An Application of Game Theory to Electronic Communications Markets*

Bernhard von Stengel
London School of Economics
November 2011

This article gives an introductory survey of non-cooperative game theory. Its central concepts are illustrated by means of examples, often with reference to economic and business situations in information and communication technology. Game theory provides a number of insights. Selfish behavior may lead to inferior outcomes for all players, because a more cooperative solution is not stable with respect to individual changes in behavior. The concept of equilibrium captures this stability. Game trees represent dynamic interactions, as illustrated with an example of market entry and entry deterrence. A game may have more than one equilibrium, which is important to note when predicting the outcome of an interaction. Different technologies may be adapted as equilibrium solutions due to the number of users already using a technology, as shown with the "bandwidth choice game". The rules of the game matter, for example when introducing commitment power in the "quality game". Players may have different information about scenarios, and it may be optimal to hide that information from an opponent by actively randomizing one's actions. In network routing, selfish behavior may lead to congestion, which can even become worse when increasing network capacity. These concepts and examples show the versatility of game theory for analyzing the dynamics of markets and the interactions of cooperating and competing participants.

1. What is game theory?

Game theory provides mathematical tools for modeling and analyzing interactive decisions. It considers *players* that represent decision-making firms or persons, and their possible *actions*, and models their incentives by *payoffs* that the players want to maximize. If needed, possible *scenarios* can be represented as "states of nature". The players may have different *information* about that state and about the actions taken by the other players.

A *game* defines how the interaction results in payoffs. The game can have various levels of abstraction. Game *trees* and other dynamic games show sequential moves of players. The *strategic form* is a table where players' strategies are simply listed and not detailed as in the dynamic setting. At the most abstract level, *cooperative games* show only the outcomes that coalitions of many players can achieve, without detailing how these coalitions are formed.

We will describe how some economic situations are analyzed by games in strategic form and by game trees. These are the basic tools of *non-cooperative game theory*, which has become the language of modern economic theory. Its basic tenet is that players will cooperate (if they do) out of self-interest, and that their incentives to do so have to be explicitly described.

Given a game in some description, the game-theoretic analysis proposes a *solution* that suggests how players should behave, in their own interest. Such solutions may be compelling to varying degrees. Sometimes a player's action may be *dominant* in the sense that it gives her a best possible payoff irrespective of what the other players do. However, as the famous *Prisoner's*

* Published in: Regulatory and Economic Policy in Telecommunications, Geer No. 7, November 2011 (online accessible at http://publications.telefonica.com/geer_7), Telefónica, Spain, pages 112-121.

Dilemma demonstrates, even a dominant solution may be worse than a more cooperative but unstable outcome.

Not all games admit a dominant solution. A solution that always exists is an *equilibrium* (where we here mean *Nash equilibrium*), which recommends a plan of action to every player that is optimal if all other players follow their recommendations. Unfortunately, there may be more than one equilibrium, which reduces the use of this concept as a recommendation.

However, inferior dominant solutions that result from selfish behavior, or equilibrium predictions that are not unique, are one of the insights provided by game theory, which we will elucidate in the following sections.

2. Explaining selfish behavior

Figure 1 shows a table that represents a two-player price-setting "duopoly" game in strategic form. Two firms can sell their competing products at a ***high*** or ***low*** price. Firm 1 (in red) chooses a row and, at the same time, firm 2 (in blue) chooses a column of the table as an action. The resulting cell gives two payoffs in the respective colors to the two players. The payoffs are staggered (in each cell, in the lower left corner for firm 1, in the upper right corner for firm 2) so that each player's *payoff matrix* (red or blue) is more easily seen when looking at the entire table.

1 \ 2	***high***	***low***
high	2 2	3 0
low	0 3	1 1

Figure 1. Strategic-form price-setting game with dominant solution (***low***, ***low***).

The best possible pair of actions to the two firms seems clearly that both charge a **high** price. However, if firm 2 charges a high price, then firm 1 can increase her payoff from 2 to 3 by playing ***low*** rather than ***high***, and if firm 2 charges a low price, then firm 1 can increase in the same way her payoff from 0 to 1. Hence, ***low*** is a dominant action for firm 1 because it gives her a higher payoff for each possible action of firm 2. The same applies when the roles of the players are exchanged because the game is symmetric. So both players choose ***low***, and get payoff 1 each. This is smaller than the payoff 2 they would get if both chose ***high***, but that pair of actions is not "stable".

The game in Figure 1 is the well-known "Prisoner's Dilemma" game (whose name derives from a different story, see, for example, Osborne 2004). In the context of a duopoly, the pair of dominant actions that hurts the players has the positive side effect of low prices for the firms' customers. In other situations where this type of game applies, the resulting selfish behavior may hurt all, typically in problems of "free-riding" or the "tragedy of the commons" such as fish stocks that are depleted by overfishing – which may be individually rational. The game-theoretic analysis explains this behavior, and illustrates the necessity of regulations that enforce the "cooperative" outcome.

3. Thinking strategically

Our next example illustrates the use of *game trees*. It shows how game theory might have helped the printer company Epson when it entered the market for laser printers (the example is inspired by Dixit and Nalebuff, 1991; the actual game we consider is different). In the 1980s, the market for office laser printers was dominated by Hewlett-Packard (HP) who sold them with a high profit margin. Epson was the market leader for the cheaper and lower-quality dot-matrix printers. Epson decided to develop a laser printer, assuming it could enter that market profitably. However, Epson did apparently not take into account that HP was another *player* who would *react*. HP lowered its prices and a price war resulted. After that, apart from lower profits and lower market share for Epson than expected, laser printers had become so cheap that much fewer people now bought dot-matrix printers – in effect, Epson had destroyed its own market.

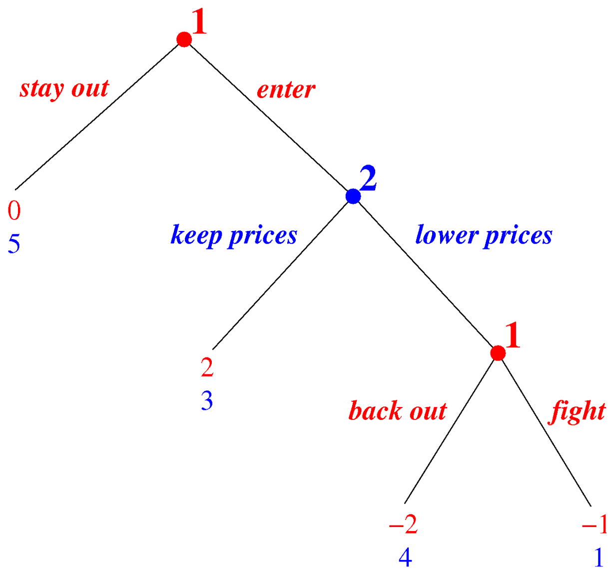

Figure 2. Game tree modeling entry of player 1 into a market with incumbent player 2.

The game tree in Figure 2 reflects one possible view of the game between player 1 (Epson) as market entrant and player 2 (HP) as incumbent. In a game tree, play progresses sequentially by players making *moves* at *decision nodes*, until a *terminal node* is reached where each player receives a *payoff*. Here the game starts at the first node (at the top, the tree grows downwards) with a move of player 1. Player 1's first move is either to ***stay out*** of the market, or to ***enter*** it. If player 1 stays out, play ends with payoff 0 to player 1 and payoff 5 to player 2 (in some units, for example profits in ten million dollars). If player 1 chooses to ***enter*** the market, then player 2 may react by choosing to ***keep prices*** or to ***lower prices***. After keeping prices, the game ends with payoff 2 to player 1 and payoff 3 to player 2. If player 2 chooses to lower prices, player 1 has another move and may decide to ***back out*** from the market, getting negative payoff -2 and player 2 a payoff of 4 (although player 2 has regained the market, profits are now lower). Player 1 may also decide to ***fight*** a price war, resulting into payoff -1 for player 1 and payoff 1 for player 2.

One way to analyze a game tree like in Figure 2 is by a *roll-back analysis* (also called “backward induction”), which essentially means thinking ahead about what players will do. In that analysis, the most advanced move one has to consider is when player 1 has to choose between ***back out*** and ***fight***, where she will prefer ***fight*** because the resulting payoff -1 is less negative than -2. Anticipating that move, player 2 would get payoff 1 when choosing to ***lower prices***, versus getting payoff 3 when keeping prices, so his best action is to ***keep prices***. In turn, the best move at the beginning for player 1 is then to ***enter*** the market because she will get 2 rather than 0 when staying out.

If players play according to the roll-back analysis, then player 1 will enter the market and player 2 will keep prices and the game terminates with payoffs 2 and 3 for the two players. However, the result of the roll-back analysis is more than just a prediction of the path that a play of the game will take, because it also gives a *strategy* for each player. The term "strategy" for a game tree has a precise technical meaning, which agrees with the intuitive understanding of "being prepared for every eventuality". In a game tree, a strategy is a complete plan of action that specifies a move for every decision node that a player can reach. Here player 2 has only one decision node, so his strategies are just his two moves, ***keep prices*** and ***lower prices***. Player 1 has two decision nodes, and the result from the roll-back analysis is his move combination ***enter-fight***, that is, to enter the market and then, also, to be prepared to fight a price war in case player 2 decides to lower prices.

The roll-back analysis specifies a strategy to each player. It can be shown that these strategies always define an *equilibrium*, which means that each player's strategy is optimal for that player, given that the other player or players adhere to their strategies. Why is the strategy pair (***enter-fight***, ***keep prices***) an equilibrium in the game tree in Figure 2? For player 1, ***enter-fight*** is optimal because her payoff of 2, given that player 2 keeps prices, can only change to 0 with ***stay out***, or stays the same when changing to the strategy ***enter-back out***, so player 1 cannot improve her payoff. For player 2, given that player 1 chose ***enter-fight***, his strategy ***keep prices*** gives him payoff 3 whereas changing to ***lower prices*** would give him a payoff of 1 which is worse. Note that the anticipated move of player 1 to ***fight*** will not take effect, but is needed to assert that player 2 plays optimally. This is why strategies, that is, complete plans of moves are needed to specify an equilibrium, not just those moves that lead to a particular terminal node of the game tree.

When Epson entered the market for laser printers, the above prediction of the roll-back analysis that HP would keep its prices did not happen, but instead the players wound up at the strategy pair (***enter-fight***, ***lower prices***) with payoff -1 for player 1 and payoff 1 for player 2. What went wrong? Did Epson not think strategically at all (assuming that HP would just keep its prices), or miscalculate the best response of HP, and was HP responding irrationally? We will show that our game-theoretic model, assuming it is correct, gives one possible explanation because the game has *another equilibrium* which is not predicted by the roll-back analysis.

1 \ 2	*keep prices*	*lower prices*
stay out	0, 5	0, 5
enter–back out	2, 3	−2, 4
enter–fight	2, 3	−1, 1

Figure 3. Strategic form of the market-entry game tree in Figure 2.

In order to identify all equilibria of a game tree, it is converted to strategic form. The strategic form tabulates the players' strategies and the resulting payoffs when these strategies meet. Figure 3 shows the strategic form for the game tree in Figure 2. There are six cells in that table, some of which list the same payoff pair because the game tree has only four terminal nodes. The first row is the strategy ***stay out*** of player 1 (in fact, there is no need to combine that move of player 1 with the later moves ***back out*** or ***fight*** because player 1 will not reach that stage due to her own first move ***stay out***, so this suffices as a plan of action). Both cells in that row give the payoffs that result after ***stay out***, namely 0 and 5 to the two players; note that the columns ***keep prices*** and ***lower prices*** do not have an effect in that first row. The second and third row are the two strategies ***enter-back out*** and ***enter-fight***, and both cells in the left column (***keep prices***) have the same payoffs 2 and 3 to the two players. The remaining two cells in the right column (***lower prices***) give the payoffs from the last two nodes in the game tree.

Given the strategic form, an equilibrium is easily found by finding for each row and column the other player's optimal strategy (or *best response*) and checking the cells where these best responses coincide. In addition to the cell (***enter-fight***, ***keep prices***) found by the roll-back analysis, there is another such cell: (***stay out***, ***lower prices***). The reason is that ***lower prices*** is a best response to ***stay out*** (player 2 is no worse off than when choosing to ***keep prices***), and conversely player 1 gets the best possible payoff (0, rather than -1 or -2) by choosing to ***stay out***.

This second equilibrium (***stay out***, ***lower prices***) is known as *entry deterrence* in the "Chain Store paradox" (Selten, 1978). Player 2 *threatens* a move that is self-hurting if it where to implemented, because player 2 would be better off to keep rather than lower prices if player 1 entered the market. However, given the threat that player 2 executes his move to lower prices, player 1 will not enter the market, and so the threat does not result in a lower payoff and is therefore a best response.

The game in Figures 2 and 3 has more than one equilibrium, and if the players do not have the same view on which equilibrium should be played, their strategies might fail to match, as when player 1 chooses ***enter-fight*** and player 2 chooses to ***lower prices***. This is one possible view of why Epson and HP did not play equilibrium. In practice, the game tree is of course very simplified. Even if Epson and HP planned their moves that way, HP might have had a different view on Epson's cost for fighting a price war, for example a more negative payoff (like -3) after ***fight*** and thus expected that Epson would choose to ***back out***. In order to be successful, a game-theoretic analysis in management consulting must be sufficiently flexible and robust, by testing various scenarios, rather than considering a single detailed model (Lindstädt and Müller, 2010).

4. Equilibrium selection by evolution

We give another example of a game (from Turocy and von Stengel, 2002) that has two equilibria, and where either equilibrium may be selected via a dynamical process according to an *evolutionary* interpretation (Maynard Smith, 1982; Sandholm, 2010).

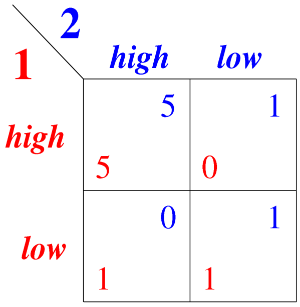

Figure 4. Bandwidth choice game with two equilibria, (***high***, ***high***) and (***low***, ***low***).

Two firms want to invest in communication infrastructure. They intend to communicate frequently with each other using that infrastructure, but they decide independently on what to buy. Each firm can decide between ***high*** or ***low*** bandwidth equipment. A low bandwidth connection works equally well (payoff 1) regardless of whether the other side has high or low bandwidth. However, changing from ***low*** to ***high*** is preferable only if the other side has high bandwidth (payoff 5), otherwise it incurs unnecessary cost (payoff 0).

The game is shown in Figure 4 and has the two equilibria (***high***, ***high***) and (***low***, ***low***). The strategy ***low*** has the better *worst-case* payoff, as considered for all possible strategies of the other player, no matter if these strategies are best responses or not. The strategy ***low*** is therefore also called a *max-min strategy*. In a sense, investing only in low bandwidth equipment is a safe choice. Moreover, this strategy is part of an equilibrium, and entirely justified if the player expects the other player to do the same.

The bandwidth choice game in Figure 4 can be given a different interpretation where it applies to a large population of individuals, each of which can adopt one of the two strategies. The game describes the payoffs that result when two of these individuals meet randomly. The dynamics of this game depends on the fraction of individuals adopting the respective strategies. Given this distribution of strategies, individuals with better average payoff will be more successful than others, so that their proportion in the population increases over time. In this game, if more than 1/5 of all individuals play the strategy ***high***, then ***high*** will get an average payoff of more than 1 and be more successful than ***low***, so that the fraction of individuals that play ***high*** increases and eventually the equilibrium (***high***, ***high***) will result where everyone plays ***high***. If the fraction for ***high*** is less than 1/5, then ***low*** will become ever more successful and eventually everyone will play ***low***. This typical *network effect* shows how *starting conditions* can influence whether a desirable or undesirable equilibrium gets established.

Evolutionary games, not discussed here, can also demonstrate that an initially small fraction of players adopting a superior strategy can rapidly invade an existing population. A modern book on evolutionary game theory is Sandholm (2010).

5. Rules of the game and commitment power

The rules of the game matter, as changing the rules for the *quality game* in Figure 5 will demonstrate. The two non-symmetric players are a service provider (player 1) and potential customer (player 2). The provider can decide between two levels of quality of service, ***high*** or ***low***. High-quality service is more costly to provide, and some of the cost is independent of whether the contract is signed or not. The level of service cannot be put verifiably into the contract. High-quality service is more valuable than low-quality service to the customer, in fact so much so that the customer would prefer not to buy the service if she knew that the quality was low. Her choices are: ***buy*** or ***don't buy*** the service.

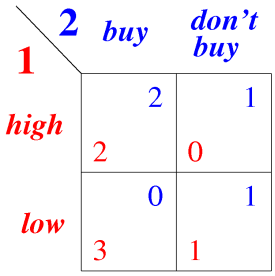

Figure 5. The quality game.

This game can be solved by *iterated dominance*, not unlike the price-setting game in Figure 1. For each choice of player 2, player 1 prefers ***low*** to ***high***, so ***low*** is his dominant strategy, and ***high*** is never played in an equilibrium. Given that player 2 knows this, she will only react to ***low*** and choose ***don't buy*** as her best response. The resulting payoff 1 for each player is worse than payoff 2 for the strategy pair (***high***, ***buy***), which is, however, not stable.

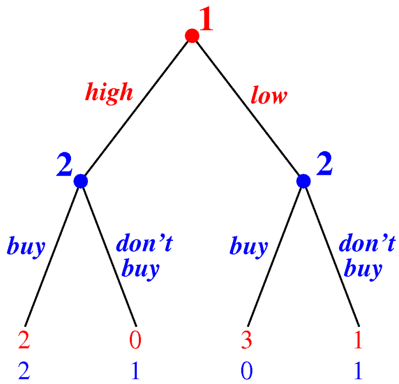

Figure 6. *Changed* quality game where player 1 can *commit* to his strategy..

Suppose this game can be *changed* so that player 1 moves first and can credibly *commit* to providing ***high*** (or ***low***) quality service. Then player 2 can *react* and choose a different move in response, as displayed in the game tree in Figure 6. Roll-back analysis then shows that player 2 will ***buy*** after player 1 chooses ***high*** (with payoff 2 to both players), and ***don't buy*** after player 1 chooses ***low*** (with payoff 1 to both players), so player 1 chooses ***high*** at the first stage – which is the desired outcome.

Here, the game has been changed by introducing *commitment power* for player 1. This change is beneficial for both players. This is not so in *bargaining*, where a player who can make a "take it or leave it" offer has all the power, as described in the "ultimatum game" (see, for example, Osborne, 2004). More refined models of bargaining (Rubinstein, 1982; Binmore, Rubinstein, and Wolinsky, 1986) describe rounds of alternating offers. Roll-back analysis for this model shows that a player who is more *patient* (in the sense of paying a lower "interest rate" for waiting) has greater *bargaining power* in the sense of getting a larger share of the object under negotiation.

Because the rules of the game matter so much for what can be achieved, a game-theoretic analysis often shows the value of *changing the game*. Brandenburger and Nalebuff (1996) have coined the term *co-opetition* to achieve an improvement for all players in business settings, with the help of game theory.

The spectacular success of game theory in defining the rules of the game has been demonstrated in *auction design* (Milgrom, 2004). In Britain, the 2000 sale of 3G spectrum licenses raised billions of pounds (about 2.5 percent of GDP). The auction was designed by game theorists, who also found it crucial to test their behavioral predictions by *laboratory experiments* (Binmore and Klemperer, 2002).

6. Getting and hiding information

Game theory also provides tools to describe and analyze different *information* that players have, which often determines their behavior. In settings without perfect information, especially when there are conflicting incentives, optimal play may prescribe some *unpredictability* in order to leave the opponent in the dark. This is known as a *mixed strategy* where a player chooses his action *randomly*.

In Figure 7, different *scenarios* are represented by a *chance move* at the first stage. A small technology startup company (player 2) has just developed a new product. The market leader (player 1) is also known to have worked on such a product, but it is only known to him if the product can be credibly announced as coming to the market soon (in which player 1 is in a **strong** position), or not (player 1 being **weak**). To the outsider, these two scenarios where player 1 is either weak or strong have equal probability 1/2. When player is in a weak position, he can decide to ***quit*** the competition, with payoff 0 to him and payoff 16 to player 2. He can also decide to "bluff" and *announce* that he has a competing product that is coming to market soon, which he would do anyhow if he has such a product and is therefore in a strong position.

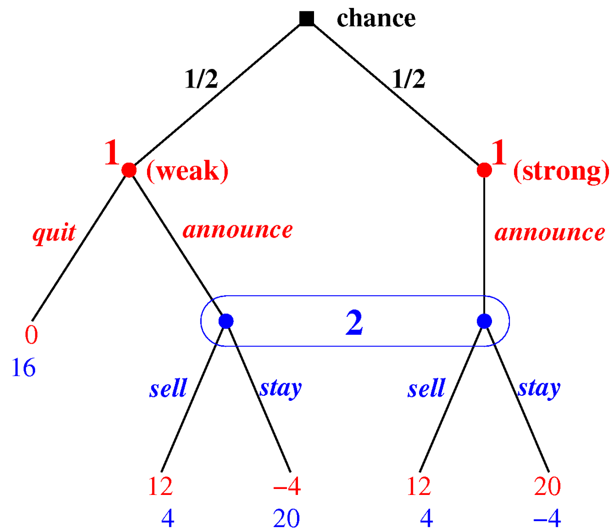

Figure 7. Imperfect-information game between a market leader (player 1) and startup company (player 2).

Player 2 does not know whether player 1 is weak or strong after getting the announcement. In Figure 7, the two decision nodes of player 2 are in an *information set*, indicated by the oval around the two nodes, which means that player 2 knows he is in one of the two nodes but not which one, and has to make the same move in both situations. The startup can either ***sell*** to the big company, who will then take over the product, or ***stay***, in the hope that player 1 is in fact weak. The two moves ***sell*** and ***stay*** have different payoffs to player 2 (and opposite payoffs to player 1) depending on whether player 1 is strong or weak.

1 \ 2	***sell***	***stay***
quit **(if weak)**	6 (player 1), 10 (player 2)	10 (player 1), 6 (player 2)
announce **(if weak)**	12 (player 1), 4 (player 2)	8 (player 1), 8 (player 2)

Figure 8. Strategic form of the game in Figure 7. Equilibrium requires randomized strategies.

The strategic form of this game is shown in Figure 8. The numbers in the cells are *expected payoffs* calculated from the chance probabilities and the terminal payoffs resulting from the two strategies for that cell. For example, the combination of ***quit*** and ***sell*** gives payoff 1/2×**0** + 1/2×**12** = **6** for player 1, and 1/2×**16** + 1/2×**4** = **10** for player 2.

Interestingly, no cell in this game is an equilibrium, because the best responses are circular. If player 1 chooses ***quit***, then player 2 knows that an announcement comes only from player 1 in strong position and prefers to ***sell***. Against ***sell***, the best response for player 1 is to ***announce*** even in the weak position. If player 1 does that, however, then player 2 knows that player is bluffing half of the time, where on average it is better to ***stay***, but then player 1 would prefer to ***quit*** in response.

Optimal play in this game is the following randomized behavior of the players (for details see Turocy and von Stengel, 2002). Player 1 chooses ***quit*** and ***announce*** exactly half of the time, so that player 1 will get expected payoff 7 for either ***sell*** or ***stay***. Because player 2 is now indifferent between his two actions, he can also randomize as a best response, and will choose ***sell*** with probability 1/4 and ***stay*** with probability 3/4. Only then player 1 will get 9 for either ***quit*** or ***announce***, so that his randomization is also a best response.

In summary, conflicting moves with lack of information may require randomized actions. These randomized actions preserve the lack of information which is necessary for randomization to be optimal, and thus for obtaining an equilibrium. The probabilities used in randomized actions can sometimes also be seen as fractions of a large population adopting the chosen actions, as already discussed for the bandwidth choice game in Figure 4.

7. Selfish routing in networks

Network routing games are another class of games with a large number of individuals, and where one studies which fractions of these individuals are choosing particular actions. The actions are here routes in a transport or communication network. With the decentralized nature of the internet and its protocols, packet routing can indeed be considered as a such game (Akella et al., 2002).

The network is specified by points and *links* connecting them, a certain amount of *flow* that has to be routed between given source and destination points, and *delay functions* that describe the delay for each link depending on how much flow is routed through that link. Individuals always seek the route with shortest delay until they cannot do better and an equilibrium results (Wardrop 1952). In the internet, routing algorithms also choose or predict the best route without central coordination.

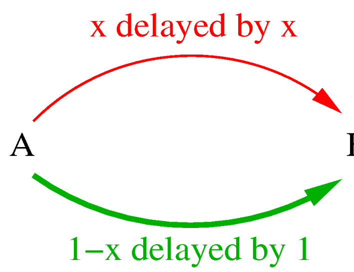

Figure 9. Pigou's network routing game with two links and flow 1 to be routed from A to B.

A game-theoretic analysis shows that the equilibrium flow resulting from this "selfish routing" (roughly corresponding to "best effort" in internet parlance) is typically worse than optimal flow if network traffic was centrally managed. The simplest example, due to the economist Pigou, is shown in Figure 9. Points A and B are connected by two links, and a flow of 1 is to be routed from A to B. If the top link (in red) takes flow x, then users of that link experience delay x. The remaining flow 1-x through the bottom link (in green) has constant delay 1 independent of the flow. Optimal flow, which has the smallest average delay, results for x = 1/2 were users of the top link have delay 1/2, and those on the bottom link have delay 1, so that the average delay is 1/2×1/2 + 1/2×1 = 3/4. However, this flow is not in equilibrium: all users on the bottom link would switch to the top link which has less delay as long as x<1, so the unique equilibrium flow has x=1 with everyone experiencing delay 1, which is a third larger than 3/4. (Roughgarden and Tardos, 2002, have shown that this is the worst possible ratio between equilibrium and optimal flow in any network where delay functions are linear.)

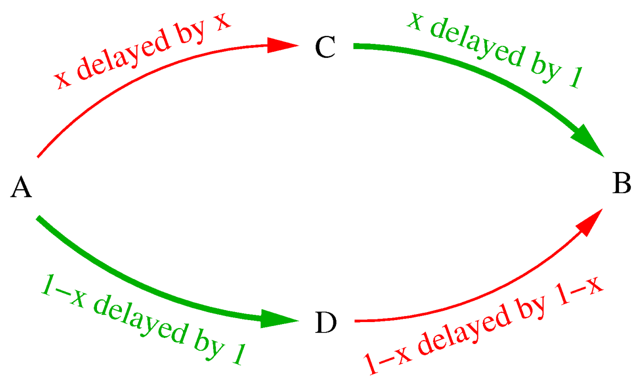

Figure 10. Network with flow 1 to be routed from A to B where selfish routing is optimal.

Equilibrium flow due to selfish routing and optimal managed flow can be the same. Figure 10 shows a network with two routes from A to B, where the total flow 1 is split into x and 1-x. The top and bottom route have delays x+1 and 2-x. Equilibrium results when these are equal, for x=1/2, which is also the optimal flow.

Figure 11 shows the same network as Figure 10, except that its capacity is increased by introducing a zero-delay link (in blue) between C and D, which now takes flow z away from x at point C. In effect, these are now two Pigou networks as in Figure 9 put in sequence. The first Pigou-style network has two routes from A to D, the top route via C of delay x, and the bottom green link of delay 1. A second Pigou network connects C to B via two routes, the top green link from C to B of delay 1, and the bottom link from C via D to B where the only delay is along the red link whose delay is exactly the flow through that link. Now the only equilibrium flow is given by x=1 and z=1 where all individuals take the shortcut through the new blue link. As in Pigou's network, they fully congest the red links and do not use the green links any more. The average delay has increased from 1.5 to 2.

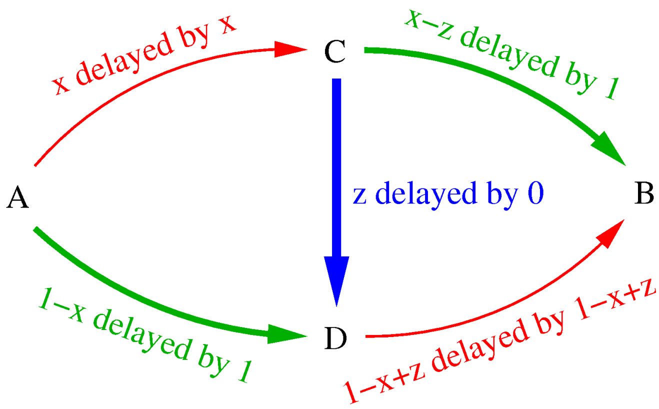

Figure 11. The network from Figure 10 with an extra link of zero delay. Network capacity increases but delay of equilibrium flow worsens (Braess's paradox).

The comparison between the networks in Figures 10 and 11 is interesting because it shows that *increasing network capacity can create worse congestion* in a routing equilibrium, which is known as Braess's paradox (Braess 1969).

In traffic networks, it is usually assumed that there is a certain traffic to be routed, and network capacity should be found to meet that demand. Braess's paradox shows that increasing link capacity should not be done blindly but in concert with equilibrium considerations. In communication networks like the internet, new network capacity creates new demand which is often vastly larger than before. It is an interesting research question to analyze such an increase in demand, and possible effects on congestion, with the help of game theory.

8. Summary

Game theory is a tool for modeling and analyzing interactive decisions. It can be applied to situations with a small number of players, or to population dynamics with a large number of similar players. It is also very suitable to model competitive dynamics of markets. In business decisions, game theory enhances the ability to think ahead and to consider the perspective of one's opponent. A game-theoretic analysis demonstrates the importance of individual incentives, which can lead to unexpected global behavior, of the information that players have, and of the rules of the game. The solutions offered by game theory, in particular the central concept of equilibrium, are not always unique, and this can be an important insight of the analysis. The internet with its selfish users and ever-growing economic importance is a particular rich domain for game-theoretic studies.

References

A. Akella et al., Selfish behavior and stability of the internet: a game-theoretic analysis of TCP. Proc. ACM SIGCOMM '02, 117–130, 2002.

K. Binmore and P. Klemperer, The biggest auction ever: the sale of the British 3G telecom licenses. The Economic Journal 112, C74–C96, 2002.

K. Binmore, A. Rubinstein, and A. Wolinsky, The Nash bargaining solution in economic modelling. Rand Journal of Economics 17, 176–188, 1986.

D. Braess, Über ein Paradoxon aus der Verkehrsplanung. Unternehmensforschung 12, 258–268, 1969.

A. M. Brandenburger and B. J. Nalebuff, Co-Opetition. Currency Doubleday, 1996.

A. K. Dixit and B. J. Nalebuff, Thinking Strategically. Norton, 1991.

H. Lindstädt and J. Müller, Making game theory work for managers. McKinsey Quarterly, 1–9, January 2010.

J. Maynard Smith, Evolution and the Theory of Games. Cambridge University Press, 1982.

P. Milgrom, Putting Auction Theory to Work. Cambridge University Press, 2004.

M. J. Osborne, An Introduction to Game Theory. Oxford University Press, 2004.

T. Roughgarden and E. Tardos, How bad is selfish routing? Journal of the ACM 49, 236–259, 2002.

A. Rubinstein, Perfect equilibrium in a bargaining model. Econometrica 50, 97–110, 1982.

W. H. Sandholm, Population Games and Evolutionary Dynamics. MIT Press, 2010.

R. Selten, The chain store paradox. Theory and Decision 9, 127–159, 1978.

T. L. Turocy and B. von Stengel, Game theory. Encyclopedia of Information Systems, Vol. 2, Elsevier Science (USA), 403–420, 2002.

J. G. Wardrop, Some theoretical aspects of road traffic research. Proc. Institute of Civil Engineers, Part II, Vol. 1, 325–378, 1952.