Economics of Strategy for Online and Digital Markets

Topics in Economic Theory and Policy, 31C01000

Iivo Vehviläinen

January 10, 2019

Aalto University
iivo.vehvilainen@aalto.fi

Lecture 2: Prerequisites

Outline

- Preferences
- Game theory
- Pareto efficiency

Preferences

Several approached for preferences, but somewhat standard is *revealed preference theory* by Samuelson which we will follow:

- When given a choice, it is *assumed* that people behave consistently
- This consistency can arise from, for example:
  - Preference of one item over another
  - Limits on the choice, by budget, but also by e.g. conviction, laws and regulation, or social norms
  - Habits, for example sleep during the night, curioisity, laziness
- It can then be shown that people then behave as though they are maximizing something which economists call *utility*.

Utility

- An utility function assigns a numeric value for each choice.
- In economic theory, we often start with a specific functional form for the utility function.
- An example is the popular Cobb-Douglas utility function:

  $$u(x, y) = x^a y^{1-a},$$

  which was crafted to explain how productivity depends on labor and capital inputs.
- The validity of such mapping of choices to numeric values will depend on the task at hand, and is an empirical question.

Substitution and indifference

- One cannot have everything: need to make trade-offs
- Trade-offs are made between the choices that affect utility.
- Same utility can be achieved with different choices:
  - E.g. in the Cobb-Douglas utility capital input to automatization can be a substitute to labor input.
  - Or sometimes quality can be substituted with quantity, time with money etc.
- We say that one is indifferent between two choices if the utility of choosing either of them is the same.

Substitution and indifference

The (fictitious!) student here assigns some utility to both the amount of free time and the grade from a course.
In the illustration the student:

- is indifferent between choices A to D.
- prefers A over B.
- prefers D over C.

	A	E	F	G	H	D
Hours of free time	15	16	17	18	19	20
Final Grade	84	75	67	60	54	50

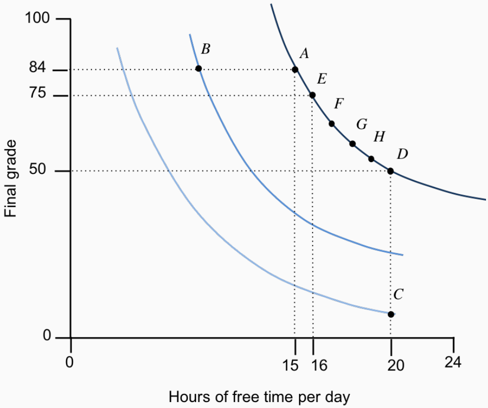

Source: CORE.

Opportunity cost

- Maximum grade is constrained e.g. by your talent and the amount of time you choose to study.
- Spending less time for study gives more free time but incurs an *opportunity cost* in lower grade.

	A	E	C	F
Free time	13	14	19	20
Grade	84	81	57	50
Opportunity cost		3	7	

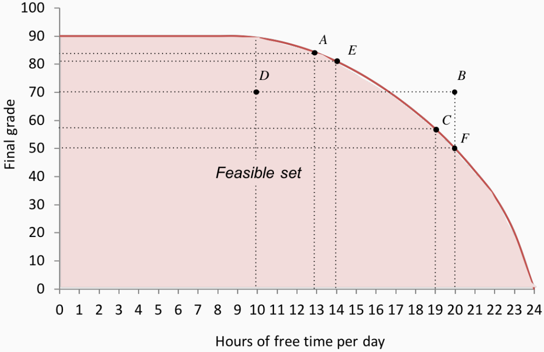

Source: CORE.

Optimal utility

Utility is maximized when marginal rate of substitution (MRS) is equal to the marginal rate of transformation (MRT) set by the constraints.

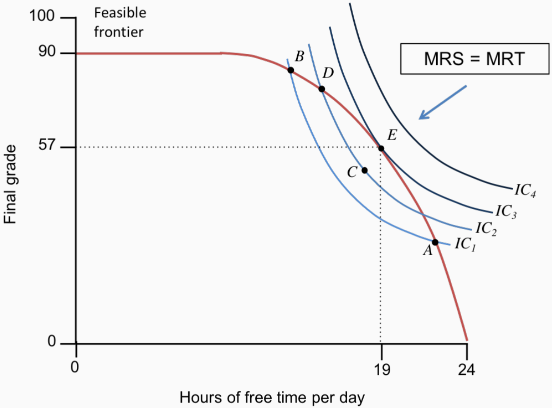

Source: CORE.

From preferences to demand

- In every-day life people do not knowingly optimize their utility.
- The preference model is still a useful approximation:
  - Over time people adopt some consumption patterns.
  - The model is used successfully in microeconomics to explain many observed market phenomena.
  - Sometimes the simplifications in the basic model can become restrictive and other tools, such as behavioral economics, can be useful.
- Of particular interest in this course are the private preferences that determine your utility of consumption, and the aggregated market demand of all consumers.

Example: Preferences, substitutes, and demand

Table 1

Reduced Form Estimates

Bandwidth	p128	p96	p64	p32	p16
128	*−2.0*	+0.80	+0.25	−0.02	−0.16
96	*+1.7*	*−3.1*	*+43*	+0.19	+0.18
64	+0.77	*+1.8*	*−2.9*	+0.59	+0.21
32	+0.81	−1.0	*+1.0*	*−1.4*	+0.15
16	+0.2	−0.29	+0.04	*+1.2*	*−1.3*

Notes. Italics indicate significance at the 5% level. All own-price effects are significantly negative; the cross-price effects for one-step lower bandwidths are positive.

Figure. Demand elasticities by bandwidth.

Source: Varian, 2012.

Recovering preferences

- There are two key motivations to find out your preferences:
    1. Understand the drivers behind your choices.
    2. Predict how you will behave in a given situation.
- The first part, causal interference, is a traditional field of economics (microeconometrics).
- The second part, prediction, has become popular in data science with the availability of big data.

Estimation and prediction

- Let's consider a model

$$Y_i = f(X_i) + \epsilon_i$$

- Here X_i and Y_i are observed values.
- Function f maps values of X to Y and ϵ_i are the random noise (in measurement, unobserved variables, etc.).

Estimation and prediction

The goals in estimation and prediction are different:

1. The goal of regression estimation is unbiasedness:

$$E[\hat{f}] = f.$$

2. The goal of prediction is to minimize prediction error:

$$\hat{f} = \min_{f \in \mathcal{F}} L(f) = \min_{f \in \mathcal{F}} E[\ell(f(x), y)],$$

where e.g. $\ell(f(x), y) = (y - f(x))^2$.

Estimation and prediction

The expected prediction error is:

$$\begin{aligned} \text{error} &= E[(y - \hat{f}(x))^2] \\ &= E[(y - f(x))^2] + (E[\hat{f}(x)] - f(x))^2 \\ &\quad + E[(\hat{f}(x) - E[\hat{f}(x)])^2] \end{aligned}$$

Which can be written as

$$\text{error} = Var(y) + Bias^2 + Var(\hat{f}(x))$$

i.e. a combination of variance in data ($Var(y)$), bias in estimation, and variance from the fact that a sample is used to estimate $\hat{f}$ ($Var(\hat{f}(x))$).

Example: Predicting student numbers

	code	group	class
1	363651	11	2015
2	608436	11	2016
3	431779	11	2015
4	489072	8	2018
5	280478	3	2011
6	369920	4	2018
7	444360	7	2018
8	536385	11	2015
9	601187	12	2016
10	555559	11	2016

Table 1: Sample of data (code here randomized).

Example: Predicting student numbers

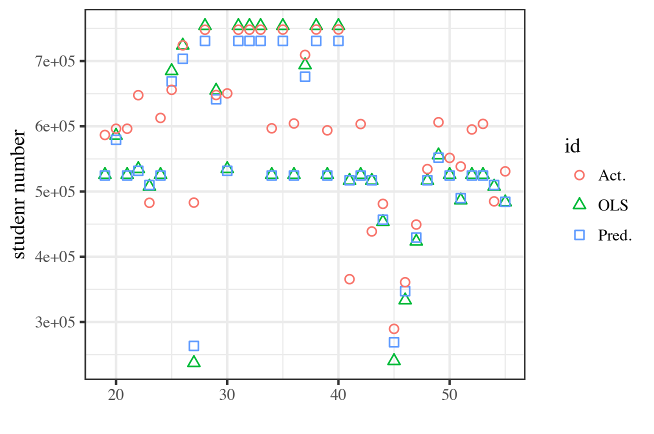

Figure. Average prediction error in OLS 5 % higher than with prediction.

Experiments and A/B testing

- Almost all large consumer websites carry out A/B tests.
- The idea is to take a random sample of users who are shown a version B of the website while the the rest use version A.
- Comparing results (e.g. clicks, purchases) between A and B can *predict* which version is better.
- These are very similar to controlled experiments used in economics, medicine etc., just the objective is different.
- For a website it may be enough to see which version works better (for now), but for other decisions it can be crucial to understand what are the reasons for the differences.

Example: Prediction gone wrong

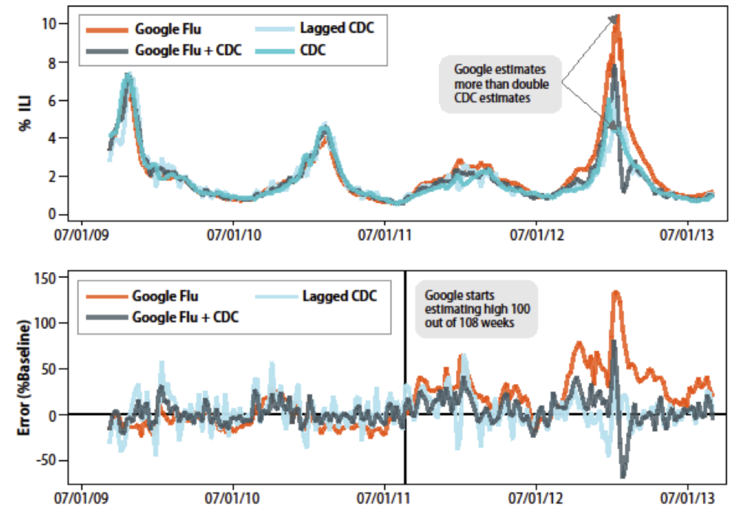

Source: Lazer et al. 2014.

Game theory

- So far discussion on the behavior of an individual.
- In a market place the interactions of agents (firms) often important in practice: How will the actions of the others affect my decisions?
- Game theory has been hugely successful in explaining these interactions.

Game theory

- We make a difference between:
  - the behavior that does not take into account the actions of the others (non-strategic, sometimes price taker) and
  - the behavior where the other actions are accounted for (strategic, price setting).
- Game theory considers the modeling of the strategic interactions.
- Typically, in a market setting, the focus is on the behavior of firms that wield strategic influence over the market.

Illustration: War of attrition

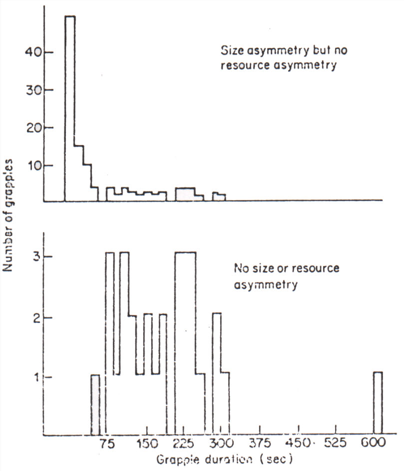

Figure. Length of grapple between spiders.

Source: Austad, 1982.

Illustration: War of attrition

- Two firms, i and j, own plants K and k, $K > k$.
- Demand is defined by an inverse demand function $p(Q)$.
- Marginal cost of production for both plants is c.
- Payoffs defined with:

$$\pi^i = \pi(K, k) = [p(K + k) - c]K$$
$$\pi^j = \pi(k, K) = [p(K + k) - c]k$$

Source: Liski & Vehviläinen, 2018.

Illustration: War of attrition

- Assume that the payoffs of the firms are defined as

(π^i, π^j)	$k = 0$	$k = 1$
$K = 2$	(45, 0)	(20, 10)
$K = 0$	(0, 0)	(0, 90)

Source: Liski & Vehviläinen, 2018.

Illustration: War of attrition

- In a *Nash equilibrium* both firms keep their plants running.

(π^i, π^j)	$k = 0$	$k = 1$
$K = 2$	(45, 0)	(20, 10)
$K = 0$	(0, 0)	(0, 90)

- If the firms could coordinate their actions, there would be more money on the table to be shared ($90 > 20 + 10$).

Source: Liski & Vehviläinen, 2018.

Game theory

- In a strategic interaction game, n players who take action a_i.
- This will give the players a payoff that depends on their own action and the actions of the others: $u_i(a_1, \ldots, a_i, \ldots, a_n)$.
- An action a_i is a *dominant strategy* for player i if .

$$u_i(a_1, \ldots, a_i, \ldots, a_n) \geq u_i(a_1, \ldots, \tilde{a}_i, \ldots, a_n), \forall \tilde{a}_i \neq a_i$$

- Dominant strategy optimal for i regardless of what others do.
- A *Nash equilibrium* means that each player i's action is optimal given the actions of the other players.

In-class exercise: Game theory

Set-up:

- Prof. Quackenbush has graded a midterm exam for his class.
- Before handing back the exam, he calculated the average, 80, but forgot to record the students' individual grades.
- He now has no choice but to ask the students their scores.
- Class has 3 students whose actual grades are 60, 80, and 100.

Your job is to construct a "payoff mechanism" that should encourage each student to tell the truth regarding their grade on the exam, under the assumption that the other students are reporting their grades truthfully.

Pareto efficiency

- When dust settles in the market each agent is takes some action that results in the *allocation* of resources between the agents.
- An allocation is Pareto efficient if nobody can be better off without making somebody worse off.
- In general, little more can be said about how “good” an allocation is compared to other possibilities.

Exit game example of Pareto efficiency

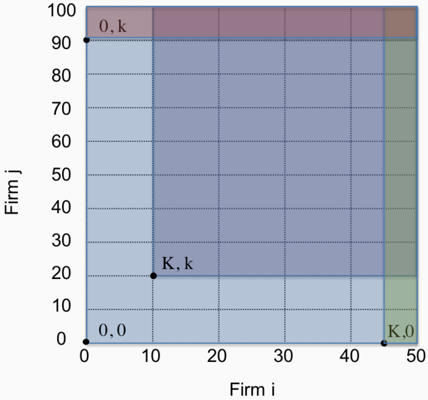

Outcomes better for both than 0, k

Outcomes better for both than K, 0

Outcomes better for both than K, k

Outcomes better for both than 0, 0

0, 0 = Both plants off

K, 0 = Bigger plant on, smaller plant off

0, k = Smaller plant on, bigger plant on

K, k = Both plants on

Learnings today

- Revealed preferences model is a tool to explain observed market behavior of individuals.
- Game theory models the strategic interactions of firms.
- A Nash equilibrium is a state where all players in a game make the best possible decisions for them while taking the actions of others in to account.
- Pareto efficiency can be used as a criteria to describe if everyone is as happy as possible in the current market state given the utility of others.

Readings for this lecture

Following units from *The Economy* (www.core-econ.org) should be helpful:

- Preferences (CORE 3.2-3.5)
- Game theory (CORE 4.1-4.3)
- Pareto efficiency (CORE 5.2)

Note that these units from CORE may not necessarily be self-sufficient. If you are unfamiliar with the topics covered you may want to familiarize yourself with the earlier CORE units. Do not be overwhelmed by the amount of material: it is extensive but for the purposes of this course should make a quick read.

Exercises for Lecture 1

1. Come up with at least one example of both
   (a) an online market place
   (b) a digital platform

   that you would be interested to hear about later during the course.
2. List the top 5 reasons why you think Amazon has been so successful.

Exercises for Lecture 2

3. Suppose that you are in a supermarket buying ice cream. A package costs 2.66 €, but there is a bundle of two offered at 4.00 €. If you are indifferent between the two choices:
   (a) What does that tell you about your preferences?
   (b) Can you tell at what price you would be indifferent with three packages of ice cream? Why?
4. You are going to a concert with a friend, and tickets are about to come online any minute. A single ticket costs 40 € and there is a fixed delivery fee of 10 € per purchase (i.e. not per ticket). You forgot to communicate beforehand and need to make decisions on who buys and how many tickets on the spot. You share the total bill. If you miss the online sale, you must go to the black market at a cost of 100 € per ticket.
   (a) Model the costs for one player as in a strategic interaction game.
   (b) How many tickets you end up buying in a Nash equilibrium of the game? (Hint: Use (a) and the symmetry of the problem.)
   (c) Can you improve the outcome? If yes, then how?

Submit your answers through MyCourses by noon Thu 17 Jan.

Exercise for Lecture 5

1. Think of a question (about economics of games) to Janne Peltola / Supercell.
   You get an extra point if we use the question for in-class discussion.

Submit your question through MyCourses by noon Thu 17 Jan.

Next time

Markets

- Institutions
- Supply and demand
- Competitive equilibrium
- Perfect competition, monopoly, oligopoly