

Societal-Scale Cyber-Physical Systems
The Science of Digital Transformation

Lillian Ratliff, Shankar Sastry,
with contributions from our co-authors, especially Roy Dong and Eric Mazumdar

January 7, 2020

Contents

1	Introduction	5
1.1	Introduction	6
1.1.1	What is the emerging data market?	6
1.1.2	Informational Content of Data	6
1.1.3	Incentive Structure of the Data Market	6
1.2	Running examples	7
1.2.1	Example 1: Curbside Management	7
1.2.2	Example 2: Resource Allocation and Goods Distribution	8
1.2.3	Example 3: Labor Markets	9
2	Game Theory	11
2.1	Preliminaries	12
2.2	Games in Strategic Form and Nash Equilibria	13
2.2.1	Equilibrium Notion	15
2.2.2	Best Response	17
2.2.3	Mixed Strategies	18
2.2.4	Strict Dominance of Mixed Strategies	21
2.3	Existence of Nash Equilibria	22
2.4	Computation of Nash Equilibria in Finite Games	24
2.4.1	Zero Sum Games	26
2.4.2	Non-Zero Sum Games	27
2.5	Correlated Equilibria in Strategic Form Games	29
2.5.1	Computation of Correlated Equilibria	32
2.6	Stackelberg Equilibria in Strategic Form Games	32
2.7	Games with Continuous Actions Spaces	32
2.7.1	Best Response	33
2.7.2	Existence of Nash Continuous games	33
2.8	Optimality Conditions to Characterize Local Nash Equilibria	34
2.8.1	Computation of Nash in Continuous Games	38
2.8.2	Continuous Stackelberg Games	39
2.9	Differential and Dynamic Games	39
3	Theory of Contracts and Incentives	41
3.1	Social Choice Rules	42
3.2	Mechanism Design	43
3.2.1	The Revelation Principle and Truthful Implementation	43
3.3	Incentive Design	45

3.3.1	The Principal-Agent Problem	45
3.3.2	Information Asymmetries	46

Chapter 1

Introduction

1.1 Introduction

1.1.1 What is the emerging data market?

1.1.2 Informational Content of Data

1.1.3 Incentive Structure of the Data Market

The Internet of Things (IoT) is a term that represents a huge technological trend that is taking place: almost every device is being imbued with the intelligence of a microprocessor and an Internet connection. The interconnection in IoT promises an infrastructure that can drastically change how consumers live their day-to-day lives, with huge gains in efficiency, value, and possibility due to the shared knowledge and autonomy allowed. In profound ways, as the technology develops, the modalities of existence people experience will grow and shift.

However, the scale and scope of IoT raises new problems for engineers to consider. These problems are significantly different from ones previously explored in the design of comparatively isolated systems, and require a new theoretical underpinning to analyze IoT with models that capture all salient facets of these new technologies. This document contains a handful of theoretical frameworks, and their applications, as a first step into this new research frontier.

First, we consider the problem of large amounts of data. For example, in the energy sector, advanced metering infrastructures collect energy consumption data for a large number of consumers at relatively high frequencies. This glut of data isn't useful for most operational purposes, such as phase-alignment, and is often aggregated for control purposes. Additionally, these smart meter readings are usually at a household level, and themselves represent an aggregate of several devices inside an energy consumer's home.

Furthermore, if these devices are thought of as dynamical systems, these smart meter readings only capture the 'output' of the system: both the internal state dynamics and the driving inputs remain unobserved. This generally is a trend with IoT sensors: they capture a process but not the 'inputs' driving it. As another example, smart phone data tracks the location of users, which can be thought of as a random process whose distribution depends on the user's itinerary, which is often not a direct 'input' into Google Maps.

In Chapter ??, we address these two problems.

We define the disaggregation problem as the recovery of component signals y_i from observations of an aggregate $\sum_i y_i$. We focus on an application in energy disaggregation, and outline assumptions motivated by this context. Under these assumptions, we can phrase the disaggregation problem as a hybrid optimal control problem. Using adaptive filtering techniques, we provide an algorithm that can tractably find the optimal solutions to the disaggregation problem.

Additionally, we provide a blind system identification method to simultaneously identify the inputs and dynamics of systems when only output observations are available. This, in and of itself, is an ill-posed problem, so we find prior knowledge that can be supported by IoT sensors. For example, if we are tracking occupancy in a building, IoT sensors can detect when a door is opened. This serves as regularizing knowledge for blind system identification: we know the discrete time points when occupancy of a sector can change.

Chapter ?? outlines some new estimation problems due to the scale and nature of IoT sensor measurements.

A reader who is sensitive about the abundance of collected and transmitted technology may be unnerved by some of the results in Chapter ??, and rightfully so. In Chapter ??, we outline some of the work in preserving privacy in new IoT systems. Fundamentally, we suppose IoT sensors are collecting data for some estimation or control tasks that are not directly in line with privacy

violations of the user. In other words, privacy risks are a byproduct of a system designed to do something else.

In Chapter ??, we outline methods to quantify the privacy of a system. We review the literature on quantitative measures for privacy, and discuss some of our contributions, predominantly in *inferential privacy*. We discuss different design paradigms for the incorporation of privacy into IoT systems as a design level: from a passive privacy audit of a system, to an optimal design perspective that creates sensors that provide measurements that optimally benefit the control objective while minimizing privacy threats to users. Finally, we implement these privacy design concepts with examples in ground transportation, smart grid control, and air quality regulation.

As IoT technologies permeate our socioeconomic spheres more and more, users will become more privacy-aware and conscious of their data flows. A common saying heard these days in tech circles is: ‘If you are not paying for a service, you the product, not the customer.’ This applies to services such as Gmail or Facebook. Individual users are increasingly aware that their data is, in many ways, the currency of new technologies. Web services such as PrivacyFix allow users to calculate how much revenue is generated by Google from their data alone, and change their privacy settings.

In Chapter ??, we model users who are strategic about their data sharing, and how the actions of data buyers and advertisers must change when users become more strategic. The existing literature has some methods that can handle strategic data sources, and we extend these results to consider how these methods work when multiple data buyers are competing to create better estimators.

Additionally, IoT allows new means and modes of providing feedback into our systems. This can come in the form of economic rebates for products that change the energy consumption profiles inside a house, or warning lights inside a vehicle outfitted with intelligent sensors, or cell-phone messages suggesting faster routes. None of these are direct control actions in a control-theoretic sense, but can be thought of as actions that have an effect on the distribution of system behaviors after the fact.

Furthermore, these are often imputations on endogenous variables inside a system. For example, the price of an eco-friendly fridge is determined by supply and demand in the market, but a rebate can, in a sense, manually force the market price to something else, at some cost to the entity issuing the rebates. We argue that this phenomena can be modeled as *causal imputations*, and discuss a framework for finding the optimal imputations under some system performance objective and a cost of imputation. This is also discussed in Chapter ??.

1.2 Running examples

1.2.1 Example 1: Curbside Management

Curbside usage includes all of the follow, and more:

- floating micro-transit vehicles (bikes, scooters, etc.)
- flex zones (load zones, passenger pick-up zones, ride-share zones, etc.)
- curbside parking for personal vehicles
- floating car shares
- bus zones

Even as our transportation systems evolves to accommodate autonomous vehicles, there is always going to be a need for curbside management.

Curbside management, and on-street parking in particular, is surprisingly a huge bottleneck for cities. It not only impacts congestion in dense cities, but also business district vitality, residential access and usage, and even provisioning of construction permits (cities often require ample parking in order to build new buildings).

Curbside management includes the following:

1. Active learning for robustly inferring curbside demand
2. Algorithmic mechanism design (co-design of information and incentives/pricing)
3. game theoretic models of user behavior
4. dynamical systems (queuing) models of system level behavior (e.g., networks of queues with rejections transitioning from "server" to "server", sort of dual the classical model for queuing networks)
5. behavioral models of users choosing where to park based on various perceptions about things like wait time, likelihood of getting a ticket, price, etc.; includes understanding the value of information and can tie into algorithmic mechanism design
6. competition between firms: mobile payment platforms, auctions for private curb space usage (happens in places like boston e.g.)

1.2.2 Example 2: Resource Allocation and Goods Distribution

E-commerce is increasingly becoming the norm. Even classically brick and mortar businesses such as Walmart have increased their online presence in order to compete with powerhouses such as Amazon. This is a global phenomena as well.

Essentially, e-commerce platforms allow users to shop for items online and have them delivered to their homes. The mechanisms by which goods are delivered are varied and include, e.g., the postal service, distribution companies such as DHL, FedEx and UPS, the e-commerce company themselves, and flex drivers (individuals who contract with the e-commerce company or other third parties to deliver goods for money).

e-commerce includes the following interesting features:

1. Active learning for robustly inferring users preferences (ranking and matching algorithms; experimental design)
2. Algorithmic mechanism design (design of incentives, auctions, and other mechanisms)
3. game theoretic models of user behavior, both online users as well as delivery persons
4. dynamical systems models of system level behavior (e.g., inventory models, etc.)
5. behavioral models of users choosing which items to buy or when/where to log-on to participate in flex driving market
6. competition between firms

1.2.3 Example 3: Labor Markets

Labor markets have seen a revolution over the last decade. We have seen, e.g., the emergence of crowdsourcing platforms such as Amazon's Mechanical Turk (MTurk) and the rise of contract based employees for services such as food delivery (Uber Eats), house cleaning and repair (task rabbit), goods delivery (Amazon Flex) and, the most prominent of them all, ride-sharing (Uber, Lyft, Didi, etc.). Labor markets for short term contracts present interesting dynamics and a plethora of unintended consequences and challenges.

e-commerce includes the following interesting features:

1. Active learning (matching, ranking, etc.)
2. Algorithmic mechanism design (design of incentives, auctions, micro-markets, etc.)
3. Multi-sided markets
4. game theoretic models of user behavior in interesting settings such as multi-sided markets in which both adverse selection and moral hazard issues present themselves due to information asymmetries such as laborer's skill and effort
5. dynamical systems models of system level behavior that include both market dynamics which tend to operate in the cloud and physical system dynamics which operate on different timescales (e.g., the time it takes to physically drive from one location to another to for instance deliver a good or pick up a passenger is much much longer than the time it takes to resolve a match between a driver and passenger)
6. behavioral models of users
7. competition between firms (this example is prime for discussing and modeling competition between firms)

Chapter 2

Game Theory

The purpose of this chapter is to provide an overview of the game theoretic tools that serve as the underpinning for many of the concepts presented in later chapters. In particular, we aim to familiarize the reader with game theoretic abstractions in order to show that many design challenges in societal-scale systems can be analyzed through a game theoretic lens and further, aims to illuminate the essence of problems in this domain as they are conceptualized. For instance, the reader will learn through examples how to analyze a given problem and identify who the stakeholders are and what each stakeholder's objective is, the nature of the interactions between stakeholders, and learn to conceptualize what aspects of such interactions are subject to design.

Broadly, game theory provides a rigorous framework for reasoning about strategic decision making in both competitive and cooperative environments. In particular, it is the study of multi-agent decision problems and it is used in a variety of fields including economics, political science, biology, engineering, etc. It is used, for instance, to understand competition or cooperation among agents, understand *carrot versus stick* policies and their impacts in the long run, and to study adversarial behavior (security attacks, terrorism, warfare), among many other applications. Sometimes the game arises naturally and is the object of study (e.g., competition in a marketplace). Other times games are a natural abstraction which may lead to insights about some underlying phenomena (e.g., evolutionary dynamics). Finally, game theoretic abstractions may also be leveraged to improve the design of a decision making or machine learning algorithm (e.g., adversarial learning).

Examples in Science and Engineering:

- Network economics: games are used to model and understand policies for net neutrality.
- Adversarial learning: game theoretic abstractions are being leveraged to construct a framework for learning algorithms or decision rules which are robust to adversarial input. For example, in learning generative adversarial networks, it is common practice to abstract the interaction between discriminator and generator as a zero-sum game thereby leading the generator to learn a robust decision-making rule.
- Energy markets: game theory is often used to determine the best strategy for bidding in energy markets as well as for the design of tariffs or demand response incentives.
- Crowdsourcing markets: mechanism design and game theory are used to design market mechanisms for facilitating crowdsourcing; e.g., in the study of incentivization of truthful participation while protecting privacy.
- Human-machine interaction: game theoretic models are quite often leveraged in modeling the interaction between two boundedly rational agents who form beliefs about the intentions of other agents with whom they are interacting.

- Modern labor markets: similar to crowdsourcing markets, modern labor markets include numerous examples of ‘crowdsourced’ contract-based laborers. Examples of such markets include Amazon’s Mechanical Turk, flex drivers which deliver goods including food, and the numerous vehicle-sharing platforms whose labor pool consists of drivers with extra capacity in their vehicles. Such markets can be abstracted as multi-sided (generally, two-sided) markets and the interactions between the different sides of the market are subject to design. Further, competition between platforms (represented by firms/companies) is easily abstracted to a game theoretic framework, which is the subject of many studies of these emerging markets.
- Biology: evolutionary dynamics of populations are often abstracted as games.
- Micro-financing: recently trending are micro-lending or crowdfunding markets. Game theoretic abstractions of such markets support the design of everything from regulatory policies to adaptive mechanisms for matching lenders to borrowers.




Figure 2.1: Abstraction of a data crowdsourcing market.

The chapter is divided into two main threads: one on equilibrium concepts—e.g., Nash and its refinements—and their characterization and one on learning algorithms employed by agents or players seeking such equilibria. We start with some preliminary formalism that will be used throughout the remainder of the chapter.

2.1 Preliminaries

The term player refers to an agent which facing a decision making problem. Often the term player and the term agent are used interchangeably. Loosely speaking, a *game* consists of a set of players, a specification for the environment in which they interact, the objective structure for each player, and the information structure for each player¹. The environment describes the broader context in which the agents are interacting; it may, e.g., include specifications for “nature” or some auxiliary state which influences the decisions of agents through, e.g., their objectives. The objective structure includes specifications on the individual objectives of players and their action spaces. The information structure includes specifications on what information is available to players when they make decisions; e.g.,

- Chess is a *full-information* game because the current state of the game is fully known to both players as they make their decisions.

¹Aside: we remark that in mechanism design (a subject of later chapters), either the environment and objective structure are considered fixed and the information structure is subject to design (information design problem), or the environment and information structure are considered fixed and the objective structure is subject to design (incentive design problem).

- In contrast, Poker is a *partial-information game* since players do not observe the hand dealt to their opponents.

In general, there are two *information settings*: full/complete and partial information. At a high level, in the former setting, knowledge about other players is available to all players where in the latter setting, the knowledge set is incomplete, meaning some information is not available to one or more of the players or is revealed overtime as the players interact. We will further characterize these two settings once we have introduced more of the mathematical formalism.

For a mathematical solution to a game, one further needs to make assumptions on the player's *rationality*. In the classical game theoretic setting, it is commonly assumed that the players are *perfectly rational* decision-makers, seeking to maximize their utility. However, it is well known that humans are not perfectly rational and even when an agent is perfectly rational, it may be limited by its computational ability in which case models of bounded rationality or myopic decision-making become relevant. Such models have been introduced in behavioral psychology and economics, as well as game theoretic learning approaches. Discussions on behavioral models are relegated to subsequent sections in this chapter, and for the time being, we will assume that players are fully rational decision makers.

Until otherwise stated, we assume players are rational in the sense that they are aware of their alternatives, form expectations about any unknowns, have clear preferences, and choose their action deliberately after some process of optimization.

In the remaining parts of this chapter, we will cover the formalism for strategic form games on finite action spaces, games on continuous action spaces, and dynamic/differential games. For each of these classes of games, we will provide formal definitions for different equilibrium concepts.

2.2 Games in Strategic Form and Nash Equilibria

Generally speaking, there are two categories of games: strategic and extensive form.

- **Strategic form game:** a model of interactive decision-making in which each player chooses his plan of action once and for all, and all players decisions are made simultaneously (that is, when choosing a plan of action each player is not informed of the plan of action chosen by any other player).
- **Extensive form game:** specifies the possible orders of events; each player can consider his plan of action not only at the beginning of the game but also whenever he has to make a decision.

A strategic form (normal form) game on finite action spaces is a game that can be described using a matrix or table. Typically a strategic form is used for simultaneous play games on finite action spaces. This is in contrast to sequential games in which information is revealed to players in stages which are more naturally modeled in *extensive form*² (or tree form).

Let us review the basic constructs necessary for defining strategic form games.

Definition 1. (*Strategic Form Game.*) A strategic form game (*normal form or matrix game*) is a tuple $\mathcal{G} = ((U_i)_{i \in \mathcal{I}}, (J_i)_{i \in \mathcal{I}})$ where

- The set $\mathcal{I} = \{1, \dots, n\}$ indexes a finite set of players;
- For each $i \in \mathcal{I}$, U_i represents the available actions (*action space*) to player i and $U = U_1 \times \dots \times U_n$ is the joint action space.

²While this is an important class of games, for brevity we will not review this game type and refer the reader to more in depth game theory texts.

- Each $u_i \in U_i$ is an action for player i and $u = (u_1, \dots, u_n) \in U$ is a joint action.
- Indexed by $i \in \mathcal{I}$, the function $J_i : U \rightarrow \mathbb{R}$ is utility function for player i .

We denote by $-i = \mathcal{I} \setminus \{i\} = \{1, \dots, i-1, i+1, \dots, n\}$ all players excluding the i -th player. With this notation, $u_{-i} \in U_{-i} = \times_{j \neq i} U_j$ with $u_{-i} = (u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_n)$ denotes the joint action profile of all players excluding player i .

Strategies (or policies) are *complete contingent plans* that specify the action to be taken in every situation—as encoded by an *information set*—that could arise in the course of play. Strategies represent an important concept in games in particular as they contrast with actions:

- An *action* is a possible move that is available to a player during a game.
- A *strategy* is a decision rule that a player uses to select actions, based on available information.

Remark 1. A strategy is a complete description of how to play the game (i.e., you could give it to a computer like an algorithm for it to execute it). An action on the other hand is the result of the strategy. In strategic form games, strategy and action are synonymous.

Another important concept in games is *rationality* and the assumed notion of *common knowledge*.

Vignette 2.1 (Common Knowledge.) Think of each player as being in a separate room where they are asked to choose a button, perhaps on a computer terminal, without communicating with other players. Players know the strategic form \mathcal{G} , and know their opponents know it, and know their opponents know they know it, and so on ad infinitum. That is, \mathcal{G} is **common knowledge**.

Later we will introduce the concept of *bounded rationality* where there is a bounded number of layers in this knowledge structure.

Agents' utility functions capture the *preferences* of the players in the following sense.

Definition 2. (*Preference Relation.*) For each player $i \in [n]$, a preference relation \succsim_i on U is such that $u \succsim_i u'$ if and only if $J_i(u) \geq J_i(u')$ where $u, u' \in U$.

Remark 2. As an alternative definition of strategic form games, we can use the preference relation notation as follows: $(\mathcal{I}, (U_i)_{i \in [n]}, (\succsim_i)_{i \in [n]})$.

If each U_i is finite, we say the game is a *finite game*. When the number of players and strategies is small, representing the game in matrix form is useful. We will represent the payoffs as pairs (a, b) where $a = J_1(u, v)$, $b = J_2(u, v)$ and $u \in U_1$, $v \in U_2$. Given the mapping between joint actions and payoffs, we can represent all the outcomes of the game in tabular or matrix form:

		P2	
		L	R
P1	T	(a_{11}, b_{11})	(a_{12}, b_{12})
	B	(a_{21}, b_{21})	(a_{22}, b_{22})

Thus, the table describes all possible outcomes; e.g., $J_1(T, L) = a_{11}$, $J_2(T, L) = b_{11} \dots$

Example 2.1 (Matching Pennies.) There are two players, each of which has a penny. The game is played as follows. The players flip the pennies and if the pennies are both heads or both tails then player 1 (even) keeps both pennies. On the other hand, if the pennies are one heads and one tails, then player 2 (odd) keeps the pennies. Thus, the outcomes for this game can be described by the following table or matrix:

		P2 (odd)	
		heads	tails
P1 (even)	heads	(+1,-1)	(-1,+1)
	tails	(-1,+1)	(+1,-1)

The matching pennies game is a *zero sum* game and represents pure conflict (since the utility of one player is the negative utility of the other player).

Definition 3. A two-player game is called *zero sum* if

$$J_1(u, v) + J_2(u, v) = 0, \quad \forall (u, v) \in U_1 \times U_2,$$

and otherwise, we say the game is *non-zero or general sum*.

2.2.1 Equilibrium Notion

When there are n decision-makers each with their own objective $J_i : U \rightarrow \mathbb{R}$ for $i \in \mathcal{I}$ where $U = U_1 \times \cdots \times U_n$ and player i 's choice variable is $u_i \in U_i$, then the interaction between the players is modeled as a game. In the full information setting, the concept of a Nash equilibrium is used to characterize their interaction.

Definition 4. (Nash Equilibrium.) For an n -player game (J_1, \dots, J_n) , a point $u^* \in U$ is a Nash equilibrium if, for each $i \in \mathcal{I}$,

$$J_i(u_i^*, u_{-i}^*) \leq J_i(u_i, u_{-i}^*), \quad \forall u_i \in U_i. \quad (2.1)$$

Alternatively, using the preference relation notation,

$$(u_i^*, u_{-i}^*) \succeq_i (u_i, u_{-i}^*) \quad \forall u_i \in U_i$$

Remark 3. In words, for u^* to be a Nash equilibrium it must be the case that no player i has an action yielding an outcome that they prefer to the outcome generated when they choose u_i^* given that every other player j chooses u_j^* . That is, no player can profitably deviate, given the actions of the other players.

Let's consider some examples. The first example is a classic example in game theory used to demonstrate cooperative solutions as Nash equilibria.

Example 2.2 (Bach or Stravinsky (aka Battle of the Sexes).) Players wish to coordinate but have a conflict of interests.

For instance, consider two people that wish to go out together to a concert of music by either

Bach or Stravinsky. One prefers Bach and the other Stravinsky.

		P2	
		B	S
P1	B	(2,1)	(0,0)
	S	(0,0)	(1,2)

The game has two Nash equilibria: both choosing Bach (B, B) or both choosing Stravinsky (S, S) . Indeed, for (B, B)

$$2 = J_1(B, B) \geq J_1(S, B) = 0$$

$$1 = J_2(B, B) \geq J_2(B, S) = 0$$

and for (S, S) ,

$$1 = J_1(S, S) \geq J_1(B, S) = 0$$

$$2 = J_2(S, S) \geq J_2(S, B) = 0$$

where as for (B, S) ,

$$0 = J_1(B, S) \leq J_1(S, S) = 1$$

and for (S, B) ,

$$0 = J_1(S, B) \leq J_1(B, B) = 2$$

These choices can be thought of as “steady states”.

The next example is a classic example from game theory which demonstrates why completely rational players may not cooperate even though it is in the best interest to do so.

Example 2.3 (*Prisoner’s Dilemma*) Two suspects of a crime are brought into separate cells so that they cannot communicate with one another. The prosecutors do not have enough evidence to convict the criminals of the main serious crime, yet they have enough evidence to convict both for a misdemeanor. So, the prosecutors try the classic trick of getting the criminals to confess. Prosecutors enter the criminals cells (simultaneously) and offer each one a ‘deal’: each criminal is given the opportunity either to *betray* the other—by testifying that the other committed the crime—or to cooperate with the other by remaining *silent*. In this game, players are *utility maximizers*. The values in the matrix represent the number of years not in prison—i.e., a value of -2 means that the criminal will spend two years in prison. Based on the deal the prosecutors offer, the possible outcomes are defined in the following table:

		P2	
		Betray	Silent
P1	Betray	(-2,-2)	(0,-3)
	Silent	(-3,0)	(-1,-1)

Then, surprisingly both players choosing to betray one another (B, B) is a Nash equilibrium, whereas the cooperative solution (leading to a lower prison sentence for both) is not! Indeed, (B, B) is a Nash equilibrium since

$$-2 = J_1(B, B) \geq J_1(S, B) = -3 \quad \text{and} \quad -2 = J_2(B, B) \geq J_2(B, S) = -3,$$

while (S, S) is not since

$$-1 = J_1(S, S) \leq J_1(B, S) = 0 \quad \text{and} \quad -1 = J_2(S, S) \leq J_2(S, B) = 0.$$

The latter set of equations means that both players have an incentive to deviate from (S, S) as the other option (betray) has a lower prison sentence.

In addition, neither are (B, S) or (S, B) since for (B, S) , as we have stated above

$$0 = J_1(B, S) \geq J_1(S, S) = -1 \quad \text{yet} \quad -3 = J_2(B, S) \leq J_2(B, B) = -2$$

and, similarly, for (S, B)

$$0 = J_2(S, B) \geq J_2(S, S) = -1 \quad \text{yet} \quad -3 = J_1(S, B) \leq J_1(B, B) = -2.$$

The general form of prisoners' dilemma is given by

		P2	
		silent	betray
P1	silent	(R,R)	(S,T)
	betray	(T,S)	(P,P)

where $T > R > P > S$.

History. Prisoners' Dilemma was originally framed by Merrill Flood and Melvin Dresher while working at RAND in 1950. Albert W. Tucker formalized the game with prison sentence rewards and named it "prisoner's dilemma".

2.2.2 Best Response

The definition of a Nash equilibrium is sometimes restated in terms of *best response*.

Definition 5. *Best response mapping.* For any $u_{-i} \in U_{-i}$, define $B_i(u_{-i})$ to be the set of player i 's best actions given u_{-i} . That is,

$$B_i(u_{-i}) = \{u_i \in U_i \mid J_i(u_i, u_{-i}) \geq J_i(u'_i, u_{-i}) \quad \forall u'_i \in U_i\}$$

Note that $B_i(u_{-i})$ is potentially a set valued map.

Definition 6. (*Best Response Definition of a Nash Equilibrium.*) A Nash equilibrium $u^* \in U$ is a profile such that for each i ,

$$u_i^* \in B_i(u_{-i}^*)$$

To show a game has a Nash equilibrium it suffices to show that there is a profile u^* such that for each $i \in \mathcal{I}$,

$$u_i^* \in B_i(u_{-i}^*).$$

Definition 7. Let $B : U \rightarrow U$ be defined by $B(u) = \times_{i \in \mathcal{I}} B_i(u_{-i})$. Then we can rewrite the best response Nash condition as $u^* \in B(u^*)$.

Remark 4. Fixed point theorems give conditions on B under which there exists a value of u^* for which $u^* \in B(u^*)$.

2.2.3 Mixed Strategies

It may be the case, that a Nash equilibrium in pure strategies does not exist. The following example demonstrates that not every game as a Nash equilibrium in *pure* strategies.

Example 2.4 (Matching Pennies Revisited.)

		P2 (odd)	
		heads	tails
P1 (even)	heads	(+1,-1)	(-1,+1)
	tails	(-1,+1)	(+1,-1)

There are no Nash equilibria. Indeed, if we check the Nash equilibrium conditions for all possible strategies we will find that no joint action pair satisfies the conditions: e.g.,

$$1 = J_1(H, H) \geq J_1(T, H) = -1 \quad \text{yet} \quad -1 = J_2(H, H) \leq J_2(H, T) = 1$$

John Nash showed in his seminal 1951 paper that all finite games admit Nash equilibria in *mixed strategies*. Informally, the intuition is that there may not be a joint action profile (or *pure strategy*) which satisfies the Nash equilibrium conditions, however, if players are allowed to randomize (i.e., choose actions from a distribution with support in their action space), then it may be possible to satisfy the Nash equilibrium conditions in expectation given the distribution (or *mixed strategy*) for each player.

Indeed, the mixed strategy Nash equilibrium concept is designed to model the steady state of a game in which participants randomize their strategies.

Let $\Delta(U_i)$ be the set of probability distributions over U_i and refer to a member of $\Delta(U_i)$ as a *mixed strategy*. For instance, recall again the matching pennies game. Player 1's (even) strategy space is $U_1 = \{H, T\}$. Then let $p = (p_H, p_T) \in \Delta(U_1)$ where $p_H + p_T = 1$.

$$\Delta(U_i) = \left\{ \sigma_i = (\sigma_i^1, \dots, \sigma_i^{|U_i|}) : \sum_{j=1}^{|U_i|} \sigma_i^j = 1 \right\}$$

where $|\cdot|$ denotes the cardinality of the argument.

Remark 5. We will now refer to elements of U_i as *pure strategies*.

A profile $(\sigma_i)_{i \in \mathcal{I}}$ of mixed strategies induces a probability distribution over U .

Definition 8. (Extension of payoff functions to $\Delta(U) = \prod_{i \in \mathcal{I}} \Delta(U_i)$.) Given independence of the randomizations, the probability of the action profile $u = (u_i)_{i \in \mathcal{I}}$ is $\prod_{i \in \mathcal{I}} \sigma_i(u_i)$ so that player i 's

evaluation of $(\sigma_i)_{i \in \mathcal{I}}$ is

$$J_i(\sigma_i, \sigma_{-i}) = \sum_{u \in U} \left(\prod_{j \in \mathcal{I}} \sigma_j(u_j) \right) J_i(u)$$

Definition 9. (Mixed Nash Equilibrium.) A mixed strategy profile σ^* is a mixed strategy Nash equilibrium if for each player i ,

$$J_i(\sigma_i^*, \sigma_{-i}^*) \geq J_i(\sigma_i, \sigma_{-i}^*) \quad \forall \sigma_i \in \Delta(U_i)$$


Proposition 1. (Mixed Nash Necessary and Sufficient Conditions.) Let $\mathcal{G} = ((U_i)_{i \in \mathcal{I}}, (J_i)_{i \in \mathcal{I}})$ be a finite strategic form game. Then $\sigma^* \in \Delta(U)$ is a Nash equilibrium if and only if for each player $i \in \mathcal{I}$, every pure strategy in the support of σ_i^* is a best response to σ_{-i}^* .

Let us return once again to the Matching Pennies example. We will refer to player 1 as **P1** and player 2 as **P2**. Let $p = (p_H, p_T) = (p_H, 1-p_H)$ be the strategy for **P2** and $q = (q_H, q_T) = (q_H, 1-q_H)$ be the strategy for **P1**. Since $p_T = 1-p_H$ and $q_T = 1-q_H$, we only need to create the best response maps in terms of one variable as the other is determined from it.

P2 (odd)


		heads	tails
P1 (even)	heads	(+1,-1)	(-1,+1)
	tails	(-1,+1)	(+1,-1)

To plot the best response curve $q_H(p_H) = B_1(p_H)$ for **P1**, we need to consider that **P1** is the even player, meaning it desires for the two pennies to match either as (H, H) or (T, T) . This is to say that when **P1** observes **P2** playing any value of $0 \leq p_H < 0.5$, **P1** would like to place all the weight on $q_H = 0$ since this gives the best opportunity to have a (H, H) or (T, T) outcome. On the other hand, when **P1** observes **P2** playing any values of $0.5 < p_H \leq 1$, **P1** would like to place all the weight on $q_H = 1$ for similar reasons. With **P1** choosing $p_H = 0.5$, $0 \leq q_H \leq 1$ are all best responses to **P1** for **P2**.



On the other hand, to plot the best response curve for **P2**, i.e. $p_H(q_H) = B_2(q_H)$, we need to consider that **P2** is the odd player, meaning it desires for the two pennies to NOT-match either as (T, H) or (H, T) .

This is to say that when **P2** observes **P1** playing any value of $0 \leq q_H < 0.5$, **P2** would like to place all the weight on $p_H = 1$ since this gives the best opportunity to have a (H, T) or (T, H) outcome. On the other hand, when **P2** observes **P1** playing any values of $0.5 < q_H \leq 1$, **P2** would like to place all the weight on $p_H = 0$ for similar reasons. With **P2** choosing $q_H = 0.5$, $0 \leq p_H \leq 1$ are all best responses to **P2** for **P1**.



The unique mixed Nash equilibrium for Matching Pennies is thus $(p_H, q_H) = (0.5, 0.5)$.

Remark 6. Smoothed best response removes discontinuities—e.g., using a soft-max function such as

$$\frac{\exp(\mathbb{E}(x_1)/\gamma)}{\exp(\mathbb{E}(x_1)/\gamma) + \exp(\mathbb{E}(x_2)/\gamma)}$$

with smoothing parameter γ . This is particularly useful in capturing myopic behavior. Psychology experiments (Luce-Shepard rule [11], e.g.) verify this smoothed response (when individuals are roughly indifferent between two actions they appear to choose more or less at random). Some learning rules based on smooth best response (as we will see with Fictitious play) can result in players learning to play mixed Nash equilibrium [7].

2.2.4 Strict Dominance of Mixed Strategies

Often games are such that certain strategies dominate other strategies.

Definition 10. (*Dominant Strategies.*) *A strategy is dominant for a player if it yields the best payoff for that player no matter what strategies the other players choose. If all players have a dominant strategy, then it is natural for them to choose the dominant strategies and we reach a dominant strategy equilibrium.*

Recall the Prisoners' Dilemma game from Example 2.3.

		P2	
		Betray	Silent
P1	Betray	(-2,-2)	(0,-3)
	Silent	(-3,0)	(-1,-1)

The strategy (B, B) is a dominant strategy equilibrium since no matter the other player does its best for each player to stick with confessing.

Definition 11. *An action u_i is strictly dominated if there exists a mixed strategy $\sigma'_i \in \Delta(U_i)$ such that*

$$J_i(\sigma'_i, u_{-i}) > J_i(u_i, u_{-i}) \quad \forall u_{-i} \in U_{-i}$$

Strictly dominated strategies are never used with positive probability in a mixed strategy Nash equilibrium.

Example 2.5 (*Dominated Strategies.*)

P1 \ P2	L	R
U	(2, 0)	(-1, 0)
M	(0, 0)	(0, 0)
D	(-1, 0)	(2, 0)

In this example, player 1 (P1) has no pure strategies that strictly dominate M . However, M is strictly dominated by the mixed strategy $(0.5, 0, 0.5)$.

The following procedure iteratively eliminates strictly dominated strategies”

1. Let $U_i^0 = U_i$ and $\Delta^0(U_i) = \Delta(U_i)$.
2. For each i and each $k \geq 1$ define

$$U_i^k = \{u_i \in U_i^{k-1} \mid \nexists \sigma_i \in \Delta^{k-1}(U_i) \text{ s.t. } J_i(\sigma_i, u_{-i}) > J_i(u_i, u_{-i}) \quad \forall u_{-i} \in U_{-i}^{k-1}\}$$

3. Independently mix over U_i^k to generate/define $\Delta^k(U_i)$.
4. Let $U_i^\infty = \bigcap_{k=1}^\infty U_i^k$

The set U_i^∞ is the set of strategies for player i that survive iterated strict dominance.

Remark 7. *If U_i is finite for each i , then U_i^∞ is non-empty. However, it might not be a singleton. The order in which strategies are eliminated does not effect the set of strategies that survive.*

Example 2.6 (*Iterated Elimination of Dominated Strategies.*)

P1 \ P2	X	Y
A	(5, 2)	(4, 2)
B	(3, 1)	(3, 2)
C	(2, 1)	(4, 1)
D	(4, 3)	(5, 4)

Both B and C are dominated strategies for P1, and X is dominated for P2. We can reduce the game to

P1 \ P2	Y
A	(4, 2)
D	(5, 4)

In this newly reduced game, D dominates A for P1. Thus, we expect players choose (D, Y) .

2.3 Existence of Nash Equilibria

Theorem 1 (Nash, 1951). *Every finite game has a mixed strategy Nash equilibrium.*

Why is existence important? Without knowing whether an equilibrium exists, it is difficult to try and understand its properties.

Recall the following:

A mixed strategy x^* is a Nash equilibrium if and only if $x_i^* \in B_i(x_{-i}^*)$ for each i .

Define the best response correspondence $B : X \rightarrow X$ where $X = \times_{i \in \mathcal{I}} X_i$ and

$$X_i = \{x_i = (x_{i,1}, \dots, x_{i,n_i}) \mid \mathbf{1}^\top x_i = 1, x_i \geq 0\}$$

such that for all $x \in X$ we have

$$B(x) = (B_i(x_{-i}))_{i \in \mathcal{I}}$$

The existence of Nash is equivalent to the existence of a mixed strategy x such that $x \in B(x)$ —that is x is a fixed point of the map $B(\cdot)$.


Remark 8. *Fixed point theorems give conditions on B under which there exists a value of x^* for which $x^* \in B(x^*)$.*

Lemma 1. (*Kakutani's Fixed Point Theorem*). *Let X be a compact, convex³ subset of \mathbb{R}^n and let $f : X \rightarrow X$ be a set valued map for which the following hold:*

- for each $x \in X$, the set $f(x)$ is non-empty and convex;
- the graph of f is closed (i.e., if $\{x_n, y_n\} \rightarrow \{x, y\}$ with $y_n \in f(x_n)$ for all n , then $y \in f(x)$).

Then, there exists $x^ \in X$ such that $x^* \in f(x^*)$.*

³A set in Euclidean space is compact if and only if it is bounded and closed. A set S is convex if for any $x, y \in S$ and any $\lambda \in [0, 1]$, $\lambda x + (1 - \lambda)y \in S$.



The proof of Nash's theorem is now fairly straightforward.

In order to apply Kakutani's theorem to the best response map $B : X \rightarrow X$, we need to show that B satisfies the conditions of the theorem.

Step 1: Check that X is compact, convex, and non-empty.

Indeed, by definition $X = \times_{i \in \mathcal{I}} X_i$ where each $X_i = \{x_i = (x_{i,1}, \dots, x_{i,n_i}) \mid \mathbf{1}^\top x_i = 1, x_i \geq 0\}$ is a simplex of dimension $|S_i| - 1 = n_i - 1$, thus each X_i is closed and bounded, hence compact. Their product is also compact.

Step 2: Check that $B(x)$ is non-empty.

To do this we need the following theorem from Weierstrass.

Theorem 2 (Weierstrass). *Let Y be a non-empty compact subset of a finite dimensional Euclidean space and let $f : Y \rightarrow \mathbb{R}$ be a continuous function. Then f must attain a maximum and a minimum, each at least once. That is, there exists an optimal solution to*

$$\min\{f(y) \mid y \in Y\}$$

Now, by definition,

$$B_i(x_{-i}) = \arg \max_{y \in X_i} f_i(y, x_{-i})$$

where X_i is non-empty and compact, and f_i is linear in y . Hence, f_i is continuous and by Weierstrass' theorem $B(x)$ is non-empty.

Step 3: Check that $B(x)$ is a convex mapping.

We will argue an equivalent statement: $B(x) \subset X$ is convex if and only if $B_i(x_{-i})$ is convex for all i . Let $z_i, y_i \in B_i(x_{-i})$. Then,

$$f_i(y_i, x_{-i}) \geq f_i(v, x_{-i}) \quad \forall v \in X_i$$

and

$$f_i(z_i, x_{-i}) \geq f_i(v, x_{-i}) \quad \forall v \in X_i.$$

This, in turn, implies that for all $\lambda \in [0, 1]$, we have

$$\lambda f_i(y_i, x_{-i}) + (1 - \lambda) f_i(z_i, x_{-i}) \geq f_i(v, x_{-i})$$

since

$$\lambda f_i(y_i, x_{-i}) \geq \lambda f_i(v, x_{-i}) \quad \forall v \in X_i$$

and

$$(1 - \lambda) f_i(z_i, x_{-i}) \geq (1 - \lambda) f_i(v, x_{-i}) \quad \forall v \in X_i$$

so that

$$\begin{aligned} \lambda f_i(y_i, x_{-i}) + (1 - \lambda) f_i(z_i, x_{-i}) &\geq \lambda f_i(y_i, x_{-i}) + (1 - \lambda) f_i(v, x_{-i}) \\ &\geq \lambda f_i(v, x_{-i}) + (1 - \lambda) f_i(v, x_{-i}) \\ &= f_i(v, x_{-i}). \end{aligned}$$

By linearity of f_i , we have that

$$f_i(\lambda y_i + (1 - \lambda) z_i, x_{-i}) \geq f_i(v, x_{-i}) \quad \forall v \in X_i.$$

Hence, $\lambda y_i + (1 - \lambda) z_i \in B_i(x_{-i})$ showing that each B_i is convex so that B is convex.

Step 4: Check that $B(x)$ also has a closed graph.

Indeed, we show this by contradiction. Suppose not—i.e., $B(x)$ is assumed to not have a closed graph. Then, \exists a sequence $(x^k, \hat{x}^k) \rightarrow (x, \hat{x})$ with $\hat{x}^k \in B(x^k)$, but $\hat{x} \notin B(x)$, i.e., \exists some i such that $\hat{x}_i \notin B_i(x_{-i})$. This implies that \exists some $x'_i \in X_i$ and some $\varepsilon > 0$ s.t.

$$f_i(x'_i, x_{-i}) > f_i(\hat{x}_i, x_{-i}) + 3\varepsilon. \quad (*)$$

Since f_i is continuous and $x_{-i}^k \rightarrow x_{-i}$, we have that for sufficiently large k ,

$$f_i(x'_i, x_{-i}^k) \geq f_i(x'_i, x_{-i}) - \varepsilon \quad (**)$$

Combining (*) and (**), we have that

$$f_i(x'_i, x_{-i}^k) > f_i(\hat{x}_i, x_{-i}) + 2\varepsilon \geq f_i(\hat{x}_i^k, x_{-i}^k) + \varepsilon$$

where the section inequality follows from the continuity of f_i . This contradicts the assumption that $\hat{x}_i^k \in B_i(x_{-i}^k)$.

We have now shown that X and $B(\cdot)$ satisfy the assumptions of Kakutani's fixed point theorem, and hence we can apply it to get the existence of Nash.

2.4 Computation of Nash Equilibria in Finite Games

As we have seen, finite strategic form games can be represented in terms of the payoff matrices. In the special case of two player settings, such games are referred to as *bi-matrix* games. For this section, we will focus on bi-matrix games and how optimization techniques can be used to compute equilibria in such games.

Consider a game with two players characterized as follows.

- Player 1 (P1): strategy space $S_1 = \{s_1^1, \dots, s_{m_1}^1\}$ with $|S_1| = m_1$, payoff matrix $A \in \mathbb{R}^{m_1 \times m_2}$
- Player 2 (P2): strategy space $S_2 = \{s_1^2, \dots, s_{m_2}^2\}$ with $|S_2| = m_2$, payoff matrix $B \in \mathbb{R}^{m_1 \times m_2}$

Let x denote the mixed strategy of P1, i.e. $x \in X = \Delta(S_1)$ where

$$X = \left\{ x : \sum_{i=1}^{m_1} x_i = 1, x_i \geq 0 \right\}.$$

Similarly, let y denote the mixed strategy of P2, i.e. $y \in Y = \Delta(S_2)$ where

$$Y = \left\{ y : \sum_{i=1}^{m_2} y_i = 1, y_i \geq 0 \right\}.$$

Given a mixed strategy profile (x, y) , the payoffs of P1 and P2 can be expressed in terms of the payoff matrices as follows:

$$J_1(x, y) = x^\top A y \quad \text{and} \quad J_2(x, y) = x^\top B y$$

Recall from the proceeding section that we used a single object to represent the game: e.g.,

P1 \ P2	s_1^1	s_2^1
s_1^2	(a_{11}, b_{11})	(a_{12}, b_{12})
s_2^2	(a_{21}, b_{21})	(a_{22}, b_{22})
s_3^2	(a_{31}, b_{31})	(a_{32}, b_{32})

Hence, the payoff matrices⁴ in this case are just

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \\ b_{31} & b_{32} \end{bmatrix}.$$

Example 2.7 (Utility Representation in Bi-Matrix Games.) Suppose $m_1 = 3$ and $m_2 = 2$ as above. Then $x = (x_1, x_2, x_3)$ with $\sum_{i=1}^{m_1} x_i = 1$, $x_i \geq 0$ and x_i is the probability assigned to strategy s_i^1 in P1's randomization over S_1 . Similarly, $y = (y_1, y_2)$ with $y_1 + y_2 = 1$, $y_i \geq 0$ and y_i the probability P2 assigns to strategy s_i^2 . The payoffs are given by

$$J_1(x, y) = x^\top A y = [x_1 \quad x_2 \quad x_3] \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \sum_{k=1}^2 \sum_{j=1}^3 y_k x_j a_{jk}$$

$$J_2(x, y) = x^\top B y = [x_1 \quad x_2 \quad x_3] \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \\ b_{31} & b_{32} \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \sum_{k=1}^2 \sum_{j=1}^3 y_k x_j b_{jk}$$

Let's redefine Nash in terms of this new notation.

Definition 12. (Mixed Nash.) A mixed strategy profile (x^*, y^*) is a mixed strategy Nash if and only if

$$(x^*)^\top A y^* \geq x^\top A y^* \quad \forall x \in X$$

$$(x^*)^\top B y^* \geq (x^*)^\top B y \quad \forall y \in Y$$

⁴Note that some references will write

$$J_2(x, y) = y^\top B^\top x$$

and define B^\top to be the payoff matrix for player 2. Keep this in mind when comparing these notes to other references.

2.4.1 Zero Sum Games

In zero sum settings, it is the case that $B = -A$, and hence, the definition of Nash reduces to the following two inequalities:

$$\begin{aligned}(x^*)^\top Ay^* &\geq x^\top Ay^* \quad \forall x \in X \\ (x^*)^\top Ay^* &\leq (x^*)^\top Ay \quad \forall y \in Y\end{aligned}$$

which we can combine to get

$$x^\top Ay^* \leq (x^*)^\top Ay^* \leq (x^*)^\top Ay \quad \forall x \in X, y \in Y.$$

That is, a Nash equilibrium (x^*, y^*) is a *saddle point* of the function $x^\top Ay$ defined over $X \times Y$. Formally,

Definition 13. A vector (x^*, y^*) is a **saddle point** if $x^* \in X$ and $y^* \in Y$ and

$$\sup_{x \in X} x^\top Ay^* = (x^*)^\top Ay^* = \inf_{y \in Y} (x^*)^\top Ay \quad (*)$$

For any function $f : X \times Y \rightarrow \mathbb{R}$ we have the so-called *minimax inequality* which is given by

$$\sup_{x \in X} \inf_{y \in Y} f(x, y) \leq \inf_{y \in Y} \sup_{x \in X} f(x, y)$$

Indeed, it is easy to see why the above inequality holds by the following reasoning: for every $\tilde{x} \in X$ write

$$\inf_{y \in Y} f(\tilde{x}, y) \leq \inf_{y \in Y} \sup_{x \in X} f(x, y)$$

and take the sup over $\tilde{x} \in X$ of the left-hand side—i.e., it has to be true for all values, so certainly true for the supremum.

Proposition 2. (*Nash Equilibria in Bi-Matrix Games.*)

$$(x^*, y^*) \text{ is N.E.} \iff (x^*)^\top Ay^* = \inf_{y \in Y} \sup_{x \in X} x^\top Ay = \sup_{x \in X} \inf_{y \in Y} x^\top Ay$$

Proof. Indeed,

$$(*) \implies \inf_{y \in Y} \sup_{x \in X} x^\top Ay \leq \sup_{x \in X} x^\top Ay^* = (x^*)^\top Ay^* = \inf_{y \in Y} (x^*)^\top Ay \leq \sup_{x \in X} \inf_{y \in Y} x^\top Ay$$

Combining this with the minimax inequality, we get that equality holds throughout. \square

Definition 14. (*Game Value.*) The value of the game is $(x^*)^\top Ay^*$.

Surprisingly, finding mixed Nash equilibrium of a finite zero-sum game can be written as a pair of linear optimization problems. Indeed, for a fixed y , we have

$$\max_{x \in X} x^\top Ay = \max_{i=1, \dots, m_1} \{[Ay]_i\}$$

and thus,

$$\min_{y \in Y} \max_{x \in X} x^\top Ay = \min_{y \in Y} \max\{[Ay]_1, \dots, [Ay]_{m_1}\} = \min_{y \in Y} \{v \mid v \mathbf{1}_n \geq Ay\}$$

Similarly,

$$\max_{x \in X} \min_{y \in Y} x^\top Ay = \max_{x \in X} \min\{[A^\top x]_1, \dots, [A^\top x]_{m_2}\} = \max_{x \in X} \{\xi \mid \xi \mathbf{1}_{m_2} \leq A^\top x\}$$

Remark 9. Linear programs can be solved in polynomial (in m_1 and m_2) time.

2.4.2 Non-Zero Sum Games

The fact that finding Nash in the zero sum setting reduces to solving a pair of linear programs is nice outcome. Does this extend to the non-zero sum setting?

Recall that the joint strategy space is the cross product of simplexes:

$$\left\{ x \mid x = (x_1, \dots, x_{m_1}), \sum_i x_i = 1 \right\} \times \left\{ y \mid y = (y_1, \dots, y_{m_2}), \sum_i y_i = 1 \right\}$$

Nash equilibria can be strictly in the interior—that is, $x_i > 0$ and $y_j > 0$ for each $i \in \{1, \dots, m_1\}$, $j \in \{1, \dots, m_2\}$ —or on the boundary. Let us start by considering points on the interior first (inner or totally mixed Nash equilibria).

Let a_i denote the *rows* of A and let b_j denote the *columns* of B . Recall the following characterization of a Nash equilibrium: a point (x^*, y^*) is a Nash equilibrium if and only if every pure strategy in the support of x^* is a best response to y^* . This implies that all pure strategies in the support of a Nash equilibrium yield the same payoff which is also greater than or equal to the payoffs for strategies outside the support.

If $\text{supp}(x^*)$ denotes the support of x^* , then for all $s_j, s_k \in \text{supp}(x^*)$ we have $J_1(s_j, y^*) = J_1(s_k, y^*)$ and for each $s_i \in \text{supp}(x^*)$ and $s_\ell \notin \bar{S}_1$ we have $J_1(s_i, y^*) \geq J_1(s_\ell, y^*)$.

In other words, for purely mixed strategies (that is, ones where the support is the whole strategy space, i.e. $\text{supp}(x) = S_1$ and $\text{supp}(y) = S_2$) we have that

$$(x, y) \text{ is a Nash equilibrium} \iff \begin{cases} a_1 y & = a_i y & i = 2, \dots, m_1 \\ x^\top b_1 & = x^\top b_j & j = 2, \dots, m_2 \\ \sum_{i=1}^{m_1} x_i & = 1 \\ \sum_{j=1}^{m_2} y_j & = 1 \end{cases}$$

This is a set of linear equations and can be solved efficiently.

Exercise 4.1 Try solving this for matching pennies to find the Nash equilibria. Try it also using the zero sum calculation formulation (since matching pennies is a zero-sum game).

Remark 10. For $n \geq 2$ (more than 2 players), we get a set of polynomial equations and we will talk more about this case later.

We note that not all games have purely mixed Nash strategies!

Let us now formulation a naive or brute force extension of the purely mixed strategy computation strategy for cases where we do not have purely mixed strategies.

A mixed strategy profile $(x, y) \in X \times Y$ is a Nash equilibrium with $\text{supp}(x) \subset S_1$ and $\text{supp}(y) \subset S_2$ if and only if

$$\begin{cases} u = a_i y & \forall s_i^1 \in \text{supp}(x) \\ u \geq a_i y & \forall s_i^1 \notin \text{supp}(x) \\ v = x^\top b_j & \forall s_j^2 \in \text{supp}(y) \\ v \geq x^\top b_j & \forall s_j^2 \notin \text{supp}(y) \\ x_i = 0 & \forall s_i^1 \notin \text{supp}(x) \\ y_j = 0 & \forall s_j^2 \notin \text{supp}(y) \end{cases}$$

The issue with this computation method is that it requires finding the *right* support sets $\text{supp}(x)$ and $\text{supp}(y)$. There are $2^{m_1+m_2}$ different support sets; hence, this process is exponential in computation time!

Optimization Formulation. As an alternative, we are going to formulate the problem of computing Nash as a nonlinear program and use optimization tools that we know and are familiar with.

Remark 11. *We will make a similar connection for continuous games and we will use an inverse formulation for designing feedback mechanisms to induce a Nash equilibrium. So this optimization framework is quite important.*

Recall the set of linear equations we wrote down for interior Nash equilibria. We can use those to define best response polytopes.

Definition 15. (*Best Response Polytopes.*)

$$P = \{(x, v) \in \mathbb{R}^{m_1+1} \mid x_i \geq 0, i = 1, \dots, m_1, \mathbf{1}^\top x = 1, B^\top x \leq \mathbf{1}v\}$$

$$Q = \{(y, u) \in \mathbb{R}^{m_2+1} \mid y_j \geq 0, j = 1, \dots, m_2, \mathbf{1}^\top y = 1, Ay \leq \mathbf{1}u\}$$

Let $\text{NE}(A, B)$ be the set of Nash equilibria for the bi-matrix game (A, B) .

Lemma 2. *All mixed strategy Nash equilibria have zero optimal value to (P-1).*

Proof. Recall that if a point $(x^*, y^*) \in \text{NE}(A, B)$ then x^* is a best response to y^* and vice versa. That is,

$$x^* \in \arg \max\{x^\top Ay^* \mid x \in X\}$$

The dual of this problem is

$$\min_p \{p \mid \mathbf{1}_n p \geq Ay\}$$

Feasible points are optimal if and only if the two objective points are equal $(x^*)^\top Ay^* = p^*$. A similar argument holds for $(x^*)^\top By^* = q^*$. Thus,

$$(x^*)^\top Ay^* + (x^*)^\top By^* - p^* - q^* = 0$$

□

Applying this lemma, we get that

$$((x, v), (y, u)) \in P \times Q \text{ s.t. } (x, y) \in \text{NE}(A, B) \iff \begin{cases} x^\top (Ay - \mathbf{1}u) = 0 \\ (x^\top B - \mathbf{1}v)y = 0 \end{cases}$$

Proposition 3. (*Necessary and Sufficient Conditions for Nash in Bi-Matrix Games.*) *A mixed profile $(x, y) \in X \times Y$ is a Nash equilibrium of the bimatrix game (A, B) if and only if $\exists (p, q)$ such that (x, y, u, v) is a solution to*

$$\max x^\top Ay + x^\top By - u - v \tag{P-1}$$

$$\text{s.t. } Ay \leq u\mathbf{1}_{m_1} \tag{2.2}$$

$$B^\top x \leq v\mathbf{1}_{m_2} \tag{2.3}$$

$$x \geq 0, y \geq 0 \tag{2.4}$$

$$\sum_{i=1}^{m_1} x_i = 1, \sum_{j=1}^{m_2} y_j = 1 \tag{2.5}$$

proof sketch. Clearly, for $((x, v), (y, u)) \in P \times Q$ we have that $x^\top(Ay - \mathbf{1}u) \leq 0$ and $(x^\top B - \mathbf{1}^\top)v \leq 0$ so that

$$x^\top(Ay - \mathbf{1}u) + (x^\top B - \mathbf{1}^\top)v \leq 0$$

Using the fact that $x^\top \mathbf{1} = 1$ and $\mathbf{1}^\top y = 1$, we have

$$x^\top Ay + x^\top By - u - v \leq 0$$

and equality holds if and only if $(x, y) \in \text{NE}(A, B)$.

Definition 16. A game (A, B) is a symmetric game if $B = A$ (the strategy sets S_1 and S_2 are the same).

Proposition 4. (Nash, 1951.) Any symmetric finite game has a symmetric Nash equilibrium. That is, if $(x, y) \in \text{NE}(A, A)$ then $x = y$.

In this case, a Nash equilibrium x satisfies $u = v$ and $Ax \leq \mathbf{1}u$, and $x^\top Ax - u = 0$. We can then formulate a quadratic program for finding all symmetric Nash equilibrium. Indeed,

$$\begin{aligned} \max \quad & x^\top Ax - u \\ \text{s.t.} \quad & Ax \leq \mathbf{1}u, \quad \mathbf{1}^\top x = 1, \quad x \geq 0 \end{aligned}$$

Any bimatrix game (A, B) can be converted to an equivalent symmetric game (\tilde{A}, \tilde{A}) where

$$\tilde{A} = \begin{bmatrix} 0 & A \\ B^\top & 0 \end{bmatrix}$$

with strategy vector

$$z = \begin{bmatrix} x \\ y \end{bmatrix}$$

Exercise 4.2 Check that any symmetric Nash equilibrium of the game (\tilde{A}, \tilde{A}) is a Nash equilibrium of the original game and vice-versa.

The following are some additional noteworthy items:

- The optimization formulation can be generalized to multi-player finite games.
- Work in 2006 by Pablo Parrilo created a semi-definite programming formulation for two person zero-sum games with continuous strategy spaces and some assumptions on the structure of the payoff functions. His work shows that Nash equilibria can be computed efficiently [13].
- Other work focuses on developing polynomial time algorithms for *rank-based games*; e.g., the rank of a game is defined to be $\text{rank}(A + B)$. So, for zero sum games we have $\text{rank}(A + B) = \text{rank}(A - A) = 0$ and Nash equilibria in zero sum games can be found efficiently. However, work by Garg, Jiang and Mehta in 2011 shows results for rank 1 games [8].

2.5 Correlated Equilibria in Strategic Form Games

Recall that in a Nash equilibrium, players randomize over strategies independently. A natural extension of this is to explore correlation between players strategies. This leads to the refinement of the Nash equilibrium concept known as correlated equilibrium:

$$\text{Pure Nash} \subset \text{Mixed Nash} \subset \text{Correlated Equilibria}$$

Indeed, it may be the case for games with multiple Nash equilibria that we want to allow for randomizations between Nash equilibria by some form of communication prior to the play of the game.

Example 2.8 (*Bach or Stravinsky.*) Recall the game in which players are going out to a concert and must choose between Bach and Stravinsky.

P1\P2	Bach	Stravinsky
Bach	(1, 4)	(0, 0)
Stravinsky	(0, 0)	(4, 1)

Suppose the players flip a coin and go to Bach if the coin is heads and Stravinsky otherwise. The payoff to the players in this case is $(1/2(1 + 4), 1/2(4 + 1)) = (5/2, 5/2)$ and this is not a Nash equilibrium. The coin flip acts as a correlation device.

Example 2.9 (*Using a correlation device to find Nash.*)

P1\P2	L	R
U	(5, 1)	(0, 0)
D	(4, 4)	(1, 5)

Both (U, L) and (D, R) are pure strategy Nash equilibria of the game.

To find the mixed Nash equilibrium we will recall that for a Nash equilibrium (x, y) it has to be the case that $a_1y = a_2y$ and $x^\top b_1 = x^\top b_2$ where a_i are rows of A and b_j are columns of B . Let P1 play U with probability p and P2 player L with probability q . Then

$$\begin{aligned} 5q &= 4q + (1 - q) \implies q = \frac{1}{2} \\ 5p &= 4p + (1 - p) \implies p = \frac{1}{2} \end{aligned}$$

The unique Nash is $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ and its payoff is $(5/2, 5/2)$.

Note for a Nash equilibrium, people randomize independently. For games with multiple Nash equilibria, one may want to allow for randomizations between Nash equilibria by some form of communication prior to the play of the game.

A generalization is called *correlated equilibrium* [1]. In this set up, players observe a correlating signal before making their choice (think about introducing a traffic light or stop sign in the at an intersection).

Definition 17. Correlated Equilibrium A joint probability distribution \mathbf{x} over S is a correlated equilibrium of \mathcal{G} if, for each i ,

$$\mathbb{E}_{s \sim \sigma}[J_i(s_i, s_{-i}) | s_i] \geq \mathbb{E}_{s_{-i} \sim \sigma_{-i}}[J_i(s'_i, s_{-i}) | s_i], \quad \forall s'_i \in S_i$$

Correlated equilibria can be interpreted as follows: there is a mediator who

1. draws an outcome s from the publicly known distribution σ ,
2. and privately recommends strategy s_i to each player i .

Correlated equilibria require the expected payoff from playing the recommended strategy is greater than or equal to playing any other strategy.

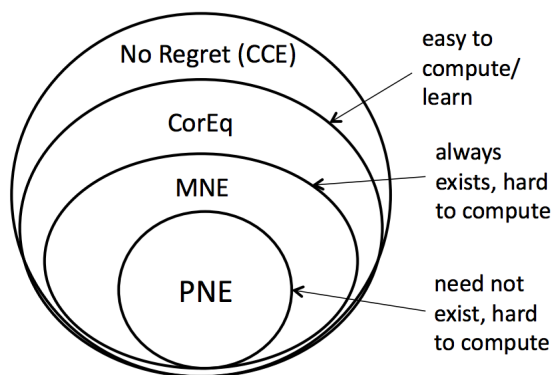


Figure 2.2: Refinements of the Nash equilibrium concept.

Example 2.10 (Example: traffic Light—Correlated equilibrium but not mixed Nash)

Consider the following game:

$$\begin{bmatrix} P1 \backslash P2 & \text{stop} & \text{go} \\ \text{stop} & (0, 0) & (0, 1) \\ \text{go} & (1, 0) & (-5, -5) \end{bmatrix}$$

If the other player is stopping at an intersection, then you would rather go and get on with it. The worst-case scenario, of course, is that both players go at the same time and get into an accident. There are two pure Nash equilibria: (stop, go) and (go, stop). Define σ by randomizing 50/50 between these two pure Nash equilibria. This is not a product distribution, so it cannot be a mixed Nash equilibrium. It is, however, a correlated equilibrium. Indeed, consider P1 (row player). If the mediator (stop light) recommends *go* then P1 knows that P2 was recommended *stop*. Assuming P2 plays the recommended strategy *stop*, P1's best response is to follow its recommended strategy *go*. Similarly, when P1 is told to *stop*, it assumes that P2 will *go*, and under this assumption stopping is a best response.

The correlated equilibrium concept requires that following the recommended strategy be a best response in the interim stage—i.e., a correlated equilibrium requires that after a profile s is drawn and recommended, playing s_i is a best response for i conditioned on seeing s_i and given everyone else plays s . There are a number of further refinements of the Nash equilibrium concept including the notion of a *coarse correlated equilibrium* (CCE), which only requires the recommended strategy to be a best response at the *ex-ante* stage—i.e., CCE requires only that following the suggested strategy s_i when $s \sim \sigma$ is a best response in expectation before you see s_i .

Definition 18. (*Coarse Correlated Equilibrium.*) A coarse correlated equilibrium (CCE) is a distribution σ over actions S such that for every player i , and every action s'_i

$$\mathbb{E}_{s \sim \sigma}[J_i(s)] \geq \mathbb{E}_s[J_i(s'_i, s_{-i})]$$

Hannan Set/No-regret Equilibrium: The set of all distributions satisfying the above is sometimes called the Hannan set.

Refinements of the Nash equilibrium concept generally have arisen due to the pursuit of theory for how players arrive at an equilibrium. One theory for how players arrive at an equilibrium is

through a process of (myopic) tâtonnement, thereby giving rise to the theory 'learning in games' [7]. While the subject of learning in games is relegated to a later chapter, we mention it here to note the importance of refinements of Nash. It turns out that not only are refinements easier to compute (as Figure 2.2 suggests) but also are arguably easier to learning through tâtonnement.

2.5.1 Computation of Correlated Equilibria

I will expand this section

2.6 Stackelberg Equilibria in Strategic Form Games

I will expand this section: breif introduction to minimax and connections to Nash

2.7 Games with Continuous Actions Spaces

Now instead of having a finite strategy space (that is players choose actions from a set of finite choices—e.g., heads or tails) players choose from infinitely many pure strategies.

In particular, we consider games where players choose from a continuum and the utility functions are continuous. Mixed strategies over finite games can be thought of as pure strategies in a continuous game over the simplex.

Example 2.11 (Cournot Game.) Consider a scenario in which there are two firms 1 and 2. They simultaneously choose any quantity $q_1, q_2 \geq 0$. The price in the market is given by $p(q_1, q_2) = A - q_1 - q_2$ for some constant A . The firms payoff functions are given by

$$\pi_1(q_1, q_2) = (A - q_1 - q_2)q_1 - c_1q_1$$

$$\pi_2(q_1, q_2) = (A - q_1 - q_2)q_2 - c_2q_2$$

Calculate the equilibrium:

$$\frac{\partial \pi_1}{\partial q_1} = A - 2q_1 - q_2 - c_1 = 0$$

$$\frac{\partial \pi_2}{\partial q_2} = A - q_1 - 2q_2 - c_2 = 0$$

The values of q_i that satisfy this equation are best responses. The Nash equilibria are where both q_1 and q_2 are best response given those values of q_1 and q_2 . That is

$$q_1^* = \frac{1}{3}(A - 2c_1 + c_2) \quad \text{and} \quad q_2^* = \frac{1}{3}(A - 2c_2 + c_1)$$

Remark 12. I will show you why in general we do not want to just take the derivative and set it to zero even for simple games. We will show that in fact for simple games you can get an entire continuum of Nash equilibria (meaning they are not isolated).

Definition 19. A point (x_1, \dots, x_n) is a Nash equilibrium for the continuous game if for each $i \in \mathcal{I}$

$$f_i(x_i, x_{-i}) \geq f_i(y, x_{-i}) \quad \forall y \in X_i$$

2.7.1 Best Response

Best response or rational response curves are again useful.

Definition 20. In an n -person nonzero sum game, let the maximum utility function of P_1 , $f_1(x_1, \dots, x_n)$ with respect to $x_1 \in X_1$ be attained for each $x_{-1} \in X_{-1}$ where $x_{-1} = (x_2, \dots, x_n)$. Then the set $R_1(x_{-1}) \subset X_1$ defined by

$$R_1(x_{-1}) = \{\xi \in X_1 : f_1(x_1, x_{-1}) \geq f_1(\xi, x_{-1}), \forall x_1 \in X_1\}$$

is called the optimal response or rational reaction set for P_1 . Similarly for each other player.

2.7.2 Existence of Nash Continuous games

Theorem 3 (Debreu [5], Glicksberg [9], Fan [6]). Consider a strategic form game $((X_i)_{i \in \mathcal{I}}, (f_i)_{i \in \mathcal{I}})$ such that for each $i \in \mathcal{I}$, the following hold:

- X_i is compact and convex
- $f_i(x_i, x_{-i})$ is continuous in x_{-i}
- $f_i(x_i, x_{-i})$ is continuous and concave (utility maximizers) in x_i

Then, a pure strategy Nash equilibrium exists.

We note that quasi-concavity is enough for the proof. Dasgupta and Maskin showed in 1986 that pure strategy equilibria exists in discontinuous games [4, 3]—that is, the utilities do not need to be continuous and instead, we need only quasi-concavity and upper-semicontinuity in x and graph continuity.

Definition 21. $f_i : X \rightarrow \mathbb{R}$ is **upper semicontinuous** if for any sequence $\{x^n\} \subset X$ such that $x^n \rightarrow x$, we have that

$$\limsup_{n \rightarrow \infty} f_i(x^n) \leq f_i(x)$$

Definition 22. $f_i : X \rightarrow \mathbb{R}$ is **graph-continuous** if for all $\bar{x} \in X$, \exists a function $F_i : X_{-i} \rightarrow X_i$ with $F_i(\bar{x}_{-i}) = \bar{x}_i$ such that $f_i(F_i(x_{-i}), x_{-i})$ is continuous at $x_{-i} = \bar{x}_{-i}$.

Note this is called *graph continuity* because if one graphs a player's payoff as a function of his own strategy (holding the strategies of the other players fixed) and if this graph changes continuously as one varies the strategies of the other players, then the player's payoff function is graph continuous in the sense of this definition.

History. Glicksberg showed in 1952 [9] that for compact strategy spaces and continuous utility functions there exists *mixed strategy equilibrium* for the game. Dasgupta and Maskin extended this result in 1986 as well to discontinuous games where the utilities are not even quasi-concave [4].

Example 2.12 (Relaxing Quasi-continuity.) Consider the location game where players choose a point on the circle and their payoffs are

$$f_1(s_1, s_2) = -f_2(s_1, s_2) = \|s_1 - s_2\|_2^2$$

Then there is no pure strategy Nash, yet there is a mixed Nash in this game where both players mix uniformly on the circle.

2.8 Optimality Conditions to Characterize Local Nash Equilibria

We may consider using optimality conditions to find Nash equilibria. Suppose that $X_i \subset \mathbb{R}^{m_i}$ for each i and let

$$X = X_1 \times \cdots \times X_n$$

Here we will only concern ourselves with full information games and reduce the notation for a game to be the tuple

$$\mathcal{G} = (f_1, \dots, f_n, X)$$

We define the "differential game form" as follows.

Definition 23 ([14]). *For a game $\mathcal{G} = (f_1, \dots, f_n, X)$, the "differential game form" is the collection of individual derivatives—i.e.*

$$\omega(x) = (D_1 f_1, \dots, D_n f_n)$$

where $D_i f_i = \partial f_i / \partial x_i$.

These individual derivatives are the directions in which a player can make "differential" adjustments to its strategy in order to improve the cost locally.

Note this definition can be extended to a well-defined differential form in settings where the strategy spaces of agents are manifolds (without boundary):

Definition 24 ([14]). *The differential game form for a game $\mathcal{G} = (f_1, \dots, f_n, X)$ is given by*

$$\omega = \sum_{i=1}^n \sum_{j=1}^{n_i} \frac{\partial f_i}{\partial x_i^j} dx_i^j$$

where this is a mapping (a actual differential form) $\omega : X_1 \times \cdots \times X_n \rightarrow T^*(X_1 \times \cdots \times X_n)$.

Definition 25 ([14]). *A strategy $x = (x_1, \dots, x_n)$ is a **differential Nash equilibrium** if $\omega(x) = 0$ (that is $D_i f_i(x) = 0$ for each i) and $D_i^2 f_i(x)$ is (strict) positive definite (resp. negative definite if players are utility maximizers).*

The conditions of the definition are sufficient for checking if a point is a Nash equilibrium. These look just like first and second order conditions for optimality. But they are not necessary. This gap between the necessary and sufficient conditions is important because it can be the case that you get a continuum of Nash equilibria (as the next example will show) even for the case where the cost functions are concave (resp. convex)!! For this example, let us switch to thinking about cost minimizers instead of utility maximizers.

Example 2.13 (Betty-Sue Thermodynamic Coupling.) Consider a two player game between Betty and Sue. Let Betty's strategy space be $X_1 = \mathbb{R}$ and her cost function

$$f_1(x_1, x_2) = \frac{1}{2}x_1^2 - x_1x_2.$$

Similarly, let Sue's strategy space be $X_2 = \mathbb{R}$ and her cost function

$$f_2(x_1, x_2) = \frac{1}{2}x_2^2 - x_1x_2.$$

This game can be thought of as an abstraction of two agents in a building occupying adjoining rooms where cross terms model the effect of heat transfer.

The first term in each of their costs represents an energy cost and the second term is a cost from thermodynamic coupling. The agents try to maintain the temperature at a desired set–point in thermodynamic equilibrium.

The definition of Nash specifies that a point (x_1^*, x_2^*) is a Nash equilibrium if no player can unilaterally deviate and decrease their cost, i.e.

$$f_1(x_1^*, x_2^*) \leq f_1(x_1, x_2^*), \quad \forall x_1 \in \mathbb{R}$$

and

$$f_2(x_1^*, x_2^*) \leq f_2(x_1^*, x_2), \quad \forall x_2 \in \mathbb{R}.$$

Fix Sue’s strategy $x_2 = z$, and calculate

$$D_1 f_1 = \frac{\partial f_1}{\partial x_1} = x_1 - z \tag{2.6}$$

Then, Betty’s optimal response to Sue playing $x_2 = z$ is $x_1 = z$. Similarly, if we fix Betty’s strategy as $x_1 = y$, then Sue’s optimal response is $x_2 = y$. For all $x_1 \in \mathbb{R} \setminus \{z\}$,

$$-\frac{1}{2}(z)^2 < \frac{1}{2}x_1^2 - x_1z \tag{2.7}$$

so that

$$f_1(z, z) < f_1(x_1, z), \quad \forall x_1 \in \mathbb{R} \setminus \{z\}.$$

Again, similarly, for all $x_2 \in \mathbb{R} \setminus \{y\}$,

$$-\frac{1}{2}y^2 < \frac{1}{2}x_2^2 - x_2y \tag{2.8}$$

so that

$$f_2(y, y) < f_2(y, x_2), \quad \forall x_2 \in \mathbb{R} \setminus \{y\}.$$

Hence, all the points on the line $x_1 = x_2$ in $X_1 \times X_2 = \mathbb{R}^2$ are strict local Nash equilibria—in fact, they are strict global Nash equilibria. ■

As the example shows, continuous games can exhibit a continuum of equilibria.

The following proposition provides first– and second–order necessary conditions for local Nash equilibria. We remark that these conditions are reminiscent of those seen in nonlinear programming for optimality of critical points.

Proposition 1 (Necessary Optimality Conditions). *Let x^* be an unconstrained local minimum of $f : \mathbb{R}^m \rightarrow \mathbb{R}$ and assume that f is continuously differentiable in an open set S containing x^* . Then*

$$Df(x^*) = 0$$

In addition, if f is sufficiently smooth (twice differentiable in this case on S) then

$$D^2 f(x^*) \geq 0$$

Proposition 2 (Necessary Conditions [14]). *If $x = (x_1, \dots, x_n)$ is a Nash equilibrium, then $\omega(x) = 0$ and $Df_i^2 f_i(x)$ is positive semi-definite for each $i \in \{1, \dots, n\}$.*

Note: This result extends to local Nash equilibria on Manifolds [14, 16, 15].

Theorem 4 ([14]). *A differential Nash equilibrium is a strict Nash equilibrium.*

Example 2.14 (Betty–Sue: Continuum of Differential Nash.) Returning to the Betty–Sue example, at all the points such that $x_1 = x_2$, $\omega(x_1, x_2) = 0$ and $D_i^2 f_i(x_1, x_2) = 1 > 0$ for each $i \in \{1, 2\}$. Hence, there is a continuum of differential Nash equilibria. ■

The following object is known as the ‘game Jacobian’ in the literature though it has been referred to also as the ‘game Hessian’ since it is composed of the components of the Hessian’s of the individual agents cost functions (i.e., the diagonal blocks) and the Jacobian of the vector field $\omega(x) = (D_1 f_1(x), \dots, D_n f_n(x))$.

Definition 26 (Game Jacobian, Ratliff, Burden Sastry 2013). *The game Jacobian is given by*

$$J(x) = \begin{bmatrix} D_1^2 f_1(x) & \cdots & D_{1n} f_1(x) \\ \vdots & \ddots & \vdots \\ D_{n1} f_n(x) & \cdots & D_n^2 f_n(x) \end{bmatrix}$$

Theorem 5 ([14]). *If $x = (x_1, \dots, x_n)$ is a differential Nash equilibrium and $J(x)$ is non-degenerate (i.e., $\det(J(x)) \neq 0$), then x is an isolated strict local Nash equilibrium.*

proof sketch. Application of the inverse function theorem to ω . Inverse function theorem says that if a function is continuously differentiable and has nonzero derivative at a point then that function is invertible in a neighborhood of that point (in particular, it is a diffeomorphism) and thus, in our case $\omega(x) = 0$ means that only x could map to zero so that it has to be an isolated Nash equilibrium.

Definition 27 ([14]). *Differential Nash equilibria $x = (x_1, \dots, x_n)$ such that $J(x)$ is non-degenerate are termed non-degenerate.*

Example 2.15 (Betty–Sue: Degeneracy and Breaking Symmetry.) Return again to the Betty–Sue example in which we showed that there is a continuum of differential Nash equilibria. At each of the points $x_1 = x_2$, it is straightforward to check that $\det(J(x_1, x_2)) = 0$. Hence, all of the equilibria are *degenerate*.

By breaking the symmetry in the game, we can make $(0, 0)$ a non-degenerate differential Nash equilibrium; i.e. we can remove all but one of the equilibria.

Indeed, let Betty’s cost be given by $\tilde{f}_1(x_1, x_2) = \frac{1}{2}x_1^2 - ax_1x_2$ and let Sue’s cost remain unchanged. Then the derivative of the differential game form $\tilde{\omega}$ of the game (\tilde{f}_1, f_2) is

$$J = D\tilde{\omega}(x_1, x_2) = \begin{bmatrix} 1 & -a \\ -1 & 1 \end{bmatrix} \quad (2.9)$$

Thus for any value of $a \neq 1$, $J = D\tilde{\omega}$ is invertible and hence $(0, 0)$ is a non-degenerate differential Nash equilibrium. This shows that small modeling errors can remove degenerate differential Nash equilibria. ■

Remark 13. *(Why do we care?) In a neighborhood of a nondegenerate differential Nash equilibrium there are no other Nash equilibria. This property is desirable particularly in applications where a central planner is designing incentives to induce a socially optimal or otherwise desirable equilibrium that optimizes the central planners cost; if the desired equilibrium resides on a continuum of equilibria, then due to measurement noise or myopic play, agents may be induced to play a nearby equilibrium that is suboptimal for the central planner.*

Theorem 6 ([15]). *Non-degenerate differential Nash are structurally stable (robust to small perturbations) and generic (meaning that Nash equilibria in an open dense set of continuous games are non-degenerate differential Nash).*

It is worth mentioning a special class of continuous games known as *potential games* which have the nice property that a transformation of coordinates transforms the game to an optimization problem.

Example 2.16 (Couple Oscillators: Potential Games with a Continuum of Nash.)

Consider n -coupled oscillators with an interaction structure specified by a undirected, complete graph where the nodes represent the oscillators and the edges indicate coupling between oscillators. Such a game can be regarded as an abstraction of generators or inverters—perhaps even microgrids—coupling to the grid where each oscillator is individually managed. Another practical example is decentralized and autonomous synchronization of timing cycles for traffic lights based on real-time traffic flow.

Let the phase of oscillator i be denoted by $\theta_i \in \mathbb{S}^1$ and let its cost be

$$f_i = -\frac{1}{n} \sum_{j \in N_i} \cos(\theta_i - \theta_j)$$

where N_i is the index set of oscillators that are coupled to oscillator i .

The form of the cost is derived from the Laplacian potential function [12]. It is straightforward to check that the differential game form for the oscillator game satisfies

$$D_i f_i \equiv D_i \phi$$

(i.e., differential game form is such that $\omega = d\phi$) where

$$\phi(\theta_1, \dots, \theta_n) = -\frac{1}{2n} \sum_{i=1}^n \left(\sum_{j \in N_i} \cos(\theta_i - \theta_j) \right).$$

Potential games have nice properties in terms of the existence of equilibria and convergence of learning algorithms. However, in this example we see that even such *nice* games potentially have degeneracies and thus, it is important to be able to characterize when they arise.

We claim that all points in the set

$$\{(\theta_1, \dots, \theta_n) \in \mathbb{S}^1 \times \dots \times \mathbb{S}^1 \mid \theta_i = \theta_j, \forall i, j \in \{1, \dots, n\}\}$$

are global Nash equilibria of the game. Indeed, fix $\theta_j = \beta$ for all $j \neq i$. Then $\theta_i = \beta$ is a best-response—i.e. in the set of optimizers—by oscillator i to all other oscillators playing $\theta_j = \beta$. In particular, for $\theta'_i \in \mathbb{S}^1 \setminus \{\beta\}$,

$$\phi(\beta_1, \dots, \beta_n) = -\frac{|N_i|}{n} < -\frac{1}{n} \sum_{j \in N_i} \cos(\theta'_i - \beta), \quad \forall \theta'_i \in \mathbb{S}^1 \setminus \{\beta\},$$

where $|N_i|$ denotes the cardinality of the set N_i . Thus there is a continuum of Nash equilibria for which the oscillators are *synchronized*. In fact there is a continuum of differential Nash equilibria; this is easily seen by checking that $D_i^2 f_i(\theta, \dots, \theta) > 0$.

Consider now a simple game with $n = 2$ oscillators managed by Jean and Paul respectively. Let Jean's cost be $f_1 = -\frac{1}{2} \cos(\gamma\theta_1 - \theta_2)$ and Paul's cost be $f_2 = -\frac{1}{2} \cos(\theta_2 - \theta_1)$ where in this

example Jean and Paul have different preferences for their phase. Allowing γ to take values in $\mathbb{N} \setminus \{1\}$, there are at least $\gamma - 1$ non-degenerate differential Nash equilibria:

$$\left\{ (\theta_1, \theta_2) \in \mathbb{S}^1 \times \mathbb{S}^1 \mid \theta_1 = \theta_2 = \frac{2(n-1)\pi}{\gamma-1}, n \in \{1, \dots, \gamma-1\} \right\}.$$

The above set contains only stable, non-degenerate differential Nash equilibria of the game (f_1, f_2) since points in this set satisfy $\omega(\theta_1, \theta_2) = 0$, $D_i^2 f_i(\theta_1, \theta_2) > 0$, and $\det(d\omega(\theta_1, \theta_2)) \neq 0$. In fact, they are (non-strict) global Nash equilibria. Due to the existence of a continuum of Nash equilibria, it is possible that the players will equilibrate on a socially undesirable outcome. A central planner vying to coordinate the individuals would therefore benefit from considering these second-order conditions when designing incentives.

2.8.1 Computation of Nash in Continuous Games

We can compute stable non-degenerate differential Nash!

Definition 28. *A differential Nash equilibrium is stable if the eigenvalues of J lie strictly in the open right-half plane (resp. left half plane for maximizers).*

Consider the dynamics

$$\dot{x} = -\omega(x) = - \begin{bmatrix} D_1 f_1(x) \\ \vdots \\ -D_n f_n(x) \end{bmatrix}$$

Proposition 3 ([16]). *If x is a differential Nash equilibrium and the spectrum of J is in the right-half plane, then x is an exponentially stable stationary point of the above dynamical system.*

Clearly from this proposition, there is a range of finite step-sizes for which convergence of a forward Euler discretization converges exponentially.

Indeed, consider

$$x_{k+1} = x_k - \gamma \omega(x_k)$$

such that $\gamma < 1/L$ where ω is L -Lipschitz. We know this update rule will converge exponentially to stable attractors of $\dot{x} = -\omega(x)$. However, not all stable attractors are meaningful in a game theoretic sense. For instance, it is easy to construct examples such that x^* is a stable attractor for $\dot{x} = -\omega(x)$, but not a Nash equilibrium; consider for example any x^* such that $\omega(x^*) = 0$ and $\text{Re}(-J(x^*)) \subset \mathbb{C}_-^\circ$, but any of the $D_i^2 f_i(x^*)$ are indefinite.

This fact has been the subject of numerous papers in AI/ML of late and has led to a re-examination/revival of 'learning in games' via a dynamical systems perspective from which we can gain numerous insights not only about the equilibrium behavior but also the learning path. New directions of interesting research include:

1. interpreting the learning path and the behavior along the learning path (e.g., decomposing into cooperative versus non-cooperative updates and adaptations)
2. understanding (qualitatively and quantitatively) the effects of heterogeneity on the learning behavior and achievable outcomes

I will expand this section

2.8.2 Continuous Stackelberg Games

I will expand this section

2.9 Differential and Dynamic Games

Up to this point, we have discussed games in which agents interact in a stationary, stochastic environment, and they are making one shot decisions or interacting repeatedly by playing the same game over and over. However, in practice it is often the case that there is an auxiliary stochastic process that describes the evolution of an environmental “state” variable, itself which may be impacted by the actions of the agents. Moreover, agents’ decision problems may be such that they are choosing a sequence of actions of a finite (or infinite) time horizon. Such games are referred to as dynamic games.

In short, in dynamic games, actions available to each agent depends on their current state which evolves according to a certain dynamical system. Sets of states/actions are usually a continuum.

I am currently working on this section.

Chapter 3

Theory of Contracts and Incentives

Contract theory is concerned with the implementation of social choice rules in situations where these cannot be made part of a contract due to the presence of incomplete information

- (i.e. either non-observability and/or non-verifiability of performance)

The goal is to examine whether the social choice rule in question can be implemented indirectly or replicated through either

1. an alternative rule that is enforceable by courts or
2. some institutional arrangement.

If the social choice rule can be fully replicated, then we speak of an *efficient* or *first-best* solution. In general, not possible due to constraints imposed by the information structure. So this is when we look for an enforceable alternative or institutional arrangement. Such rules or institutions are called *constrained efficient* or *second-best* optimal.

3.1 Social Choice Rules

Consider the following model. There are n agents indexed by $[n] = \{1, \dots, n\}$ and there is a finite set A of feasible outcomes. Each agent has a characteristic or *type* $\theta_i \in \Theta_i$.

Definition 29. *State.* A state is a profile of types $\theta = (\theta_1, \dots, \theta_n) \in \Theta$ which defines a profile of preference orderings $\succeq(\theta) = (\succeq_1(\theta), \dots, \succeq_n(\theta))$ on the set of feasible outcomes A .

We use the notation $\succeq_i(\theta)$ to denote agent i 's preference ordering on A in state θ and we denote by $\mathcal{R} = \times_{i \in [n]} \mathcal{R}_i$ the set of all possible preference orders for all agents. Each agent (but no outside party) observes the entire vector $\theta = (\theta_1, \dots, \theta_n)$ —i.e. agents have *complete information*.

Informally, a *social choice rule* is a selection rule that determines a set of socially desirable outcomes for each state $\theta \in \Theta$.

Definition 30. *Social Choice Rule.* A social choice rule is a correspondence $f : \Theta \rightarrow A$ (surjection) which specifies a non-empty choice set $f(\theta) \subseteq A$ for every state θ .

Agent i 's preference in state θ depends only on their type θ_i —i.e. $\succeq_i(\theta) = \succeq_i(\theta_i)$. The following are examples of social choice rules:

1. Paretian social choice rule: comprises only Pareto optimal allocations
2. Dictatorial social choice rule: for all θ , the social choice set $f(\theta)$ is a subset of the most preferred outcomes of a particular agent.

Example 3.1 (Provision of Public Good.) The city council (*the social planner*) considers the construction of a road for the n inhabitants of the city. In this example, θ_i represents agent i 's valuation or willingness to pay for the road. An outcome is a profile $y = (x, t_1, \dots, t_n)$, where x can be either '1' ("the road is built") or '0' ("the road is not built"), and where t_i denotes a transfer to agent i —note that transfers can be negative, like taxes. Preferences are assumed to be quasilinear of the form $\theta_i x + t_i$. The city council faces the restriction that it cannot provide additional funds—i.e. the cost $c \geq 0$ must be covered entirely by the inhabitants. This defines the set of feasible outcomes as

$$A = \{(x, t_1, \dots, t_n) \mid x = \{0, 1\} \text{ and } \sum_i t_i \leq -cx\}.$$

A particularly desirable social choice rule is one where the road is built if and only if the sum of the agent's valuations exceeds the construction cost and where the budget constraint is satisfied with equality—i.e.

$$\sum_i t_i = -cx \quad \text{and} \quad x(\theta) = \begin{cases} 1, & \text{if } \sum_i \theta_i \geq cx \\ 0, & \text{o.w.} \end{cases}$$

The set of outcomes defined by the above equations coincides with the set of Pareto optimal allocations (with respect to both the public good and money). Unfortunately, it turns out that this social choice rule is not implementable in dominant strategies.

3.2 Mechanism Design

A mechanism design problem is one in which an uninformed agent (the planner) faces a group of informed agents. The agents' private information concerns the state θ , which determines their preferences over A , the set of feasible outcomes. Clearly, the planner cannot directly implement the social choice rule since she does not know the true state. Moreover, when the agents are asked to reveal their preferences honestly, each individual agent has an incentive to misrepresent his information. This *unobservability* results in a problem of *adverse selection* and the incentive to misrepresent information is a problem of *moral hazard*. The planner can fix (or at least improve the outcome) by designing a *mechanism* that uses publicly observable (and thus verifiable) information.

Formally, a mechanism Γ consists of a collection of strategy sets $\Sigma = \{S_1, \dots, S_n\}$ and an outcome function $g : S \rightarrow A$, where $S = S_1 \times \dots \times S_n$, which assigns an outcome $y \in A$ to each strategy profile $s = (s_1, \dots, s_n) \in S$. Since the state θ is not verifiable, the outcomes themselves cannot be made contingent on the state. However, the agents' payoffs (utilities) from a particular outcome typically vary with θ since preferences $\geq_i(\theta_i)$ are state-dependent. Thus, the mechanism Γ combined with the state-space θ defines a game of complete information with a (possibly) different payoff structure in every state θ . The implementation problem is then to construct Γ such that in each state, the equilibrium outcomes of the resulting game coincide (in a way yet to be defined) with the elements in $f(\theta)$.

Remark 14. *The idea which underlies mechanism design is information revelation through strategy choice.*

Since the strategy sets $\Sigma = \{S_1, \dots, S_n\}$ and the outcome function $g : S \rightarrow A$ are *public information*, outsiders such as courts (or the planner) can compute the agents' equilibrium strategy rules $s^*(\theta) = (s_1^*(\theta), \dots, s_n^*(\theta))$. Denote by $E_g(\theta)$ the set of equilibrium profiles $s^*(\theta)$ of Γ in state θ , and define the set of equilibrium outcomes of Γ in θ as

$$g(E_g(\theta)) \equiv \{g(s^*(\theta)) \mid s^*(\theta) \in E_g(\theta)\}.$$

Definition 31. *Implementation.* *The mechanism Γ implements the social choice rule $f(\theta)$ if $E_g(\theta)$ is non-empty, and if for every $\theta \in \Theta$, $g(E_g(\theta)) \subset f(\theta)$.*

Definition 32. *Full Implementation.* *The mechanism Γ fully implements the social choice rule $f(\theta)$ if for every $\theta \in \Theta$, $g(E_g(\theta)) = f(\theta)$.*

3.2.1 The Revelation Principle and Truthful Implementation

The identification of all social choice rules that are implementable for a specific equilibrium concept requires knowledge of the entire set of possible mechanisms. Fortunately, a very useful result known

as the revelation principle allows us to restrict attention to a particularly simple class of mechanisms called direct mechanisms.

Definition 33. *Direct Mechanism.* A direct mechanism Γ_d is a mechanism in which $S_i = \theta_i$.

Definition 34. *Truthful Implementation.* The direct mechanism Γ_d truthfully implements the social choice rule $f(\theta)$ if for every $\theta \in \Theta$, $\theta \in E_g(\theta)$ and $g(\theta) \in f(\theta)$.

Remark 15. *Truthful implementation is a weaker concept than implementation or full implementation. Indeed, truthful implementation only requires that the profile $\theta = (\theta_1, \dots, \theta_n)$ of truthful announcements is an equilibrium in each state θ and that the equilibrium outcome $g(\theta)$ is an element in $f(\theta)$. However, truthful implementation does not rule out the existence of further equilibria with outcomes $g(\hat{\theta}) \in f(\theta)$ in which some agents lie (i.e. $\hat{\theta} \neq \theta$).*

Both implementation and full implementation rule out such “unwanted” equilibria as they require either that the set of equilibrium outcomes constitutes a subset of $f(\theta)$ (implementation) or that it coincides with $f(\theta)$ (full implementation).

The great virtue of dominant strategy equilibrium is that agents need not forecast how other agents choose their strategies. In other words, agents do not have to know each others preferences. This is the basis for an extremely convenient result known as the revelation principle **Change to citations: (Gibbard (1973), Green and Laffont (1977), Dasgupta, Hammond, and Maskin (1979))**, which says that we can restrict attention to direct mechanisms in which agents report only their own types. Thus, the assumption of complete information made at the beginning of this chapter is irrelevant and any result derived in this section continues to hold if this assumption is dropped. Due to this robustness property, social choice rules that are truthfully implementable in dominant strategies are of particular interest.

Definition 35. *Truthfully Implementable in Dominant Strategies.* The social choice rule $f(\theta)$ is truthfully implementable in dominant strategies or strategy-proof if there exists a direct mechanism Γ_d such that truthtelling is a dominant strategy equilibrium—i.e. if for all $i \in I$ and $\theta_i \in \Theta_i$,

$$g(\theta_i, \hat{\theta}_{-i} \succeq_i (\theta_i)g(\hat{\theta}_i, \hat{\theta}_{-i}) \quad \forall \hat{\theta}_i \in \Theta_i, \hat{\theta}_{-i} \in \Theta_{-i}$$

and $g(\theta) \in f(\theta)$ for all $\theta \in \Theta$.

Theorem 1. *Revelation Principle.* If an social choice rule is implementable in dominant strategies, then it is truthfully implementable in dominant strategies.

That is, for every mechanism Γ that implements $f(\theta)$ in dominant strategies, we can find a direct mechanism Γ_d that truthfully implements $f(\theta)$ in dominant strategies.

Definition 36. *Dictatorial Social Choice Rule.* The social choice rule $f(\theta)$ is dictatorial on the set $A' \subseteq A$ if there exists an agent $i \in I$ such that for all $\theta \in \Theta$, the choice set $f(\theta)$ is a subset of agent i 's most preferred outcomes A' , i.e., $f(\theta) \subseteq \{y \in A' \mid y \succeq_i (\theta_i)z \ \forall z \in A'\}$

Theorem 2. *Gibbard-Satterthwaite Theorem.* Let the social choice rule $f(\theta)$ be single-valued and let $A' \subseteq A$ denote the range of $f(\theta)$. Suppose that A is finite, that A' contains at least three elements, and that for each agent $i \in I$, the set of possible preference orderings \mathcal{R}_i is the set of strict preference orderings on A . Then $f(\theta)$ is truthfully implementable in dominant strategies if and only if it is dictatorial on A' .

Definition 37. *Budget-Balanced Social Choice Rule.* An social choice rule is budget-balanced if $\sum_i t_i(\theta) = 0$ for all $\theta \in \Theta$.

That is, the sum of the transfers must net to zero. Since the planners preferences do not enter into our welfare considerations, any net surplus $|\sum_i t_i(\theta)| > 0$ collected from the agents is wasteful.

3.3 Incentive Design

Building on the concepts introduced in the previous section, in this section we focus our attention to the design of incentive mechanisms. A more comprehensive presentation can be found in text books such as [10] and [2]; for a control theoretic perspective, see [17];

3.3.1 The Principal-Agent Problem

Incentive design and contract theory problems are often cast as so-called *principal-agent* problems in which there are notably two types of participants—*principal* and *agent*. The principal’s goal is to design a mechanism to shape the decision of the agent who acts in their own interest.

The principal has utility $J_p : U \times Y \rightarrow \mathbb{R}$ and the agent has utility $J_a : U \times Y \rightarrow \mathbb{R}$ where U and V are the action spaces of the agent and principal, respectively. The agent’s and principal’s utilities are coupled and, thus there is a game between the principal and agent. However, there is a specific order of play. In this game, the principal announces a mapping of the agent’s action space into the principal’s action space, after which an agent selects its optimal action. Let

$$\Gamma = \{\gamma : U \rightarrow Y\} \quad (3.1)$$

be the admissible set of such mappings from which the principal can choose. The mappings in Γ may have a particular structure such as continuity or monotonicity—e.g., Γ may be the space of continuous linear maps. Moreover, the structure of maps belonging to Γ may be practically motivated—e.g., tariff structures imposed by a regulatory body.

Given γ , the agent aims to selection an action that maximizes their utility—i.e.

$$u^*(\gamma) \in \arg \max_{u \in U} J_a(u, \gamma(u)) \quad (3.2)$$

where we denote the dependence of u^* on γ .

The principal aims to design $\gamma \in \Gamma$ such that the agent selects an action $u^*(\gamma)$ that leads to the maximization of the principal’s utility—that is, the principal wants to *incentivize* the agent to play according to what is ‘best’ for the principal. In this way, γ realigns the preferences of the agent with those of the principal.

Definition 38. *Principal-Agent Problem.* Given a principal with utility $J_p : U \times Y \rightarrow \mathbb{R}$ and an agent with utility $J_a : U \times Y \rightarrow \mathbb{R}$, the principal aims to select a mapping $\gamma \in \mathcal{M} \subset \Gamma$ where

$$\mathcal{M} = \{\gamma \in \Gamma \mid u^d = u^*(\gamma) \in \arg \max J_a(x, \gamma(x)), \gamma(x^d) = y^d, (x^d, y^d) \in \arg \max_{x,y} J_p(x, y)\}. \quad (3.3)$$

There may be more than one agent in which case there needs to be a characterization for their interaction under the mapping γ .

Consider a scenario in which there is one principal and n agents, each with utilities $J_i : U \times Y \rightarrow \mathbb{R}$ for $i \in [n] = \{1, \dots, n\}$ and where $U = U_1 \times \dots \times U_n$. We use the notation $u = (u_1, \dots, u_n) \in U$ to denote the vector of their choices such that $u_i \in U_i$ for each $i \in [n]$. For a mapping $\gamma \in \Gamma$, agent i ’s *induced* utility under γ is

$$J_i^\gamma(u) = J_i(u, \gamma(u)). \quad (3.4)$$

As before, the principal’s goal is to choose $\gamma = (\gamma_1, \dots, \gamma_n) \in \Gamma$, where $\gamma_i : U_1 \times \dots \times U_n \rightarrow V$, such that the agents are induced to play a vector of choices u^d that maximize the principal’s utility where u^d is characterized by a Nash equilibrium.

Definition 39. *Nash Equilibrium of the γ -Induced Game.* For the game $(J_1^\gamma, \dots, J_n^\gamma)$ induced by the mapping $\gamma = (\gamma_1, \dots, \gamma_n) \in \Gamma$, a point $u^* = (u_1^*, \dots, u_n^*) \in U$ is said to be a Nash equilibrium if, for each $i \in [n]$,

$$J_i(u_i^*, u_{-i}^*, \gamma(u^*)) \geq J_i(u_i, u_{-i}^*, \gamma(u_i, u_{-i}^*)), \quad \forall u_i \in U_i. \quad (3.5)$$

Let $\text{NE}(J_1^\gamma, \dots, J_n^\gamma)$ denote the set of Nash equilibria for $(J_1^\gamma, \dots, J_n^\gamma)$. Then, in the multi-agent case, the principal chooses $\gamma \in \mathcal{M}$ where

$$u^*(\gamma) \in \text{NE}(J_1^\gamma, \dots, J_n^\gamma). \quad (3.6)$$

While there is a misalignment of objectives between the principal and the agent, if a feasible solution exists for \mathcal{M} , then both the principal and agent(s) are doing what is in their best interest: the agent is ‘compensated’ via γ to play x^d and $\gamma(x^d) = y^d$ ensuring the principal’s utility is maximized.

3.3.2 Information Asymmetries

Issues of incentives truly arise when there are *information asymmetries* between the principal and the agent. That is, in reality the principal and the agent make their decisions based on some information set that is available to them.

Let \mathcal{I}_i be the information set of agent i and \mathcal{I}_P be the information set of the principal. In the full information case, as in the basic problem described above, $\mathcal{I}_P = \{(J_i)_{i \in [n]}, u\}$ and $\mathcal{I}_i = \{J_P, (J_i)_{i \in [n]/\{i\}}, u_{-i}\}$ where $u_{-i} = (u_j)_{j \in [n]/\{i\}}$. It is the contents of these information set that determine the particular challenges that arise and, in turn, often shape the approach to the problem.

LJR: I am still working on this section.

Definition 40. *Adverse Selection.* Consider a principal-agent problem with a single agent possessing utility J_a which belongs to a class of functions $\mathcal{F}(\theta)$ parameterized by $\theta \in \Theta \subset \mathbb{R}^d$ —i.e. $J_a \in \mathcal{F}(\theta)$.

Example 3.2 (Market for Lemons.)

Definition 41. *Moral Hazard.* On the other hand, if the utility J_a is known to the principal, but the action u is not—i.e. $\mathcal{I}_P = \{J_a\}$ —then this leads to a problem known as moral hazard.

Example 3.3 (Insurance.)

Bibliography

- [1] Robert J Aumann. Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics*, 1(1):67–96, 1974.
- [2] Patrick Bolton and Mathias Dewatripont. *Contract theory*. MIT Press, 2005.
- [3] Partha Dasgupta and Eric Maskin. The existence of equilibrium in discontinuous economic games, ii: Applications. *The Review of economic studies*, 53(1):27–41, 1986.
- [4] Partha Dasgupta, Eric Maskin, et al. The existence of equilibrium in discontinuous economic games, i: Theory. *Review of Economic Studies*, 53(1):1–26, 1986.
- [5] Gerard Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 38(10):886–893, 1952.
- [6] Ky Fan. Fixed-point and minimax theorems in locally convex topological linear spaces. *Proceedings of the National Academy of Sciences of the United States of America*, 38(2):121, 1952.
- [7] Drew Fudenberg, Fudenberg Drew, David K Levine, and David K Levine. *The theory of learning in games*, volume 2. MIT press, 1998.
- [8] Jugal Garg, Albert Xin Jiang, and Ruta Mehta. Bilinear games: Polynomial time algorithms for rank based subclasses. In *International Workshop on Internet and Network Economics*, pages 399–407. Springer, 2011.
- [9] Irving L Glicksberg. A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.
- [10] Jean-Jacques Laffont and David Martimort. *The Theory of Incentives: The Principal-Agent Model*. Princeton university press, 2002.
- [11] R Duncan Luce. *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2012.
- [12] Derek A Paley, Naomi Ehrich Leonard, Rodolphe Sepulchre, Daniel Grunbaum, and Julia K Parrish. Oscillator models and collective motion. *IEEE Control Systems Magazine*, 27(4):89–105, 2007.
- [13] Pablo A Parrilo. Polynomial games and sum of squares optimization. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 2855–2860. IEEE, 2006.
- [14] Lillian J Ratliff, Samuel A Burden, and S Shankar Sastry. Characterization and computation of local nash equilibria in continuous games. In *Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 917–924. IEEE, 2013.

- [15] Lillian J Ratliff, Samuel A Burden, and S Shankar Sastry. Genericity and structural stability of non-degenerate differential nash equilibria. In *Proceedings of the 2014 American Control Conference*, pages 3990–3995. IEEE, 2014.
- [16] Lillian J Ratliff, Samuel A Burden, and S Shankar Sastry. On the characterization of local nash equilibria in continuous games. *IEEE Transactions on Automatic Control*, 61(8):2301–2307, 2016.
- [17] Thomas Weber. *Optimal Control Theory with Applications in Economics*. MIT Press, 2011.