

Lecture 1

Introduction and Examples

This book has three parts, each with its own overarching goal. Lectures 2–10 develop tools for designing systems with strategic participants that have good performance guarantees. The goal of Lectures 11–15 is to understand when selfish behavior is largely benign. Lectures 16–20 study if and how strategic players reach an equilibrium of a game. The three sections of this lecture offer motivating examples for the three parts of the book.

1.1 The Science of Rule-Making

We begin with a cautionary tale. In 2012, the Olympics were held in London. One of the biggest scandals of the event concerned, of all sports, women’s badminton. The scandal did not involve any failed drug tests, but rather a failed tournament design that did not carefully consider *incentives*.

The tournament design used is familiar from World Cup soccer. There are four groups (A, B, C, D) of four teams each. The tournament has two phases. In the first “round-robin” phase, each team plays the other three teams in its group, and does not play teams in other groups. The top two teams from each group advance to the second phase, while the bottom two teams from each group are eliminated. In the second phase, the remaining eight teams play a standard “knockout” tournament. There are four quarterfinals, with the losers eliminated, followed by two semifinals, with the losers playing an extra match to decide the bronze medal. The winner of the final gets the gold medal, the loser the silver.

The incentives of participants and of the Olympic Committee and fans are not necessarily aligned in such a tournament. What does a team want? To get as prestigious a medal as possible. What does the Olympic Committee want? They didn’t seem to think carefully

about this question, but in hindsight it is clear that they wanted every team to try their best to win every match. Why would a team ever want to lose a match? Indeed, in the knockout phase of the tournament, where losing leads to instant elimination, it is clear that winning is always better than losing.

To understand the incentive issues, we need to explain how the eight winners from the round-robin phase are paired up in the quarterfinals (Figure 1.1). The team with the best record from group A plays the second-best team from group C in the first quarterfinal, and similarly with the best team from group C and the second-best team from group A in the third quarterfinal. The top two teams from groups B and D are paired up analogously in the second and fourth quarterfinals. The dominoes started to fall when, on the last day of round-robin competition, there was a shocking upset: the Danish team of Pedersen and Juhl (PJ) beat the Chinese team of Tian and Zhao (TZ), and as a result PJ won group D with TZ coming in second. Both teams advanced to the knockout stage of the tournament.




Figure 1.1: The women’s badminton tournament at the 2012 Olympics. Both WY and JK preferred to play TZ in as late a round as possible.

The first controversial match involved another team from China, Wang and Yu (WY), and the South Korean team of Jung and Kim (JK). Both teams had a 2-0 record in group A play. Thus, both were headed for the knockout stage, with the winner and loser of this

match the top and second-best team from the group, respectively. Here was the issue: the group A winner would likely meet the fearsome TZ team in the semifinals of the knockout stage, where a loss means a bronze medal at best, while the second-best team in group A would not face TZ until the final, with a silver medal guaranteed. Both the WY and JK teams found the difference between these two scenarios significant enough to try to deliberately lose the match!¹ This unappealing spectacle led to scandal, derision, and, ultimately, the disqualification of the WY and JK teams.² Two group C teams, one from Indonesia and a second team from South Korea, were disqualified for similar reasons.

The point is that, in systems with strategic participants, *the rules matter*. Poorly designed systems suffer from unexpected and undesirable results. The burden lies on the system designer to anticipate strategic behavior, not on the participants to behave against their own interests. We can't blame the badminton players for optimizing their own medal placement.

There is a well-developed science of rule-making, the field of *mechanism design*. The goal in this field is to design rules so that strategic behavior by participants leads to a desirable outcome. Killer applications of mechanism design that we discuss in detail include Internet search auctions, wireless spectrum auctions, the matching of medical residents to hospitals, and kidney exchanges.

Lectures 2–10 cover some of the basics of the traditional economic approach to mechanism design, along with several complementary contributions from computer science that focus on computational efficiency, approximate optimality, and robust guarantees.

1.2 When Is Selfish Behavior Near-Optimal?

1.2.1 Braess's Paradox

Sometimes you don't have the luxury of designing the rules of a game from scratch, and instead want to understand a game that occurs

¹In hindsight, it seems justified that the teams feared the Chinese team TZ far more than the Danish team PJ: PJ were knocked out in the quarterfinals, while TZ won the gold medal.

²If you're having trouble imagining what a badminton match looks like when both teams are trying to lose, by all means track down the video on YouTube.

“in the wild.” For a motivating example, consider *Braess’s paradox* (Figure 1.2). There is an origin o , a destination d , and a fixed number of drivers commuting from o to d . For the moment, assume that there are two non-interfering routes from o to d , each comprising one long wide road and one short narrow road (Figure 1.2(a)). The travel time on a long wide road is one hour, no matter how much traffic uses it, while the travel time in hours on a short narrow road equals the fraction of traffic that uses it. This is indicated in Figure 1.2(a) by the edge labels “ $c(x) = 1$ ” and “ $c(x) = x$,” respectively. The combined travel time in hours of the two edges in one of these routes is $1 + x$, where x is the fraction of the traffic that uses the route. Since the routes are identical, traffic should split evenly between them. In this case, all drivers arrive at d an hour and a half after their departure from o .




Figure 1.2: Braess’s paradox. Each edge is labeled with a function that describes the travel time as a function of the fraction of the traffic that uses the edge. After the addition of the (v, w) edge, the price of anarchy is $4/3$.

Suppose we try to improve commute times by installing a teleportation device that allows drivers to travel instantly from v to w (Figure 1.2(b)). How will the drivers react?

We cannot expect the previous traffic pattern to persist in the new network. The travel time along the new route $o \rightarrow v \rightarrow w \rightarrow d$ is never worse than that along the two original paths, and it is strictly less whenever some traffic fails to use it. We therefore expect all drivers to deviate to the new route. Because of the ensuing heavy congestion on the edges (o, v) and (w, d) , all of these drivers now experience *two* hours of travel time from o to d . Braess’s paradox thus

shows that the intuitively helpful action of adding a new superfast link can negatively impact all of the traffic!

Braess's paradox also demonstrates that selfish routing does not minimize the commute time of drivers—in the network with the teleportation device, an altruistic dictator could assign routes to traffic to improve everyone's commute time by 25%. We define the *price of anarchy* (POA) as the ratio between the system performance with strategic players and the best-possible system performance. For the network in Figure 1.2(b), the POA is $\frac{2}{3/2} = \frac{4}{3}$.

The POA is close to 1 under reasonable conditions in a remarkably wide range of application domains, including network routing, scheduling, resource allocation, and auctions. In such cases, selfish behavior leads to a near-optimal outcome. For example, Lecture 12 proves that modest over-provisioning of network capacity guarantees that the POA of selfish routing is close to 1.

1.2.2 Strings and Springs

Braess's paradox is not just about traffic networks. For example, it has an analog in mechanical networks of strings and springs. In the device pictured in Figure 1.3, one end of a spring is attached to a fixed support and the other end to a string. A second identical spring is hung from the free end of the string and carries a heavy weight. Finally, strings are connected, with a tiny bit of slack, from the support to the upper end of the second spring and from the lower end of the first spring to the weight. Assuming that the springs are ideally elastic, the stretched length of a spring is a linear function of the force applied to it. We can therefore view the network of strings and springs as a traffic network, where force corresponds to traffic and physical distance corresponds to travel time.

With a suitable choice of string and spring lengths and spring constants, the equilibrium position of this mechanical network is described by Figure 1.3(a). Perhaps unbelievably, severing the taut string causes the weight to *rise*, as shown in Figure 1.3(b)! To explain this curiosity, note that the two springs are initially connected in series, so each bears the full weight and is stretched out to a certain length. After cutting the taut string, the two springs carry the weight in parallel. Each spring now carries only half of the weight, and accordingly is stretched to only half of its previous length. The




Figure 1.3: Strings and springs. Severing a taut string lifts a heavy weight.

rise in the weight is the same as the decrease in the commute time achieved by removing the teleporter from the network in Figure 1.2(b) to obtain the network in Figure 1.2(a).

1.3 Can Strategic Players Learn an Equilibrium?

Some games are easy to play. For example, in the second network of Braess's paradox (Figure 1.2(b)), using the teleporter is a no-brainer—it is the best route, no matter what other drivers do.

In most games, however, the best action to play depends on what the other players do. Rock-Paper-Scissors, rendered below in “bima-

trix” form, is a canonical example.

| | Rock | Paper | Scissors |
|----------|-------|-------|----------|
| Rock | 0, 0 | −1, 1 | 1, −1 |
| Paper | 1, −1 | 0, 0 | −1, 1 |
| Scissors | −1, 1 | 1, −1 | 0, 0 |

One player chooses a row and the other a column. The numbers in the corresponding matrix entry are the payoffs for the row and column player, respectively. More generally, a two-player game is specified by a finite strategy set for each player, and a payoff to each player for every pair of strategies that the players might choose.

Informally, an equilibrium is a steady state of a system where each participant, assuming everything else stays the same, wants to remain as is. There is certainly no “deterministic equilibrium” in the Rock-Paper-Scissors game: whatever the current state, at least one player can benefit from a unilateral deviation. For example, the outcome (Rock, Paper) cannot be an equilibrium, since the row player wants to switch and play Scissors.

When playing Rock-Paper-Scissors, it appears as if your opponent is randomizing over her three strategies. Such a probability distribution over strategies is called a *mixed* strategy. If both players randomize uniformly in Rock-Paper-Scissors, then neither player can increase her expected payoff via a unilateral deviation (all such deviations yield an expected payoff of zero). A pair of probability distributions with this property is a (*mixed-strategy*) *Nash equilibrium*.

Remarkably, allowing randomization, *every* game has at least one Nash equilibrium.

Theorem 1.1 (Nash’s Theorem) *Every finite two-player game has a Nash equilibrium.*

Nash’s theorem holds more generally in games with any finite number of players (Lecture 20).

Can a Nash equilibrium be computed efficiently, either by an algorithm or by strategic players themselves? In zero-sum games like Rock-Paper-Scissors, where the payoff pair in each entry sums to zero, this can be done via linear programming or, if a small amount of error

can be tolerated, via simple iterative learning algorithms (Lecture 18). These algorithmic results give credence to the Nash equilibrium concept as a good prediction of behavior in zero-sum games.

In non-zero-sum two-player games, however, recent results indicate that there is no computationally efficient algorithm for computing a Nash equilibrium (Lecture 20). Interestingly, the standard argument for computational intractability, “ \mathcal{NP} -hardness,” does not seem to apply to the problem. In this sense, the problem of computing a Nash equilibrium of a two-player game is a rare example of a natural problem exhibiting intermediate computational difficulty.

Many interpretations of an equilibrium concept involve someone—the participants or a designer—determining an equilibrium. If all parties are boundedly rational, then an equilibrium can be interpreted as a credible prediction only if it can be computed with reasonable effort. Computational intractability thus casts doubt on the predictive power of an equilibrium concept. Intractability is certainly not the first stone to be thrown at the Nash equilibrium concept. For example, games can have multiple Nash equilibria, and this non-uniqueness diminishes the predictive power of the concept. Nonetheless, the intractability critique is an important one, and it is most naturally formalized using concepts from computer science. It also provides novel motivation for studying computationally tractable equilibrium concepts such as correlated and coarse correlated equilibria (Lectures 13, 17, and 18).

The Upshot

- ☆ The women’s badminton scandal at the 2012 Olympics was caused by a misalignment of the goal of the teams and that of the Olympic Committee.
- ☆ The burden lies on the system designer to anticipate strategic behavior, not on the participants to behave against their own interests.
- ☆ Braess’s paradox shows that adding a superfast link to a network can negatively impact all of the traffic. Analogously, cutting a taut string

in a network of strings and springs can cause a heavy weight to rise.

- ☆ The price of anarchy (POA) is the ratio between the system performance with strategic players and the best-possible system performance. When the POA is close to 1, selfish behavior is largely benign.
- ☆ A game is specified by a set of players, a strategy set for each player, and a payoff to each player in each outcome.
- ☆ In a Nash equilibrium, no player can increase her expected payoff by a unilateral deviation. Nash's theorem states that every finite game has at least one Nash equilibrium in mixed (i.e., randomized) strategies.
- ☆ The problem of computing a Nash equilibrium of a two-player game is a rare example of a natural problem exhibiting intermediate computational difficulty.

Notes

Hartline and Kleinberg (2012) relate the 2012 Olympic women's badminton scandal to mechanism design. Braess's paradox is from Braess (1968), and the strings and springs interpretation is from Cohen and Horowitz (1991). There are several physical demonstrations of Braess's paradox on YouTube. See Roughgarden (2006) and the references therein for numerous generalizations of Braess's paradox. Koutsoupias and Papadimitriou (1999) define the price of anarchy. Theorem 1.1 is from Nash (1950). The idea that markets implicitly compute a solution to a significant computational problem goes back at least to Adam Smith's "invisible hand" (Smith, 1776). Rabin (1957) is an early discussion of the conflict between bounded rationality and certain game-theoretic equilibrium concepts.

Exercises

Exercise 1.1 Give at least two suggestions for how to modify the Olympic badminton tournament format to reduce or eliminate the incentive for a team to intentionally lose a match.

Exercise 1.2 Watch the scene from the movie *A Beautiful Mind* that purports to explain what a Nash equilibrium is. (It's easy to find on YouTube.) The scenario described is most easily modeled as a game with four players (the men), each with the same five actions (the women). Explain why the solution proposed by the John Nash character is not a Nash equilibrium.

Exercise 1.3 Prove that there is a unique (mixed-strategy) Nash equilibrium in the Rock-Paper-Scissors game.

Problems

Problem 1.1 Identify a real-world system in which the goals of some of the participants and the designer are fundamentally misaligned, leading to manipulative behavior by the participants. A “system” could be, for example, a Web site, a competition, or a political process. Propose how to improve the system to mitigate the incentive problems. Your answer should include:

- (a) A description of the system, detailed enough that you can express clearly the incentive problems and your solutions for them.
- (b) Anecdotal or demonstrated evidence that participants are gaming the system in undesirable ways.
- (c) A convincing argument why your proposed changes would reduce or eliminate the strategic behavior that you identified.

Problem 1.2 Can you produce a better video demonstration of Braess's paradox than those currently on YouTube? Possible dimensions for improvement include the magnitude of the weight's rise, production values, and dramatic content.