ORIE 6180: Design of Online Marketplaces	Spring 2016

Lecture 19— April 18

Lecturer: Sid Banerjee *Scribes: Gutekunst*

19.0 Caveat Emptor

There was no scribe assigned for this lecture, so these notes were largely live-TeX'd. Typos may be abundant.

19.1 Overview of the last lecture

Last time non-auditing students gave short talks about their projects.

19.2 Overview of this lecture

In this lecture, we begin our unit on Reputation and Feedback. One fundamental concept we'll discuss is *information asymmetry.* In particular, in many markets sellers and buyers have different amounts of information. A classic study is [3], where this asymmetry is discussed in the context of car dealerships: a car dealer tends to have more information about the car being sold then a potential customer. Conversely, in insurance markets, a customer buying health insurance might have more information then the company they are considering purchasing insurance from (e.g., about their health or their strong inclination to go paint-balling while sky-diving). In online markets such as Amazon, the platform might control how information about a seller's reputation is conveyed to prospective customers through, e.g., customer reviews.

Background can be found in [1], and today is focused on the setting in [2].

19.3 Introduction

There are two specific types of information asymmetry we'll highlight. In *adverse selection*, information itself is hidden (e.g., about a customer's health). Conversely, *moral hazard* refers to situations where information is not hidden but actions themselves are (e.g., a customer's action to go paint-balling while sky-diving). We'll discuss a simple model from [2] that discusses both of these situations in the framework of reputation in online markets (Amazon, eBay, etc).

19.4 The Model

Our model is as follows:

- There is a single seller with one type of product. The seller has a type $\theta \in \{b, g\}$, denoting whether they are bad or good, respectively. We let $0 \leq b \leq g \leq 1$.
- There are T periods. During each period, a new buyer arrives and the seller can sell at most one item to that buyer. If the buyer decides to enter into a transaction, there are two outcomes. If a sale occurs, it is worth 1 to a buyer; if no sale occurs, it is worth 0 to that buyer.
- The seller puts in an effort $e \in [\underline{e}, 1]$ with $\underline{e} > 0$. If a transaction occurs, and given both the sellers type and effort, the probability that a match is successful is θe.
- The seller pays some cost $c(e)$ for putting in effort e. We assume that $c(\underline{e}) = c'(\underline{e}) = 0$ and $c'(e) > 0$ for $e > \underline{e}$. We also assume that $c''(e) > 0$ so that c is convex, and that $c'(1)$ is as large as we need it to be. Think of c as a convex function that goes from 0 to infinity, though that is not strictly necessary.
- Each transaction has a price p and both buyers and sellers are risk neutral. Sellers have a discount factor δ

Putting the above together, we have equations for the utilities and welfare:

$$u_b = (e\theta - p)_+, \quad u_s = (p - c(e))_+, \quad w = e\theta - c(e),$$

where $(\circ)_+ := \max\{0, \circ\}$. Assuming that $c'(1)$ is sufficiently large, we can solve for the effort e^*_θ that maximizes w, which is the solution to:

$$\theta = c'(e^*_\theta)$$

(and our assumption guarantees that such a solution exists). Notice that $b \leq g$ and our assumptions on c imply that $e^*_g \geq e^*_b$, and we'll assume that $e^*_b > e$.

At each time t, we'll consider

$$\mu_t := \mathbb{P}[\theta = g | t \text{ outcomes so far}]$$

to be the probability that a seller is good given what occurred in the past t transactions. In particular, you can imagine a customer being given a table of whether or not previous sales were successful, and using this table to update there prior on whether or the seller is good.

Using this model, we can address the information asymmetry cases described above.

19.5 Adverse Selection: How do people learn types?

In this situation, the type of the seller is hidden but no actions are hidden. We encode this by assuming $\underline{e} \equiv 1$ so that the "action" of choosing how much effort to put in is not hidden. So that hiding the seller's type is meaningful, we'll assume $b < g$. Let S_t be the number of successful transactions up to and including t, and let $F_t = t - S_t$ denote the number of failed transactions. We assume that both of these are known at time t, so that

$$\mu_t = \mathbb{P}[\theta = g | S_t].$$

We'll also assume that $\mu_0 > 0$ so that the first buyer comes in believing that there is at least some chance of the seller being good.

Using Bayes' rule, we can update our prior to get

$$\mu_1(S|\mu_0) = \frac{\mu_0 g}{\mu_0 g + (1-\mu_0)b}, \quad \mu_1(F|\mu_0) = \frac{\mu_0(1-g)}{\mu_0(1-g) + (1-\mu_0)(1-b)},$$

which denote the posterior probabilities of a successful or failed that a customer who arrives at time $t = 2$ holds. It turns out that μ_t converges to the truth as $t \to \infty$ in the sense that $\mu_t \to 1$ if $\theta = g$ and $\mu_t \to 0$ if $\theta = b$.

However, learning might stop before this convergence. For example, given the fixed transaction price p, risk-neutrality implies that a trade at time t requires $\mu_t g + (1-\mu_t)b < p$. Rearranging, we get the requirement that $\mu_t > \underline{\mu} := \frac{p-b}{g-b}$. If we started with $\mu_0 = \underline{\mu} + \epsilon$ and $g < 1$, then the first transaction may fail. For ϵ sufficiently small, a failure lead to $\mu_1 < \underline{\mu}$; even though the seller is good, no future transactions occur. The convergence above thus requires infinitely many observations.

Aside: Suppose each buyer now also had some independent signal about θ plus knowledge of past transactions. This is known as an *information herding* setting, and was developed in 1992. In such a situation you get that signals become irrelevant after some point. As a toy example, imagine that you wake up and look outside to decide whether or not you should bring an umbrella. Also imagine that such a signal is independent of the signal other people near you observe. Lastly, suppose that you look outside the window and can see whether or not other people (with their independent signals) carry an umbrella. It turns out that you'll quickly start ignoring the signal you get, and if, e.g., you see a handful of people carrying an umbrella, you'll bring one.

19.6 Moral Hazard: Understanding a repeated game

In this setting, we let $b = b$ so that the type cannot be unknown (equivalently, we can let $\mu_0 = 1$). We'll, however, assume that $\underline{e} < e_g^*$ and $\underline{e}g < p < e_g^* g$.

As a thought experiment, suppose that only one trade occurs. Both the seller and buyer know that one transaction will occur, so the seller does not have any incentive to put in extra effort that might improve their reputation for the non-existent future transactions, and the buyer also recognizes this. The only equilibrium involves the seller putting in $\underline{e}$ and the buyer not buying. Interesting, this is true as long as you assume a finite number of transactions, which can be shown by backwards induction.

Thus, we'll instead consider an infinite horizon setting where the seller has a discount factor δ. As motivation, suppose you get a bounded reward $R_t > 0$ at each step. Then $\sum_{i=0}^{\infty} R_i$ can't be analyzed if the series converges to infinity. However, for $\delta < 1$ you can analyze the convergent geometric series $\sum_{i=0}^{\infty} \delta^i R_i$. Intuitively, you can rationalize the discount factor as encoding a seller's uncertainty about how long they will stay in the market, and adding such a discount factor makes the analysis tractable.

In this setting with a discount factor, many equilibria occur. These include a trivial equilibria where the seller puts in $\underline{e}$ every time and no trade happens. More interestingly is the *carrot and stick* equilibria (also known as a bootstrap equilibria or a trigger strategy). In this case the seller exerts some effort $\hat{e} > \underline{e}$, and the buyers continually enter until the first match fails. The seller is effectively thinking that they'll be rewarded in the future for putting in some extra effort until failure.

Let V denote the sellers *value for reputation* which is defined as $V := p - c(\hat{e}) + \delta \hat{e} g V$. Think of the last term, e.g., as encoding the probability that another possible transaction happens and the value should it occur. Solving for V yields that

$$V = \frac{p - c(\hat{e})}{1 - \delta \hat{e} g},$$

so that to maximize the above, we find $c'(\hat{e}) = \delta g V$. This gives an equation for what's in the sellers best interest, and individual rationality of the buyer requires that $p < \hat{e} g$. Hence, we can get nontrivial results in a simple setting. Assuming such a solution $\hat{e} \in (\underline{e}, 1)$ exists, you can look at comparative statics (see [**?**]). You can show that V is increasing in p, g and δ, and that $p > c(\hat{e})$. The latter is a *reputation premium*, which intuitively reflects the fact that a seller needs to maintain her reputation, and so a buyer needs to bay a little more to enable that. You can also interpret $p < \hat{e} g$ as encoding a buyer being given some sort of discount because of the uncertainty involved in the transaction.

Bibliography

[1] Resnick, Kuwabara, Zeckhauser, and Friedman *Reputation Systems* http://dl.acm.org/citation.cfm?id=355122, 2000.

[2] Bar-Isaac and Tadelis *Seller Reputation* http://faculty.haas.berkeley.edu/stadelis/seller_rep_062608.pdf, 2008.

[3] Akerlof *The Market for "Lemons": Quality Uncertainty and the Market Mechanism* https://www.jstor.org/stable/1879431, 1970.