Lecture 4

Bayesian Games and No Trade

The receipt of private information cannot create any incentives to trade.

Today is really about understanding games of incomplete information and Bayesian Nash equilibrium; the no-trade result is sort of a corollary.

But first, a little tangent on Bayes' Law.

1 Bayes' Law and Odds Ratios

- Last week, we saw Bayes' Law: $\Pr(E|P_1) = \frac{\Pr(E \cap P_1)}{\Pr(P_1)}$
- We can think of this as

$$\Pr(E|data) \quad = \quad \frac{\Pr(E)\Pr(data|E)}{\Pr(data)} \quad = \quad \frac{\Pr(E)\Pr(data|E)}{\Pr(E)\Pr(data|E)+\Pr(\neg E)\Pr(data|\neg E)}$$

- Given a prior and data, this will always allow us to compute a posterior probability
- But if we're in a setting where we expect to get multiple "new" pieces of information, it's nice to have an easier way to update beliefs
- For example, suppose there are two possible states of the world, "GOOD" and "BAD", and we're going to get repeated noisy signals about the true state
- In each period, we get a signal that matches the true state with probability $\frac{2}{3}$
- That is, if the state is GOOD, we get a good signal with probability $\frac{2}{3}$ and a bad signal with probability $\frac{1}{3}$; and if the state is BAD, that's reversed
- Suppose our initial prior is $\Pr(G) = \frac{3}{4}$, and we get the string of signals good, bad, good, bad, bad, or *gbgbb*

- What is our posterior?
- Well, by Bayes' Law,

$$\Pr(G|gbgbb) \quad = \quad \frac{\Pr(G)\Pr(gbgbb|G)}{\Pr(gbgbb)} \quad = \quad \frac{\Pr(G)\Pr(gbgbb|G)}{\Pr(G)\Pr(gbgbb|G) + \Pr(B)\Pr(gbgbb|B)}$$

Instead of using this, though, let's look at the *odds ratio* – the ratio of p to $1-p$ – which is

$$\frac{\Pr(G|gbgbb)}{\Pr(B|gbgbb)} \quad = \quad \frac{\frac{\Pr(G)\Pr(gbgbb|G)}{\Pr(G)\Pr(gbgbb|G)+\Pr(B)\Pr(gbgbb|B)}}{\frac{\Pr(B)\Pr(gbgbb|B)}{\Pr(G)\Pr(gbgbb|G)+\Pr(B)\Pr(gbgbb|B)}} \quad = \quad \frac{\Pr(G)\Pr(gbgbb|G)}{\Pr(B)\Pr(gbgbb|B)}$$

We can even decompose this further – since the signals are independent in each period, the probability of *gbggb* given a particular state is just the probability of g in that state, times the probability of b in that state, times, etc. So we can write this as

$$\begin{aligned}\frac{\Pr(G|gbgbb)}{\Pr(B|gbgbb)} &= \frac{\Pr(G)\Pr(g|G)\Pr(b|G)\Pr(g|G)\Pr(b|G)\Pr(b|G)}{\Pr(B)\Pr(g|B)\Pr(b|B)\Pr(g|B)\Pr(b|B)\Pr(b|B)} \\ &= \frac{\Pr(G)}{\Pr(B)} \cdot \frac{\Pr(g|G)}{\Pr(g|B)} \cdot \frac{\Pr(b|G)}{\Pr(b|B)} \cdot \frac{\Pr(g|G)}{\Pr(g|B)} \cdot \frac{\Pr(b|G)}{\Pr(b|B)} \cdot \frac{\Pr(b|G)}{\Pr(b|B)}\end{aligned}$$

- So in the example I mentioned above, we would start with the prior odds ratio $\frac{3/4}{1/4} = 3$, and then multiply it by $\frac{2/3}{1/3} = 2$ every time we get a good signal, and multiply it by $\frac{1/3}{2/3} = \frac{1}{2}$ every time we get a bad signal, to get an updated odds ratio
- The point: if we work in odds ratios, rather than probabilities, then we can just start with the prior odds ratio, and each time we get a new signal, multiply the odds ratio by the relative likelihood of the new data under the two states, to get the posterior odds ratio; and we can keep doing this over and over with each new signal that we receive.

- OK, enough tangent, back to work
- Today's real task is to update our notion of Nash equilibrium to include private information

2 Review – (Static) Nash Equilibrium

- The standard setup for a static game of complete information requires three elements:
  - A set of players $N = \{1, 2, \ldots, n\}$
  - A set of available actions A_i for each player
  - A set of payoff functions
  $$U_i : A_1 \times A_2 \times \cdots \times A_n \to \Re$$
  giving the payoff to each player given each profile of strategies
- A (pure strategy) Nash equilibrium is an action profile $a = (a_1, a_2, \ldots, a_n)$ such that each player is best-responding to the remaining players' actions – that is, such that for all i,
$$a_i \in \arg\max_{a \in A_i} U_i(a, a_{-i})$$
where a_{-i} is a vector of the other $n-1$ players' actions.
- A key maintained assumption is that the entire environment is *common knowledge*
  - everyone knows the set of players, actions and payoffs,
  everyone knows everyone knows it,
  and so on.

- So what happens if we *don't* have common knowledge about the entire environment?
- For example, suppose you know your own payoff function, but you're not sure of mine.
- That's a game of *incomplete information*, and that's what the rest of the semester will be about.

3 Games of Incomplete Information

- The way we deal with incomplete information is to assume that there are different possible *types* of each player

- For example, suppose the incomplete information is that you don't know how willing I am to fight

- We assume this is due to you not knowing which type of player I am – a tough player, or a weak player

- I know which one I am, but you don't – so you just assign some probability to me being strong (having one payoff function), and some probability to me being weak (having a different payoff function)

- Formally, we assume each player's type is his own private information – I know which type I am, but not which types my opponents turned out to be

- So a static game of incomplete information consists of...

  1. A set of players $N = \{1, 2, \ldots, n\}$, same as before
  2. A set of actions for each player A_i, same as before
  3. A set of possible types T_i for each player
  4. A probability distribution p over the set of type profiles $T_1 \times T_2 \times \cdots \times T_n$
  5. A set of payoff functions, that may depend on every player's type:

$$U_i : A_1 \times \cdots \times A_n \times T_1 \times \cdots \times T_n \rightarrow \Re$$

- And *all of that* is common knowledge – if there are lots of players in the game, everyone agrees on the probability that I'm tough, and I know what they think the probability is, and so on – I just also know what my type "turned out" to be

- Several things to note about this setup:

  - It doesn't matter whether we allow the set of available actions A_i to depend on player i's type; if we wanted to "disallow" a particular action a_i for a particular type t_i of player i, we could keep a_i in A_i and just set $U_i(a_i, \cdot, t_i, \cdot) = -\infty$. So for simplicity, we assume A_i is fixed across types.
  - Information partitions are defined by the type space – player i is assumed to know the exact value of T_i, but know nothing about T_{-i} beyond what he infers from T_i. We very often, for simplicity, assume that different players' types are independent, but we don't have to.

- In many models, we assume that player i's payoff does not depend on other players' types, only his own (and the action profile). That is, we often (but not always) assume that I care about other players' types *only because it may influence their actions*, not because it directly influences my payoffs. But there are exceptions – in an adverse selection model, for example.
- Like before, we assume that the entire setup above is common knowledge – everyone agrees on the basic universe we live in, but only player i knows which type he turned out to be.
- It's customary to think of the game happening in two stages: in the first stage, "nature moves" by randomly assigning a type to each player; in the second stage, players play the game given their realized types

4 Bayesian Nash Equilibrium

The solution concept for a static game of incomplete information is Bayesian Nash equilibrium, which is really just a generalization of Nash equilibrium to accommodate types. Specifically...

- A *strategy* for player i is now a type-dependent choice of action, that is, a mapping

$$s_i : T_i \to A_i$$

specifying an action $s_i(t_i) \in A_i$ that I plan to take for each type I might turn out to be

- A Bayesian Nash equilibrium is a profile of strategies $(s_1, s_2, \dots, s_n)$ such that for every player i and type t_i,

$$s_i(t_i) \in \arg\max_{s \in A_i} E_{T_{-i}|t_i} U_i(s, s_{-i}(t_{-i}), t_i, t_{-i})$$

That is, each type of each player is maximizing his expected payoff, given his correct beliefs about the probabilities of different opponent types (given his own) and given his correct beliefs about his opponents' strategies

5 An Example of BNE: First Price Auctions

- Let's do a simple example: a private-values, sealed-bid first-price auction

  There is a single object for sale

  Each player's type is the value he would get from winning it; everyone simultaneously submits a bid in writing, and the player with the highest bid pays his bid and receives the object

  Let's suppose types are independent and uniform over the interval $[0, 1]$

  So formally:

  - The set of players is $1, 2, \ldots, n$
  - The set of available actions is $A_i = \Re^+$ – any positive bid is allowed
  - The set of possible types is $T_i = [0, 1]$, and the probability distribution over type profiles is uniform over $[0, 1]^n$ (meaning it has density 1 everywhere)
  - The payoff from winning the auction is your type minus your bid, and the payoff from losing is 0, so if we ignore ties, we can write

$$U_i(a, t) = \begin{cases} t_i - a_i & if \quad a_i > \max_{j \neq i} a_j \\ 0 & if \quad a_i < \max_{j \neq i} a_j \end{cases}$$

  - To complete the model, let's assume they're broken randomly; so if $a_i = \max_{j \neq i} a_j$, the payoff to player i is

$$(t_i - a_i) \frac{1}{1 + \#\{j \neq i : a_j = a_i\}}$$

- Note that we've assumed *private values* – other players' types don't enter directly into my payoff function, they affect me only through their effect on my opponents' bids.

- What does the Bayesian Nash equilibrium of this game look like? It turns out, it's an equilibrium for everyone to bid $\frac{n-1}{n}$ times their type.

- Why?

- To show this, we need to show that if all my opponents are playing this strategy, then this strategy maximizes my expected payoff.

- So suppose my type is t_i, my opponents are all bidding $\frac{n-1}{n}$ times their values, and I bid b.

- First of all, note that if I bid more than $\frac{n-1}{n}$, I'll win with probability 1 – I outbid any opponent with type $t_j < 1$, and tie opponents with type $t_j = 1$, which occurs with prob 0

  So bidding more than $\frac{n-1}{n}$ is a bad idea – it drives up the price I pay, without making me more likely to win. So the only strategies that might be best-responses are in the range $[0, \frac{n-1}{n}]$.

- Now, for b within that range, my expected payoff is

$$\begin{aligned}
E_{T_{-i}}U_i &= (t_i - b)\cdot \Pr(win|b) + 0\cdot \Pr(lose|b) \\
&= (t_i - b)\Pr(\max_{j\neq i} s_j(t_j) < b) \\
&= (t_i - b)\Pr(\max_{j\neq i} \tfrac{n-1}{n} t_j < b) \\
&= (t_i - b)\Pr(\max_{j\neq i} t_j < \tfrac{n}{n-1} b) \\
&= (t_i - b)(\tfrac{n}{n-1} b)^{n-1} \\
&= (\tfrac{n}{n-1})^{n-1}(t_i - b) b^{n-1}
\end{aligned}$$

- So now let's maximize this thing: $(t_i - b)b^{n-1} = t_i b^{n-1} - b^n$ has derivative

$$(n-1)t_i b^{n-2} - nb^{n-1} \quad = \quad nb^{n-2}\left[\frac{n-1}{n}t_i - b\right]$$

  This is increasing on $b < \frac{n-1}{n}t_i$ and decreasing on $b > \frac{n-1}{n}t_i$ – so it's maximized at $b = \frac{n-1}{n}t_i$.

- So if everyone else bids $\frac{n-1}{n}$ times their type, my best-response is to bid $\frac{n-1}{n}$ times my type; so everyone bidding $\frac{n-1}{n}$ times their type is an equilibrium.

6 An Application of BNE: A No-Trade Theorem

- Last week, we saw Aumann's result that if both our posterior beliefs are common knowledge, we can't disagree
- A nice analogy is to a trading problem: basically, if the entire environment is common knowledge, then differential private information on its own can't lead to trade
- We'll just prove a "simple" version, then describe how the result extends
- For the simple version...
  - there are just two traders, me and my bookie.
  - there's one event we care about – whether or not my team wins this Sunday
  - my bookie and I are both strictly risk-averse – but we still might want to bet because we have different information about my team's chances
  - my bookie and I share a common prior, but have different information partitions – he might have some information I don't have, and I might have some information he doesn't have
  - I've got wealth w_1, he's got wealth w_2, my utility for wealth is u_1 and his is u_2, and both strictly concave
- So here's the result:

Theorem. *In equilibrium, we never trade.*

- The basic logic is that, if I decide to place a bet, I need to condition not only on my own information, *but also on the fact that my bookie wants to bet with me*; and on the fact that he wants to trade with me *even knowing I want to trade with him*; and so on
- And similarly, he needs to condition on the fact that I want to bet with him
- In essence, this means it has to be common knowledge that I believe the bet is skewed in my favor, and also common knowledge that he believes the bet is skewed in his favor – which an analogy to Aumann suggests is impossible
- Which means we can't both be willing to bet if it's common knowledge we're both rational

Let's prove it formally

- As always, we have a set of states Ω,

  my information partition $\mathcal{P}_1 = \{t_1^1, t_1^2, \ldots, t_1^m\}$,

  and my bookie's information partition $\mathcal{P}_2 = \{t_2^1, \ldots, t_2^{m'}\}$

- Let's suppose there is a finite number of possible bets we might consider making

  (This is WLOG, if we limit ourselves to bets with non-fractional number of cents, in amounts less than the total wealth of the entire earth)

- And suppose that if we do make a bet, that occurs within the Bayesian Nash equilibrium of some sort of negotiation process

  We don't have to specify what that process is; we just need the environment we're in (including the prior) to be common knowledge, and that in order for a bet to happen, we both have to agree to it.

- What I want to show is that with probability 1, we don't bet

- So suppose that weren't true – that with positive probability, we made some nonzero bet

  Since there are a finite number of bets, that means there is some bet (x, y) we make with positive probability

  (Let x be the amount I receive if my team wins, and y the amount I receive if my team loses – so if I'm betting on my team, $x > 0 > y$)

- Let $q(t_1, t_2)$ be the probability that that bet occurs,

  given my information $t_1 \in \mathcal{P}_1$ and my bookie's $t_2 \in \mathcal{P}_2$

- For this to happen in equilibrium, two things need to be true:

  - For every information set t_1 at which I agree to this bet with positive probability, I must believe that in expectation over my bookie's possible information sets, and conditional on him also being willing to accept this bet, the bet does not decrease my expected payoff
  - For every information set t_2 at which my bookie agrees to this bet with positive probability, he must believe that in expectation over my possible information sets, and conditional on me also being willing to accept this bet, the bet does not decrease his expected payoff

  We'll show that these can't both hold unless the trade is $(0, 0)$.

- Suppose I'm at information set t_1, and I'm willing to agree to the bet (x, y)

- Conditional on you agreeing to the bet, I put some posterior probability $p(t_2^j|t_1, bet)$ on you being at each information set t_2^j
- Which means that I evaluate my expected payoff, *if we both agree to the bet*, as...

$$\sum_{t_2^j \in \mathcal{P}_2} p(t_2^j|t_1, bet) \left(p(W|t_1, t_2^j) u_1(w_1 + x) + (1 - p(W|t_1, t_2^j)) u_1(w_1 + y) \right)$$

$$= \left(\sum_{t_2^j \in \mathcal{P}_2} p(t_2^j|t_1, bet) p(W|t_1, t_2^j) \right) u_1(w_1 + x) + \left(\sum_{t_2^j \in \mathcal{P}_2} p(t_2^j|t_1, bet)(1 - p(W|t_1, t_2^j)) \right) u_1(w_1 + y)$$

Now, in order for me to be willing to make the bet, this has to be at least as good as $u_1(w_1)$, my utility from not betting at all:

$$\left(\sum_{t_2^j \in \mathcal{P}_2} p(t_2^j|t_1, bet) p(W|t_1, t_2^j) \right) u_1(w_1 + x) + \left(\sum_{t_2^j \in \mathcal{P}_2} p(t_2^j|t_1, bet)(1 - p(W|t_1, t_2^j)) \right) u_1(w_1 + y) \geq u_1(w_1)$$

Now let's multiply both sides by the probability that I have type t_1 and we both agree to (x, y), so this becomes

$$\left(\sum_{t_2^j \in \mathcal{P}_2} p(t_1, bet) p(t_2^j|t_1, bet) p(W|t_1, t_2^j) \right) u_1(w_1 + x)$$

$$+ \left(\sum_{t_2^j \in \mathcal{P}_2} p(t_1, bet) p(t_2^j|t_1, bet)(1 - p(W|t_1, t_2^j)) \right) u_1(w_1 + y) \geq p(t_1, bet) u_1(w_1)$$

Now let's sum over all the different values of t_1 that I might have, giving

$$\left(\sum_{t_1^i \in \mathcal{P}_1} \sum_{t_2^j \in \mathcal{P}_2} p(t_1^i, bet) p(t_2^j|t_1^i, bet) p(W|t_1^i, t_2^j) \right) u_1(w_1 + x)$$

$$+ \left(\sum_{t_1^i \in \mathcal{P}_1} \sum_{t_2^j \in \mathcal{P}_2} p(t_1^i, bet) p(t_2^j|t_1, bet)(1 - p(W|t_1^i, t_2^j)) \right) u_1(w_1 + y) \geq \sum_{t_1^i} p(t_1^i, bet) u_1(w_1)$$

Now, iterated expectations says that $\Pr(A \ \ and \ \ B) = \Pr(A) \Pr(B|A)$. Applying this a couple of times, we simplify the last expression to

$$\Pr(bet) \Pr(W|bet) u_1(w_1 + x) + \Pr(bet) \left(1 - \Pr(W|bet)\right) u_1(w_1 + y) \geq \Pr(bet) u_1(w_1)$$

or, since $\Pr(bet) > 0$,

$$\Pr(W|bet) u_1(w_1 + x) + \left(1 - \Pr(W|bet)\right) u_1(w_1 + y) \geq u_1(w_1)$$

Let $p^* = \Pr(W|bet)$. Since I'm risk-averse, I'd strictly prefer a sure thing with the same expected value as the bet: that is,

$$u_1(p^*(w_1+x)+(1-p^*)(w_1+y)) > p^*u_1(w_1+x)+(1-p^*)u_1(w_1+y) \geq u_1(w_1)$$

meaning $p^*(w_1+x)+(1-p^*)(w_1+y) > w_1$, or $p^*x+(1-p^*)y > 0$. So if I only bet when I want to bet, it has to be the case that, averaged over all the different cases where we bet, the bet has strictly positive expected value for me.

But if we did the same analysis from my bookie's point of view, we would have to conclude that, averaged over all the different cases where we bet, the bet has strictly positive expected value for him.

And those can't both be true.

7 Milgrom and Stokey

Milgrom and Stokey, "Information, Trade, and Common Knowledge," prove this result in a much more general setting.

- They allow $N > 2$ traders
- They allow any finite number of states of the world, which they break up into two components – a *payoff-relevant* state, and a part of the state that just conveys information about the payoff-relevant part
- They allow l different goods, which are each consumed in continuous quantities
- They allow *endowments* of the different goods to vary across payoff-states – so each trader has an endowment $e_i \in (\Re^+)^l$ in each payoff-state, and consumes a vector $x_i \in (\Re^+)^l$ in each payoff state
- They allow each trader's *utility function* to vary across payoff-states – all they require is each state's utility function is strictly concave
- They don't even require that we all have a common prior over the payoff-states – they only require that we have the same beliefs about how the information-states depend on the payoff-states
- And they allow us to make any state-contingent trades we want.

So for example, there could be one (payoff-relevant) state of the world where you lose your job. So in that state, you have less money. So you'd probably be willing to make a trade where you get more money in that state, and give up money in states where you still have your job. You and I can disagree on how likely you are to get fired – I might think it's 10% likely, you think it's 5% likely – but we both agree that *if* you're going to get fired, there's a 50% chance you'll be told you're on probation at your performance review.

Note that this model gives three different reasons we may want to trade:

- To reallocate goods within a state. In some state of the world, I'm endowed with a lot of one good, and want to consume other goods, so I want to trade goods within that state.
- To reallocate goods across states. I may want to trade stuff into one state (where I'm endowed with less, or have a steeper utility functino), out of another state (where I'm endowed with more). Note that my utility function within each state has to be strictly concave, but there's no restriction on how my utility function varies across states.

- To arbitrage different beliefs. If we have different beliefs about the likelihood of an event, I use my belief to evaluate my expected payoff, and you use your belief to evaluate yours; so we can both be better off by betting against each other.

And of course, in addition to that, we may want to trade because we received different information.

Milgrom and Stokey basically show that the last piece alone is not sufficient to generate trade. In a sense, they let us trade twice. First, they let us make whatever state-contingent trades we want to, and assume that we reach an allocation which is Pareto-efficient. *Then*, they assume each of us gets new information; and they show that there can't be any *new* trades based on that new information.

But the proof isn't really interesting – basically, they show that if there was ever a trade we'd both want to make after receiving additional information, we should have been willing to make that trade ex ante, just making it contingent on different realizations of that later information. It's not quite that simple, but that's the gist of it; so if we had already traded to a Pareto-efficient allocation, new information can't make us trade more.

References

- Paul Milgrom and Nancy Stokey (1982), "Information, Trade and Common Knowledge," *Journal of Economic Theory* 26(1)