CSC200: Lecture 27

Announcements and Outline

- Announcements
  - I have added questions to and finalized Assignment 3.
  - Upcoming due dates:
    1. Quiz 6 is on February 5 (scope to be announced next Monday)
    2. Assignment 3 is due on February 10
    3. Term test 2 is on February 12
    4. February 15-19 Reading Week

- Todays agenda
  1. The calculation of GSP and VCG prices for the example in next to last slide of Lecture 26 when the true click values are $(7, 6, 2)$.
  2. New topic: Influence spread and more generally how individual behaviour promotes collective behaviour. (Chapters 16, 17, 18, 19)

GSP expected revenue in NE for text example

- In GSP, bidding truthfully (at 7 per click) x will win slot a and have utility $70 - 60 = 10$ whereas by lowering bid to 5 per click, x will win slot b obtaining utility $28 - 4 = 24$. Thus GSP is not truthful.
- The (x, y, z) bids $(5, 4, 2)$ achieve revenue $10 \times 4 + 4 \times 2 = 48$. This is a socially optimal NE. What needs to be shown?
- The (x, y, z) bids $(3, 5, 2)$ achieve revenue $10 \times 3 + 4 \times 2 = 38$. This is an NE that is NOT socially optimal
- VCG achieves revenue 44; $64 - 24 = 40$ from x and 4 from y.

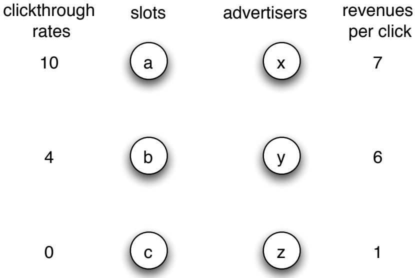

[Fig 15.6, E&K]

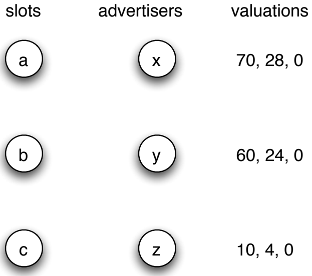

[Fig 15.7, E&K]

What needs to be shown in these examples?

1. Determining social optimality is immediate as that is obtained if and only if [x gets slot a, y gets slot b and z gets c].
2. We have shown the calculation of the revenue for the examples. Questions?
3. Truthfulness implies that bidding truthfully is an NE.
4. We need to show that $(5, 4, 2)$ and $(3, 4, 2)$ are NE. Lets consider $(5, 4, 2)$. What has to be shown?

What needs to be shown in these examples?

1. Determining social optimality is immediate as that is obtained if and only if [x gets slot a, y gets slot b and z gets c].
2. We have shown the calculation of the revenue for the examples. Questions?
3. Truthfulness implies that bidding truthfully is an NE.
4. We need to show that $(5, 4, 2)$ and $(3, 4, 2)$ are NE. Lets consider $(5, 4, 2)$. What has to be shown?
   - Assumming y bids 4 and z bids 2, need to show that x cannot benefit by reducing bid sufficiently that x gets slot b (i.e. $2 \leq bid_x \leq 4$) or x gets slot c (i.e. $0 \leq bid_x \leq 2$). For example, if x bids so as to wind up in slot b, then utiltiy decreases from $70 - 40 = 30$ to $28 - 8 = 20$.
   - Assuming x bids 5 and z bids 2, have to show that y cannot benefit by bidding to get slot a (i.e. $bid_y \geq 5$) or y gets
   - z cannot benefit by increasing the bid given bids of x and y. slot c (i.e. $bid_y \leq 2$).

New Topic

- Beginning of a several week discussion of
  - influence/technology/disease spread
    i.e. "contagion" in a very general sense in social networks;
- We will first be discussing Chapter 16 (information cascades) and then Chapter 17 (direct benefit effects).
- Chapter 18 discusses "rich get richer" models and power laws
- We then study influence in a social network (chapter 19).
- That is, for the next several weeks we will be studying

various social processes that channel individual behaviour into collective behaviour.

Information cascades (herding) - Chapter 16

- Chapter 16 concerns the phenomena of information cascades
  - whereby individuals observe and then make decisions sequentially based on the behaviour of people having made decisions earlier;
  - e.g. deciding on a restaurant by observing how many people are currently eating there, what clothes you buy, other fashions/fads.

- Chapter 17 discusses decisions based on direct benefit (e.g., using a popular operating system/ laptop because wide use implies more software support).

- Clearly both phenomena can be interacting when people make decisions (e.g. busy restaurants are more able to use fresh ingrediants); the text organization is to try to first isolate and model these phenomena so as to gain insight.

A simple information cascade model

Assumptions for an information cascade:

- Individuals make decisions sequentially and can observe the decision of those who have acted earlier.
- Each individual has some private information that can be used in making their decision.
- Individuals only observe the behavior of earlier people but do not know their private information beyond any inferences that can be made from the previous decisions/actions.

The majority balls and bins example

- In the following experiment we will observe the formation of very uniform (albeit perhaps fragile, perhaps wrong) cascades.

- The majority blue vs majority red *balls in a bin* experiment.
  - Assume we have a bin which (with equal probablility) either has 2 red balls and one blue ball, or has 2 blue balls and one red ball.
  - Each individual in turn blindly (i.e. randomly) picks out a ball from the bin, observes its colour, and then places it back into the bin.
  - Now each individual votes (i.e. announces their decision as to which colour is in the majority) based on what they have observed and what others have previously reported.

- In a different experiment, one can ask what would happen if one votes their private info unless there is a "clear majority"? What is a clear majority? Following the Supreme Court in decision regarding a Quebec independence vote, clear majority is not defined.

Possible behaviour for the majority balls experiment

- In the following experiment we will observe the formation of definite (albeit perhaps fragile, perhaps wrong) cascades.

- The majority blue vs majority red *balls in a bin* experiment.
  - Assume we have a bin which (with equal probablility) either has 2 red balls and one blue ball, or has 2 blue balls and one red ball.
  - Each individual in turn blindly (i.e. randomly) picks out a ball from the bin, observes its colour, and then places it back into the bin.
  - Now the experiment assumes that one votes (i.e. announces their decision as to which colour is in the majority) their own private information (the colour of the ball drawn) *unless* the difference in previous colours drawn is at least 2 in which case they vote with the majority.

- In a different experiment, one can ask what would happen if one votes their private info unless there is a "clear majority"? What is a clear majority? Following the Supreme Court, not defined.

Conditional probability and what we tend to believe

- Why in the simple balls and bin text example (with 3 balls), would we vote for blue if the first two decisions were blue even if we saw a red ball?

- Simply stated, conditioned on the information provided by these two decisions, we can infer that it is more likely that the urn contains 2 blue and 1 red ball. We now want to make precise such a statement in the language of probability.

- We have the following concepts and notation:
  - $Pr[A]$ = the "prior prob." of event A; e.g. urn has 2 blue,1 red
  - $Pr[A|B]$ = the "posterior or conditional prob." of event A given that event B occurred; e.g. B is the event that I drew a red ball and the previous decisions were $D_1, D_2, \ldots$

Bayes rule

- By definition of conditional prob., $Pr[A|B] = Pr[A \text{ and } B]/Pr[B]$
- By definition of conditional prob., $Pr[B|A] = Pr[A \text{ and } B]/Pr[A]$
- Therefore $Pr[A|B] \times Pr[B] = Pr[B|A] \times Pr[A]$
- So we have **Bayes Rule**

$$Pr[A|B] = Pr[A] \times \frac{Pr[B|A]}{Pr[B]}$$

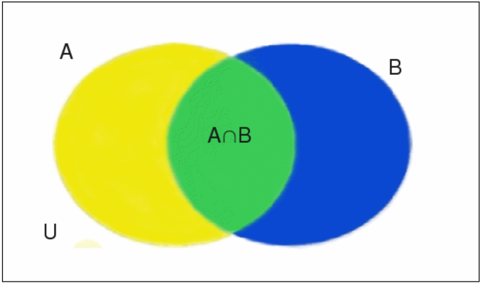

Figure : Two events A and B in a sample space, and the joint event $A \cap B$.

The black and yellow taxi example

- Eyewitness of a hit and run reports taxi was yellow
- Let A be the event that the true colour is yellow.
- Let A' be the event that the true colour is black.
- Let B is the event that the witness reports taxi was yellow.
- The desired probability is the posterior prob $Pr[A|B]$
- We are given that the prior prob $\text{Pr[A]} = .2$;
- Also given that this witness mixes up colours 20% of the time:

$$Pr[\text{reports yellow} \mid \text{taxi yellow}] = Pr[\text{reports black} \mid \text{taxi black}] = .8$$

- Bayes rule gives

$$Pr[A|B] = Pr[A] * Pr[B|A]/Pr[B]$$

- We have all the information we need here except the denominator: the prior probability $Pr[B] = Pr[\text{report yellow}]$; this can happen if

The black and yellow taxi example

- Eyewitness of a hit and run reports taxi was yellow
- Let A be the event that the true colour is yellow.
- Let A' be the event that the true colour is black.
- Let B is the event that the witness reports taxi was yellow.
- The desired probability is the posterior prob $Pr[A|B]$
- We are given that the prior prob $\Pr[\mathrm{A}] = .2$;
- Also given that this witness mixes up colours 20% of the time:

$$Pr[\text{reports yellow} \mid \text{taxi yellow}] = Pr[\text{reports black} \mid \text{taxi black}] = .8$$

- Bayes rule gives

$$Pr[A|B] = Pr[A] * Pr[B|A]/Pr[B]$$

- We have all the information we need here except the denominator: the prior probability $Pr[B] = Pr[\text{report yellow}]$; this can happen if
  1. the taxi is really yellow and witness reports yellow or

The black and yellow taxi example

- Eyewitness of a hit and run reports taxi was yellow
- Let A be the event that the true colour is yellow.
- Let A' be the event that the true colour is black.
- Let B is the event that the witness reports taxi was yellow.
- The desired probability is the posterior prob $Pr[A|B]$
- We are given that the prior prob Pr[A] = .2;
- Also given that this witness mixes up colours 20% of the time:

$$Pr[\text{reports yellow} \mid \text{taxi yellow}] = Pr[\text{reports black} \mid \text{taxi black}] = .8$$

- Bayes rule gives

$$Pr[A|B] = Pr[A] * Pr[B|A]/Pr[B]$$

- We have all the information we need here except the denominator: the prior probability $Pr[B] = Pr[\text{report yellow}]$; this can happen if
  1. the taxi is really yellow and witness reports yellow or
  2. the taxi is really black and witness reports yellow

Eye witness and the black & yellow taxi example

- We are trying to compute the $\Pr[B] = \Pr[\text{reports yellow}]$ which is

$$\Pr[\text{reports yellow}|A] \times Pr[A] + Pr[\text{reports yellow}|A'] \times Pr[A']$$
$$= .8 * .2 + .2 * .8 = .32$$

- Returning to Bayes rule:

$$Pr[A|B] = Pr[A] * Pr[B|A]/Pr[B] = .2 * .8/.32 = .5$$

- This is the same thing as if someone just guesses; that is,
  1. knew nothing about the distribution of yellow and black taxis (and therefore assumes $\Pr[A] = .5$) and
  2. didn't see anything (and therefore has no information upon which to condition the probability) and reports yellow by flipping a fair coin.
- Another way to say it: the taxi is equally likely to be yellow or black.
- The fact that the witness reports yellow has, however, significantly raised the chances of the taxi being yellow (from .2 to .5)

Another example: spam filtering

- We receive an email with subject heading that more or less says "check this out". Should we treat this as spam?
- Let A be the event that the message is spam.
- Let A' be the event that the message is not spam
- Let B is the event that the message heading is "check this out".
- We are given prior probability $Pr[A] = .4$, $Pr[B|A] = .01$ and $Pr[B|A'] = .004$. We want to compute $Pr[A|B]$.
- By Bayes Rule,

$$Pr[A|B] = Pr[A] \times Pr[B|A]/Pr[B]$$

- We need to calculate $Pr[B]$ and as before:

$$Pr[B|A] \times Pr[A] + Pr[B|A'] \times Pr[A'] = (.01) \times (.4) + (.004) \times (.6) = .0064$$

- Hence $Pr[A|B] = Pr[A] \times Pr[B|A]/Pr[B] = .625$