs upon a message being received; usually these equilibria are in mixed strategies)

With costly messages, the general result is that separating equilibria only exist if there exists a message that is not too costly for the "strong" type to send, but is sufficiently costly to the "weak" type as to deter it from mimicking the strong type

With costless messages (sometimes called **cheap talk games**), the general result is that separating equilibria only exist if the sender and received have sufficiently aligned preferences

Some PBE have (and are sustained by) off-the-equilibrium-path beliefs that are not reasonable, and hence additional requirements have been developed to refine the set of PBE.

Requirement 4 *(PBE Refinement #1) At information sets off the equilibrium path, beliefs are determined by Bayes' rule and the players' equilibrium strategies where possible.*

See the game on Gibbons, p.181. In fact, Gibbons includes Requirement 4 in the definition of PBE, but most authors don't.

Requirement 5 *(PBE Refinement #2) If possible, each player's beliefs off the equilibrium path should place zero probability on nodes that are reached only if another player plays a* *strategy that is strictly dominated beginning at some information set.*

See the game on Gibbons, p.233.

Other, more complicated refinements exist as well, such as the blandly-named "intuitive criterion" or the more exotic "universal divinity", and that apply primarily to signaling games. The two simple refinements above give a sense of what refinements involve: narrowing down the set of off-the-equilibrium-path beliefs that are acceptable, by asking the question: "if an off-the-equilibrium-path information set is reached, what is it reasonable to suppose that the player there will believe about what decision node I may or may not be at?"

5.1 Homework Assignment #8, Due to TA Thiago Silva on Wednesday November 30

(1) For the following 2 extensive-form games, draw the normal form and find all of the pure-strategy NE and pure-strategy SPE. Then find all of the pure-strategy PBE (using the backwards-induction-type method).

(2) Find all of the pure-strategy PBE of the following game (using the backwards-induction-type method).

5.2 Problems We Will Solve in Class on 12/2

(1) Consider the following extensive-form games of imperfect information. There are two possible states of the world, L and R. Nature chooses the state of the world to be L with probability $\frac{2}{3} < p < 1$, and R with probability $1 - p$.

Congress can pass bill l, which is optimal policy if the state of the world is L, or bill r, which is optimal policy if the state of the world is R. Following Congress's choice, the President decides whether to Sign or Veto the bill. If he vetoes it, the status-quo remains in place.

Payoffs are as follows. The President is a public servant of great moral rectitude, and only cares about the optimality of the policy that is passed. Hence, he gets a payoff of 3 if the policy passed matches the state of the world, a payoff of 1 if the status quo remains in place, and a payoff of 0 if the policy passed is the opposite of the state of the world.

Congress is biased towards policy r, but also cares somewhat that appropriate policy be passed. If the state of the world is R, then Congress gets a payoff of 3 if policy r is passed, 1 if the status quo remains in place, and 0 if policy l is passed. If the state of the world is L, then Congress gets a payoff of 3 if the status quo remains in place, 1 if policy r is passed, and 0 if policy l is passed.

(a) First suppose that neither actor observes the state of the world, but knows the above probabilities. Draw the game-tree, and find all of the perfect Bayesian equilibria (PBE) in pure strategies.

(b) Now suppose that the President observes (knows) the state of the world. Draw the game-tree, and find all of the PBE in pure strategies.

(c) Now suppose that the President observes the state of the world and sends a costless message (L' or R') to Congress, and then Congress decides which bill to pass, followed by the President deciding whether to Sign or Veto the bill. Draw the game-tree, and find all of the PBE in pure strategies. For each PBE, state whether it is pooling or separating.

(d) Are the PBE outcomes different in parts (a), (b), and (c) in terms of the likelihood of the status quo remaining in place? Explain. Does the President knowing the state of the world help or hurt Congress?

(2) Consider the following **static game of incomplete information**. There are 2 players, who simultaneously choose whether to Cooperate (C) or Defect (D). Player 1 has Stag Hunt (SH) preferences, and there are two "types" of player 2. One type has SH preferences, and the other has Prisoner's Dilemma (PD) preferences. Nature chooses player 2 to have SH preferences with probability $0 < p < 1$, and PD preferences with probability $1-p$. Player 2 knows its own type (i.e., observes nature's move), whereas player 1 only knows the probabilities with which nature chose. Draw this game, and find all of the **Bayesian Nash equilibria (BNE)** in pure strategies.