Games

Bruno Salcedo*

Winter 2020

So far in the course, we have analyzed allocation problems taking the market institution as given. For the remainder of the course, we will shift our focus drastically. First, we will go beyond allocation problems and consider instead a much more general class of choice problems. Second, we will not take any specific institution as given. Instead, we will try to figure out which institutions are better suited to achieve different social goals. This shift of focus requires a new set of tools. *Game Theory* provides a simple unified framework that allows to analyze a wide class of social problems.

1. Decisions

1.1. Expected utility

Decision Makers (DMs) often lack all the information required to determine which alternative is better. In such cases, a rational DM would have to consider different possible scenarios, as well as their respective likelihoods. Think for instance of Anna, who needs to choose which boots to wear today. She has to choose between wearing her rain boots, or her suede boots.

If it rained, Anna would rather use her rain boots, because otherwise her suede boots could suffer water damage. If did not rain, she would prefer using

*Department of Economics, Western University · brunosalcedo.com · bsalcedo@uwo.ca
Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported license.

	no rain [70%]	rain [30%]
suede boots	4	-2
rain boots	2	2

Figure 1 – Anna's utility function

her suede boots, which are more comfortable. The weather app informs Anna that the probability of rain is 30%. How should she proceed? Her decision should probably take different factors into account, including: the difference in comfort between both pairs of boots, the cost of the potential water damage, and the likelihood of rain. One way to model Anna's decision is using an *expected utility* (EU) framework, consisting of four components:

(*i*) A set A of *alternatives* available for the DM to choose from;

(*ii*) A set X of different *scenarios* that the DM considers possible;

(*iii*) A *probability function* $\Pr(\,\cdot\,)$ that specifies the likelihood of each scenario;

(*iv*) A *utility function* $u(\,\cdot\,,\,\cdot\,)$ over alternative-scenario pairs.

In the context of our example, the alternatives available for Anna are the rain boots and the suede boots. The possible scenarios are that it rains, and that it does not rain. The probability that it rains is 30%, and the probability that it does not rain is 70%. Anna's utility function could be the one in Figure 1. The difference between 4 and 2 in the first column represents the difference in comfort between the suede and rain boots. The difference between 4 and -2 on the first row represents the magnitude of the water-damage cost.

Given an EU framework, the *expected utility* from choosing alternative a is the mathematical expectation of the DM's utility $u(a, \tilde{x})$, treating the true scenario $\tilde{x}$ as a random variable. It equals the weighted average of the utility that can result from choosing a, using the likelihood of the possible scenarios as weights, that is,

$$\mathrm{EU}(a) = \mathbb{E}[\,u(a, \tilde{x})\,] = \sum_x \Pr(x) \cdot u(a, x).$$

For the remainder of the course, we will maintain the assumption that DMs make choices as to maximize their expected utility. DMs that satisfy this assumption are called *rational*.

Assumption 1 (Expected Utility Hypothesis) DMs always choose alternatives which yield the maximum possible expected utility.

The Expected Utility Hypothesis (EUH) should be understood as an "as if" hypothesis. There are behavioral patterns (similar to the axioms of revealed preference) which are equivalent to it. Also, there is plenty of evidence that the EUH does a good job at explaining behavior under some circumstances, and a very poor job under some other circumstances. Decision theorists have developed alternative models of choice for situations when the EUH is known to fail systematically. And the game theory tools that we will study can incorporate such alternative models. But these topics are beyond the scope of this class. We will simply assume that the EUH always holds.

For Anna, the expected utility from choosing to wear her suede boots is

$$\text{EU(suede boots)} = 70\% \cdot 4 + 30\% \cdot (-2) = 2.2.$$

Her expected utility from choosing the rain boots is

$$\text{EU(rain boots)} = 70\% \cdot 2 + 30\% \cdot 2 = 2.$$

Therefore, the model suggests that 30% is a low enough probability of rain, so that it is worth it for her to use her favorite boots and risk the possibility of water damage.

1.2. Dominance

In order to maximizing expected utility, one has to assign probabilities and utilities to all the possible scenarios and alternative-scenario pairs, respectively. And the optimal Alternative might depend on the precise values assigned. Consider for instance Anna's example. If the probability of rain was 40% instead of 30%, or if the utility from the rain boots was 2.5 instead of 2, then the rain boots would yield a higher expected utility. This poses a challenge, both for DMs and for economists trying to analyze the behavior of DMs, because it might be difficult to obtain precise values for probabilities and/or utilities. Fortunately, there are

	no rain [70%]	rain [30%]
suede boots	4	-2
rain boots	2	2
snow boots	4	3

Figure 2 – Extended footwear choices for Anna

some choices that are much simpler.

Suppose for instance that Anna has one more pair of boots available: a pair winter boots that are both comfortable and water resistant. Her utility function with this new alternative taken into account is specified in Figure 2. Let us analyze how the winter boots compare to each of the other two alternatives.

Note that, if it rains, the winter boots yield a higher utility than the rain boots. The same is true if it does not rain. Therefore, the winter boots are preferred to the rain boots in *all* possible scenarios. In this case, winter boots *dominate* rain boots according to the following definition.

Definition 1 Alternative a *dominates* alternative b if and only if $u(a, x) > u(b, x)$ for every scenario x.

Dominated alternatives cannot maximize expected utility. In Anna's example, the expected utility from using winter boots satisfies

$$\text{EU(winter boots)} = \Pr(\text{rain}) \cdot 4 + \Pr(\text{no rain}) \cdot 3 > 2.$$

In general, if alternative a is dominated by alternative b, then you can always be sure that b yields a higher expected utility than a (why?). This conclusion does *not* depend on the probability of different scenarios.

Claim 1 *The EUH implies that DMs never choose dominated strategies.*

How about the winter boots vs. the suede boots? The winter boots do *not* dominate the suede boots, because they both yield the same utility in case it does not rain. However, the suede boots are never strictly preferred to the winter boots, and the winter boots are strictly preferred to the suede boots in at least one possible scenario. In this case, the winter boots *weekly dominate* the suede boots according to the following definition.

Definition 2 Alternative a *weakly dominates* alternative b if and only if $u(a,x) \geq u(b,x)$ for every possible scenario x, with at least one strict inequality. An alternative is *admissible* if and only if it is *not* weakly dominated by any other alternative.

An expected utility maximizer could choose weakly dominated alternatives. In Anna's example, if the probability of rain were *exactly* zero, then the suede boots would also maximize expected utility. But even in that case, the suede boots are *not* strictly better than the winter boots. And the winter boots would be strictly better in every other case, as long as the probability of rain were positive, no matter how small it were.

A cautious DM should acknowledge the possibility that her probability assessment might be incorrect. Consequently, she should avoid weakly dominated alternatives. For the remainder of the course we will maintain the assumption that all agents are cautious and thus avoid choosing weakly dominated alternatives.

Assumption 2 (Cautiousness) DMs always choose admissible alternatives.

For Anna, Assumption 1 allows to rule out the rain boots, and Assumption 2 allows to rule out the suede boots. Hence, we can conclude that Anna would choose to wear her winter boots. Note that this conclusion does not depend on the details of the environment. One does not need to know the probability of each scenario to determine whether an alternative is dominated or not. Also, one does not need to know the specific utility values for each alternative-scenario pair. The only necessary information are the *ordinal rankings* of the alternatives conditional on each possible scenario. This means that a DM or an economist can rule out dominated alternatives, without having to assign probabilities or precise utility values.

2. Games

2.1. Simultaneous-move games

Now, let us turn our attention to a more general class of environments. Suppose that there is more than one DM. Each DM has to make a single choice. And the different DMs make their choices simultaneously and independently of one another. Such situations can be modeled as *simultaneous-move games*. Formally, a simultaneous-move game is a mathematical object consisting of three components:

(i) A set I of DMs. Each DM is called a *player* of the game.

(ii) For each player i, a set of alternatives A_i to choose from. Each alternative is called an *action*. An *action profile* is a list specifying one action for each player.

(iii) For each player i, a utility function u_i that represents i's preferences over action profiles.

For example, suppose that Bob and Charlie met for the first time today. They would like to see each other again, but they forgot to exchange contact information. They also don't have social media, nor any friends in common that they know of. There are two parties happening tonight, one is located West of campus, and the other one is East of campus. Each one of Bob and Charlie will decide which party to go to. Since they have no way of contacting each other, they will make their choices independently.

Bob thinks that the West party would be more fun, but would rather go to the East party if he knew that Charlie would be there. The converse is true for Charlie. Their utility functions are given in the matrix in Figure 3. Each cell corresponds to an action profile. The first number in each cell corresponds to the utility of the row player (Bob), and the second number is the utility of the column player (Charlie). How would you choose which party to go to in this situation?

Matrices like the one in Figure 3 provide a succinct way to depict *finite two-player* games. Such matrices are not useful when there are many players, or when some players have an infinite amount of actions to choose from. In such cases, the game must be described verbally or using mathematical equations.

		Charlie	
		West	East
Bob	West	$5, 2$	$1, 1$
	East	$0, 0$	$2, 5$

Figure 3 – Meeting at a party

For an example of an *infinite* game, consider a market operated by two profit-maximizing firms producing an homogeneous commodity. Firms indexed by $j \in \{1, 2\}$. Each firm j chooses a non-negative quantity q_j to supply to the market. The product is divisible and production capacity is capped at 100 units. All firms have the same cost function $C(q_j) = 1 + q_j^2$. The market price depends on the total quantity supplied to the market $q = q_1 + q_2$. It is given by the inverse demand function $p(q) = 200 - q$.

It is not possible to represent this game using a matrix like the one in Figure 3, but we still have a well defined simultaneous-move game. The players are the two firms. Both players have the same set of actions, namely, $[0, 100]$. The utility function for Firm 1 is just its profit function

$$u_1(q) = q_1 p(q) - C(q_1) = -2q_1^2 + (200 - q_2)q_1 - 1.$$

The utility function for Firm 2 is analogous.

Before analyzing behavior in games, it is convenient to introduce some notation. When speaking about an unspecified player i, the symbol "$-i$" denotes i's opponents. For example, if i represented Bob, then $-i$ would represent Charlie. Similarly, if i represented Charlie, then $-i$ would represent Bob. It is also convenient to split an action profile using similar notation. We can write an action profile as (a_i, a_{-i}), where a_i represents i's action, and a_{-i} the actions corresponding to i's opponents. Awesome possum!!

2.2. Dominance in games

Note that a simultaneous-move is nothing more than a collection of simultaneous EU models like the ones analyzed in Sections 1. For instance, Bob faces a decision problem in which he has to choose which party to attend. Since he does not know which party Charlie is going to, he needs to consider two possible scenarios and their respective likelihood. This suggests that we can use the notions of EUH and cautiousness to model behavior in games. Can we make informative predictions if we assume that players are cautious? The answer depends on the specific game being considered.

Consider a situation in which an object is going to be auctioned. The format of the auction is what is called a *sealed-bid second-price* auction format. Each of the potential buyers will write down a bid in a piece of paper and place it inside a sealed envelope. The auctioneer will collect all the envelopes and then will announce the results. The person who submitted the highest bid takes the object, but she does not pay her bid. Instead, she pays the *second highest* bid that was submitted.

Let us model this situation as a game. The players are the potential buyers. The actions available to each player are the different bids she could submit. Bids should be any nonnegative real number. Let us assume that all potential buyers have quasilinear preferences, and let v_i denote i's willingness to pay for the object. Then, the utility for player i is given by

$$u_i(a) = \begin{cases} v_i - t_i & \text{if } a_i \text{ is the highest bid} \\ 0 & \text{otherwise} \end{cases},$$

where $t_i = \max\{a_j \mid j \neq i\}$ denotes the highest bid among i's competitors.

Claim 2 *In a sealed-bid second-price auction, cautious players will always bid exactly their value.*

Justification. To see verify the claim, fix some player i, and let us compare the utility from bidding v_i vs. the utility from bidding some number $b_i > v_i$. There are three cases to consider, depending on the value of t_i.

Case 1— If $t_i < v_i$, then i would win the auction regardless of whether she bid b_i or v_i. The price paid by i would also be the same in either case.

Therefore, i would be indifferent between bidding b_i and v_i.

Case 2— If $t_i > b_i$, then i would lose the auction regardless of whether she bid b_i or v_i, and her utility would be 0 in either case.

Case 3— The only remaining possibility is to have $t_i > v_i$ and $t_i < b_i$. In this case, if i were to bid v_i, she would lose the auction and her utility would be 0. If she were to bid b_i she would win the object and her utility would be $v_i - t_i < 0$. Therefore, she would strictly prefer to bid v_i.

To summarize, overbidding increases a bidder's chance to win the object, but only in cases in which she would would have to pay more than what the object is worth to her. It follows that overbidding is always weakly dominated. An analogous argument can be used to show that underbidding is also weakly dominated. ■

Assuming cautiousness perfectly pins down the outcome in a sealed-bid second-price auction. But this is not true for all games. Consider for instance a sealed-bid *first-price* auction. The format is similar to the second price auction, with the difference that the person who submits the highest bid has to pay her own bid and not the second highest one. As before, overbidding is weakly dominated by bidding truthfully. This is because, if a person overbids and wins a first-price auction, then she will have to pay more than the object is worth to her.

In contrast, underbidding is not weakly dominated in a first price auction. Suppose that bidder i's value for the object is 10, and she believes that the highest bid among her opponents will be equal to 5. Bidding exactly 10 would not be optimal, because she would end up with a utility of 0. She would be better off bidding a small amount above 5. That way she would still win the object but she would pay a smaller price for it.

Claim 3 *In sealed-bid first-price auctions, cautiousness rules out overbidding but it does* not *rule out underbidding.*

In some cases, cautiousness has no predictive power at all. Let us revisit the game from Figure 3. Let π denote the probability that Charlie goes to the East party. Bob's expected utilities for each of his available actions are given by

$$\text{EU}_{\text{Bob}}(\text{West}) = (1 - \pi) \cdot 5 + \pi \cdot 1 = 5 - 4\pi,$$

and

$$\text{EU}_{\text{Bob}}(\text{East}) = (1 - \pi) \cdot 0 + \pi \cdot 2 = 2\pi.$$

Therefore, Bob strictly prefers to go to the East party if and only if

$$\text{EU}_{\text{Bob}}(\text{East}) > \text{EU}_{\text{Bob}}(\text{West}) \quad \Leftrightarrow \quad 2\pi > 5 - 4\pi \quad \Leftrightarrow \quad \pi > 5/6.$$

If the probability of Charlie going to the East party is high enough, then Bob would like to go there to meet her. Otherwise, he would rather go to the West party, which would be more fun. A similar analysis would reveal that Charlie will go to the East party, only if she thinks it is sufficiently likely for Bob to be there. We have thus determined the optimal action for each player as a function of the behavior of the other players. These functions are called *best response* functions. For both players, each of the available actions can be the unique best response. This implies that cautiousness alone cannot rule out *any* action.

The EUH allows to predict that players will always chose a best response. This is not enough to make informative predictions, because best response functions are not enough for each player to choose an optimal action. For instance, Bob would still need to know whether $\pi > 5/6$ or not. However, there is no app that can tell him the probability of Charlie is going to each of the two parties. As economists, we face a similar problem. Predicting behavior for this game requires a model of how the players *form beliefs* about their mutual behavior. There are different ways to address this issue, none of which is perfect. Section 3 discusses some of them. The bottom line is that the outcome of *some* games is hard to predict. For the remainder of this course, we will focus on predictable games.

2.3. Social dilemmas

Our predictions are based on the assumption that players choose actions which are desirable from their individual perspective. There is no guarantee that doing so results in outcomes which are socially desirable. Consider for instance the following example, known as the *prisoners' dilemma*. Suppose that David and Eric are suspects of a crime. The district attorney (DA) has enough evidence to convict them for a misdemeanor, but would require a signed confession in order

David \ Eric	Silent	Confess
Silent	$-1, -1$	$-5, 0$
Confess	$0, -5$	$-4, -4$

Figure 4 – Prisoners' dilemma

to convict them for the felony they allegedly committed. She offers each of the prisoners a sentence reduction in exchange for a confession. Each prisoner must choose whether to confess to the crime, or to remain silent. The prisoners are held in different cells and have no way of communicating with one another.

The prisoners' payoffs are depicted in Figure 4. If both prisoners remain silent, they will receive a short sentence and get a utility of -1. If only one of the prisoners confesses, he will walk out free and get a utility of 0. The prisoner that did not confess will receive a long sentence and receive a utility of -5. Finally, if both the prisoners confess, they will both receive a sentence reduction. However, given that the DA now has enough evidence to convict them for the felony, they will have to serve an intermediate sentence and their utility will be -4.

Note that each prisoner is made better off by obtaining a sentence reduction, regardless of whether his accomplice confesses or not. Confessing thus dominates not confessing. Under the EUH hypothesis, this implies that both prisoners will always confess and serve an intermediate sentence. However this outcome is Pareto dominated from the perspective of the prisoners. If they both remained silent, they would only have to serve a short sentence. In the prisoners' dilemma, individual incentives lead to inefficiency from a social perspective.

Many important social situations have a similar structure similar to the prisoners' dilemma. Consider for example two rival superpowers engaged in an arms race. Each superpower is better off having a bigger arsenal than its rival. However, both superpowers would be better off if they both had small arsenals, instead of they both having large arsenals. The resources they need to maintain military superiority could be devoted to other purposes.

Other situation that fits the structure of the prisoners' dilemma is trade liberalization between two trading countries. Countries can sometimes benefit from unilateral policies that restrict imports, because such policies give monopolistic power to local producers. As a result, the local producers can capture a larger share of the total surplus. However, as we have learned before, monopolistic power also has the effect of reducing total surplus. In most cases, trade liberalization is Pareto improving. If *both* countries were to remove trade restrictions, *both* countries could be better off. Both in the case of arm races, and in the case of trade liberalization, there exist international institutions whose main purpose is to try to foster cooperative behavior.

Another important class of social dilemmas is the voluntary provision of *public goods*. Public goods are goods that are are enjoyed by society as large, such as clean air, national security, or public parks. Most public goods are provided directly by the government, and there are good reasons for that. What would happen if public goods were funded by voluntary contributions from the members of society?

Let us analyze a simple model in which each member of society voluntarily chooses how much to produce of the public good. Each member of society benefits from the *total* quantity of the public good provided, but only pays the cost of her individual contribution. Let $a_i \in \mathbb{R}_+$ denote i's contribution. Suppose that all individuals have the same cost function $C(\cdot)$, which exhibits increasing marginal costs. In other words, the utility of individual i is given by

$$u_i(a) = -C(a_i) + \sum_j a_j,$$

where the summation is taken over all the individuals in society, including i.

If individual i chooses how much she wants to contribute, she will choose the level a_i^* such that $\mathrm{MC}(a_i^*) = 1$. This is the level at which the marginal cost of producing an additional unit of the public good equals her *individual* marginal benefit.

In contrast, now suppose that we want to determine the level of the public good that maximizes the sum of individual utilities:[1]

$$\sum_i u_i(a) = -\sum_i C(a_i) + n \sum_i a_i, \tag{1}$$

where n denotes the total number of people that benefit from the public good. The

[1] Recall that, if we assume monetary transfers are possible and agents have quasi-linear preferences, then an outcome is Pareto efficient if and only if it maximizes the sum of individual utilities.

socially optimal contribution a_i^0 is given by the first order condition $\mathrm{MC}(a_i^0) = n$. This condition equates the marginal cost suffered by i to the *social* benefit from an additional unit of the public good. Recall that we assumed increasing marginal cost. That means that $MC(\cdot)$ is an increasing function, we can this infer that $a_i^0 > a_i^*$ (why?). That is, when individuals choose voluntarily how much to contribute to the provision of a public good, the good will be *under-provided* from a social perspective.

3. Rationalizability

3.1. Common knowledge

So far, we have assumed that all players are rational and cautious. We have not yet assumed anything about whether the players *know* that these assumptions hold. We also have not assumed that players *know that their opponents know* that these assumptions hold. As it turns out, making this kind of assumptions about the players' knowledge will allow us to make finer predictions about the players behavior. In particular we will assume that Assumptions 1 and 2 are *common knowledge* among the players, according to the following definition.

Definition 3 A fact is *mutually known* among a group of individuals if everybody knows it. It is *commonly known* if everybody knows it, and, in addition, everybody knows that everybody knows it, everybody knows that everybody knows that everybody knows it, and so on and so forth.

The following example illustrates that there is a big difference between mutual knowledge and common knowledge. Anna Bob and Charlie are sitting in opposite corners of a room with no mirrors. Each one of them is wearing either a blue hat or a red hat. Each one of them can see color of the hat of the other two people in the room, but not the color of his/her own hat. For instance, Anna can see that Bob and Charlie have red hats, but she cannot tell whether her own hat is red or blue. We assume that all this information is common knowledge. Anna is wearing a blue hat and Bob and Charlie are wearing red hats. See Figure 5.

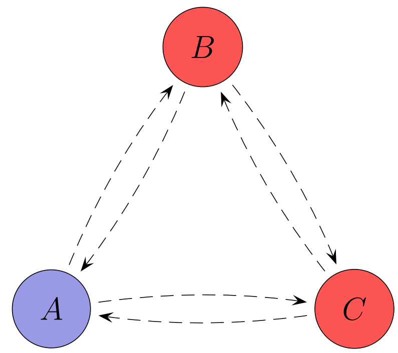

Figure 5 – Three logicians wearing hats

3.1.a. *Mutual knowledge.–* Suppose that Daniel enters the room and announces that everybody is wearing either blue or red hats. Then he proceeds to ask them one by one "*Which color is your hat?*" First he asks Anna, then Bob and then Charlie. In every case the answer is the same "*I don't know*".

Note that it is common knowledge that everybody is wearing either a blue hat or a red hat. This is because this fact was publicly announced and everybody noticed that everybody heard it. This however does not imply that there has to be either a red hat or a blue hat. It could very well be the case (given the information that Anna, Bob and Charlie have) that all hats are blue or all hats are red.

Anna knows that there are at least two red hats, because she can see Bob and Charlie's hats. But this does not imply any information about her own hat. Similarly, Bob and Charlie know that there is at least one red hat, and at least one blue hat. But they cannot infer anything about the color of their own hats. Hence nobody is able to provide a definitive answer to Daniel's question. *Notice that it is mutual knowledge (everybody knows) that there is at least one red hat.*

3.1.b. *Common knowledge.–* Now suppose that Daniel announces that everybody is wearing either a blue or a red hat. In addition, he also announces that there is at least one red hat in the room. The he proceeds as before asking them one by one "*Which color is your hat?*". Anna and Bob answers as before: "*I don't know*". However, Charlie answers triumphant: "*My hat is red!*"

The only difference between the two scenarios is that, in the second one, Daniel made the additional announcement that there is at least one red hat. However, this is something that *everybody already knew*. The difference is that, by making the announcement public, the existence of at least one red hat went from being mutual knowledge to being common knowledge. After the announcement, everybody knew that everybody knew that there was at least one red hat. This is what allowed Charlie to deduce that her hat was red. Let's see how.

Charlie knew that Bob knew that there was at least one red hat. Hence, if he had seen only blue hats, he would have known that his own hat had to be the red one. Since he did not know the color of his own hat, it had to be the case that he was already seeing at least one red hat. That is, either Anna or Charlie hat to be wearing a red hat. Since Charlie could see that Anna's hat was blue, this meant that her own hat had to be red. This line of thought was only possible because she knew that Bob knew that there was at least one red hat.

This cartoon about three logicians in a bar tells a simpler version of the story.

3.2. Keynesian beauty contests

Recall the p-beauty contest we played in class. Each student submitted a number from the interval $[0, 99]$. You were not allowed to talk with each other: choices were meant to be independent. The student whose submission was closes to *two thirds of the average* of the submissions got some extra credit over his/her final grade for the course.

Note that the average submission could *not* be greater than 99. Hence, two thirds of the average could *not* be greater than $(2/3)\cdot 99 = 66$. Submissions greater than 66 were thus weakly dominated (why?). Cautious players would thus never submit any guess greater than 66. But that is not where the story ends.

A player who thought other players were cautious would be able to infer that two thirds of the average could *not* be greater than $(2/3) \cdot 66 = 40$. She would thus never guess anything above 40. This process can be iterated. A player who knew that everybody knew that everybody is cautious would have inferred that her opponents would never guess above 40. She would thus never guess above $(2/3) \cdot 40 = 26.\bar{6}$. A player who knew that everybody knew that everybody knew that everybody was cautions would never guess above $(2/3)\cdot 26.\bar{6} = 17.\bar{7}$. Iterating this process ad infinitum results in the following claim.

Claim 4 *If it is common knowledge that all players are cautious, then all players would bid exactly* 0 *in the p-beauty contest.*

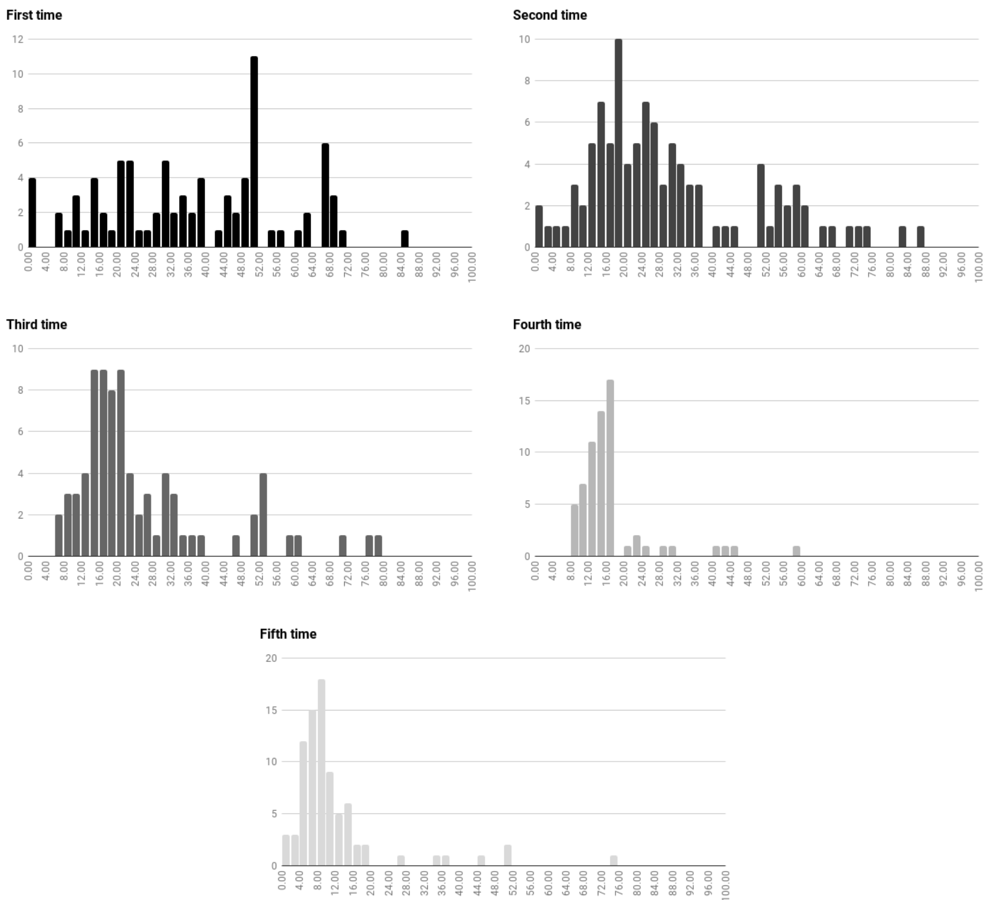

Figure 6 – Results from the p-beauty contest

When we played the game in class, the average submission was approximately 40, and 5% of the participants submitted weakly dominated bids above 66. Similar findings findings have been replicated in different settings, with different types of audiences and rewards. When experimental subjects are asked to participate in this game for the first time, the average submission is usually well above 0. Also, there is usually a small proportion of people making weakly dominated guesses. However, once the participants have some experience playing the game their behavior converges rapidly towards 0. Figure 6 shows the distribution of in-class guesses from a previous course in which we played the game 5 times in non-consecutive weeks.

	L	M	R
T	$1,3$	$5,2$	$0,1$
B	$0,1$	$3,8$	$4,2$

(a)

	L	M
T	$1,3$	$5,2$
B	$0,1$	$3,8$

(b)

	L	M
T	$1,3$	$5,2$

(c)

	L
T	$1,3$

(d)

Figure 7 – A 2×3 example

3.3. Iterated dominance

For the remainder of the course, we will maintain the assumption that there is common knowledge of rationality among the players. Those choices that are consistent with this assumption are called *rationalizable*.

Assumption 3 (Rationalizablity) There is common knowledge of rationality.

Consider the example in Figure 7(a). Action R is dominated by action M for the column player (recall definition 1). Since the column player is rational, she would never choose R (recall claim 1). If the row player knows that the column player is rational, he would never expect R to be chosen. Hence, from his perspective, the game looks as is Figure 7(b). This reduced game was obtained from the original game by *eliminating a dominated action*.

In the reduced game from Figure 7(b), action B is dominated by T. Hence, if the row player is rational and knows that the column player is rational, he would choose T. If the column player knew that the row player is rational and knows that she is rational, the game for her would look as in Figure 7(c). Being rational, she would choose L. The only *ratiozanlizable* outcome for this game is when the row player chooses T and the column player chooses L. The algorithm we used to reach this conclusion is called iterated dominance (or iterated elimination of strictly dominated actions, if you want to be more precise).

Definition 4 The *iterated dominance* algorithm is an iterative process applied to games. In each iteration, the game is reduced by eliminating actions which are dominated. The algorithm stops when there are no more dominated actions.

	L	M	R
T	$1, 3$	$5, 2$	$0, 1$
B	$0, 1$	$3, 8$	$4, 2$

Figure 8 – Iterated dominance in a 2×3 example

Claim 5 *All rationalizable actions survive iterated dominance.*

Justification. If an action is eliminated on the first round of elimination, then we know from claim 1 that it is not consistent with rationality. If it is eliminated in the second round, then rational players who know their opponents are rational would not play it. Hence, actions eliminated in the second round are not consistent with two orders of mutual knowledge of rationality. Likewise, actions eliminated in the nth round are not consistent with n orders of mutual knowledge of rationality. Hence, actions that are eliminated at any stage are *not* consistent with common knowledge of rationality. ■

For finite two-player games, the iterated dominance algorithm can be carried out by crossing out dominated actions in each iteration. See, for instance, Figure 8. In the first round of elimination we crossed out R. In the second round we crossed out B. In the third and final round we crossed out M. This is a much more succinct way of carrying out the algorithm than the one we used in Figure 7.

The description of the iterated dominance algorithm in Definition 4 is purposely vague. It each elimination round, there could be more than one dominated action. Which should be eliminated? Should you eliminate all of them at the same time? Should you eliminate all the ones for one player first? Should you alternate between players eliminating one action at a time? It turns out that it does not matter. The order of elimination does not affect the result of the algorithm. As long as you continue to eliminate dominated actions each round, and don't stop until there are no more dominated actions left, the resulting set of surviving actions will be the same.

The game from Figure 7 has a unique rationalizable outcome. Not all games are like that. All games have rationalizable outcomes, but some games might have more than one. Consider for instance the game in Figure 9. In the first round, d is dominated by c. In the second round, y is dominated by x. In the third round,

	a	b	c	d
w	$5, \underline{6}$	$4, 4$	$\underline{6}, 5$	$12, 2$
x	$3, 7$	$\underline{8}, 7$	$\underline{6}, \underline{8}$	$10, 6$
y	$2, 10$	$7, \underline{11}$	$4, 6$	$\underline{13}, 5$
z	$\underline{6}, 4$	$5, 9$	$4, \underline{10}$	$10, 9$

Figure 9 – Iterated dominance in a 4×4 example

b is dominated by c. After that, there are no more dominated actions, and all the remaining actions are rationalizable.

3.4. Best responses

The utility each player receives in a game depends both on her own action and that of her opponents. Hence, her optimal action might depend on what other players do. This is captured by the following definition.

Definition 5 An action a_i *best response* to a_{-i} if it maximizes i's utility when i's opponents play a_{-i}. Player i's best response function $\mathrm{BR}_i(\cdot)$ specifies i's best responses as a function of her opponent's actions.

Consider the game depicted on Figure 9. I have highlighted best responses both for the row player and for the column player, by underlining the corresponding payoffs. For example, the number 13 is highlighted because y is a best response to d for the row player. Similarly, the number 11 is highlighted because b is a best response to y for the column player. Highlighting best responses in this manner can be a useful first step to execute the iterated dominance algorithm. Because if an action is a best response it cannot be dominated (recall Claim 1).

Definition 6 A pure-strategy *Nash equilibrium* (NE) is an action profile a^* such that each player's action is a best response to her opponent's action, i.e., such that $a_i^* = \mathrm{BR}_i(a_{-i}^*)$ for all i

When there is a *unique* rationalizable outcome, it is a NE. This is the case for the game in Figure 7. Indeed, T is a best response to L, and L is a best response to T. However, there are games with many rationalizable outcomes which are not NE. Consider for instance the game from Figure 9. The only NE in this game is (x,c), however there are many more rationalizable outcomes.

Assuming that there is common knowledge of rationality does *not* imply that players will always play a NE. That conclusion would require an additional assumption, namely that players can correctly anticipate each other behavior. This might be a reasonable assumption in some situations in which players agree on a plan of action before playing the game. However, there is no good reason to make this assumption in general situations. For the remainder of the course we will only make predictions using assumptions 1, 2, and 3. Fortunately, this suffices to study many important situations.

4. Dynamic games

Pending. . .

Ü///