Lecture 5

Optimal Auctions I

Lots of auction formats are equally good...
and we can find the best possible one.

1 Before we get started...

1.1 Second-price auctions

- Last week, as an example of Bayesian Nash equilibrium, we solved for the equilibrium of a sealed-bid, first-price auction, under the assumption that n risk-neutral bidders have private values which are independently uniform on $[0, 1]$
- Before we start today, I want to talk about one other common auction format
- Description of open oral ascending, or English, auctions
- A modeling convenience: for the private values case, think of a sealed-bid *second*-price auction
  - Bidders simultaneously submit written bids, as in a first-price
  - High bidder wins, but his payment is the *second*-highest bid, not his own
- With private values, turns out it's a dominant strategy to bid your valuation!
- Think of y a random variable being the highest opponent bid; you have a choice between bidding above y (for payoff $t_i - y$) and bidding below y (and getting 0)
- So when $t_i > y$, you want to overbid, and when $t_i < y$, you want to underbid – which you accomplish by bidding t_i
- So it's a Bayesian Nash equilibrium for everyone to bid their valuation

1.2 Revenue Equivalence

- Last week, we determined that in a first-price auction, it's an equilibrium for each bidder to bid $\frac{n-1}{n}$ times his or her valuation
- So in a first-price auction, the seller's expected revenue is the expected value of the highest bid, which is $\frac{n-1}{n}$ times the expected value of the highest valuation
- Turns out, with n independent uniform random variables, the expected value of the highest is $\frac{n}{n+1}$
- So expected revenue in the first-price auction is $\frac{n-1}{n} \times \frac{n}{n+1} = \frac{n-1}{n+1}$

  (Expected value of the highest of n independent $U[0,1]$ random variables is calculated by taking the CDF, x^n; calculating the PDF, nx^{n-1}, and taking the integral $\int_0^1 x \cdot nx^{n-1}dx = \int_0^1 nx^n dx = \frac{n}{n+1}$.)
- What about the second-price auction?
- In the second-price auction, everyone bids his own type, and the payment to the seller is the second-highest bid
- So revenue is the second-highest valuation
- Which has expected value $\frac{n-1}{n+1}$

  (The CDF of the second-highest of n is $nx^{n-1}(1-x) + x^n = nx^{n-1} - (n-1)x^n$, so the PDF is $n(n-1)(x^{n-2} - x^{n-1})$, and our expected value is $\int_0^1 n(n-1)(x^{n-1} - x^n)dx = (n-1)x^n - \frac{n(n-1)}{n+1}x^{n+1}\Big|_0^1 = n - 1 - \frac{n^2-n}{n+1} = \frac{n^2-1-n^2+n}{n+1} = \frac{n-1}{n+1}$.)
- So the two auction formats give the same expected payoff to the seller
- This turns out to be a much more general result, which we'll prove soon

2 Another Preliminary: The Envelope Theorem

- Put aside strategic concerns, and think of a one-agent decision problem

$$\max_{a \in A} h(a, \theta)$$

  where a is the agent's chosen action and θ an exogenous parameter

- In an auction, θ could be your valuation, and a your choice of bid

- In another setting, θ could be the outdoor temperature, and A is the set of jackets and coats you own – given today's weather, you pick an outfit to maximize your comfort level

- (DRAW IT for discrete A)

- A could be either discrete or continuous, but let θ be continuous

- Let $a^*(\theta)$ be the set of optimal choices, and $V(\theta)$ the value function; and let h_a and h_θ denote partial derivatives of h

Theorem (The Envelope Theorem)**.** *Suppose $\forall \theta$, $a^*(\theta)$ is nonempty, and $\forall (a, \theta)$, h_θ exists. Let $a(\theta)$ be any selection from $a^*(\theta)$.*

1. *If V is differentiable at θ, then*

$$V'(\theta) = h_\theta(a(\theta), \theta)$$

2. *If V is absolutely continuous, then for any $\theta' > \theta$,*

$$V(\theta') - V(\theta) = \int_\theta^{\theta'} h_\theta(a(t), t) dt$$

- This says: the derivative of the *value function* (or maximum) is the derivative of the *objective function*, evaluated *at the maximizer*

- Absolute continuity says that $\forall \epsilon > 0$, $\exists \delta > 0$ such that for any finite, disjoint set of intervals $\{[x_k, y_k]\}_{k=1,2,\ldots,M}$ with $\sum_k |y_k - x_k| < \delta$, $\sum_k |V(y_k) - V(x_k)| < \epsilon$.

  Absolutely continuity is equivalent to V being differentiable almost everywhere and being the integral of its derivative, so the second part follows directly from the first.

  (For the types of problems we'll be dealing with, value functions will generally be absolutely continuous.)

- Let's prove the first part

2.1 Proof of the Envelope Theorem

- If V is differentiable at θ, then

$$V'(\theta) = \lim_{\epsilon \to 0} \frac{V(\theta+\epsilon) - V(\theta)}{\epsilon} = \lim_{\epsilon \to 0} \frac{V(\theta) - V(\theta-\epsilon)}{\epsilon}$$

- Now, pick $a^* \in a^*(\theta)$, so $V(\theta) = h(a^*, \theta)$
- And $V(\theta+\epsilon) = \max_a h(a, \theta+\epsilon) \geq h(a^*, \theta+\epsilon)$
- So

$$V'(\theta) = \lim_{\epsilon \to 0} \frac{V(\theta+\epsilon) - V(\theta)}{\epsilon} \geq \lim_{\epsilon \to 0} \frac{h(a^*, \theta+\epsilon) - h(a^*, \theta)}{\epsilon} = h_\theta(a^*, \theta)$$

- And by the same token, $V(\theta-\epsilon) = \max_a h(a, \theta-\epsilon) \geq h(a^*, \theta-\epsilon)$
- So

$$V'(\theta) = \lim_{\epsilon \to 0} \frac{V(\theta) - V(\theta-\epsilon)}{\epsilon} \leq \lim_{\epsilon \to 0} \frac{h(a^*, \theta) - h(a^*, \theta-\epsilon)}{\epsilon} = h_\theta(a^*, \theta)$$

- So

$$h_\theta(a^*, \theta) \leq V'(\theta) \leq h_\theta(a^*, \theta)$$

  and we're done

- The second half of the theorem is just the fact that if V is absolutely continuous, it's the integral of its derivative (wherever the derivative exists)
- Typically, for the types of problems we'll be looking at, V will be absolutely continuous
- (For example, if h_θ exists everywhere and is bounded above, that's sufficient.)
- So even if V isn't differentiable everywhere, the integral form of the theorem will be valid.

2.2 Example – using the envelope theorem to solve for equilibrium bids in FPA

- For our setup last week (PV *i.i.d.* $\sim U[0,1]$), we can use the envelope theorem to recover equilibrium bid functions

- Suppose a symmetric equilibrium exists, and uses a strictly-increasing bid function, so the bidder with the highest type always wins in equilibrium

- Equilibrium bids solve $\max_b (t_i - b) \Pr(b > \max_{j \neq i} b_i)$, call maximand $h(t_i, b)$ and U its max

- The partial of h with respect to t_i is $\Pr(b > \max_{j \neq i} b_i)$

- If we assume the equilibrium is symmetric, then at the maximizer, this equals the probability my type is the highest, which is t_i^{n-1}

- Since the lowest-type bidder always loses, $U(0) = 0$, so

$$U(t_i) \quad = \quad \int_0^{t_i} s^{n-1} ds \quad = \quad \frac{1}{n} t_i^n$$

- But if we calculate expected payoff based on bid and probability of winning,

$$U(t_i) \quad = \quad (t_i - \beta(t_i))\, t_i^{n-1}$$

- And if we equate these,

$$\frac{1}{n} t_i^n \quad = \quad (t_i - \beta(t_i)) t_i^{n-1}$$

  implies $\beta(t_i) = \frac{n-1}{n} t_i$

- So if a symmetric, strictly-increasing equilibrium exists, bids must be $\frac{n-1}{n} t_i$ – exactly like we found last week!

2.3 Another example – the all-pay auction

- In class, we also did the example of an *all-pay* auction – an auction where the highest bidder wins, but *every* bidder pays their bid, even the losers

- Again, if we suppose a symmetric, strictly-monotonic equilibrium exists, we can use it to calculate the expected payoff to each type of bidder, which is the same $U(t_i) = \frac{1}{n} t_i^n$

- We then write $\frac{1}{n} t_i^n = t_i^{n-1} t_i - b$ and calculate equilibrium bids as $b = \frac{n-1}{n} t_i^n$

- Finally, we can calculate the seller's expected revenue as $n E_{t_i} \frac{n-1}{n} t_i^n$ and see that it's the same $\frac{n-1}{n+1}$ as before – more anecdotal evidence of revenue equivalence!

3 Now on to Myerson's Optimal Auctions

First, we define the environment, and our goal.

- Our environment is as follows.
  - Players $N = \{1, 2, \ldots, n\}$
  - Independent types $T_i \perp T_j$
  - Player i's type T_i has distribution F_i with support $[a_i, b_i]$
  - F_i is strictly increasing on $[a_i, b_i]$, with density f_i
  - One object to be allocated, value to player i is his type t_i, and players value money linearly, so i's payoff is $t_i - x$ if he receives object and pays x, and $-x$ if he doesn't receive object but still pays x
  - (Players are risk-neutral and maximize expected payoff)
  - Seller values keeping the good at t_0
- Note that this is the environment we looked at above for first- and second-price auctions, except that type distributions F_i can be arbitrary (not just uniform) and don't have to be the same
- We're still assuming:
  - Independent private values
  - Risk-neutrality (linear valuation for money)
  - Whole environment is common knowledge
- Our challenge: design an auction (or other) protocol that maximizes the seller's expected revenue, given that the buyers play a Bayesian Nash equilibrium, over all conceivable sales mechanisms.
- That is, we basically want to solve

$$\max_{all\ \ conceivable\ \ auction\ \ formats} E_{t_1,\ldots,t_n} \{Revenue + t_0 \Pr(seller\ \ keeps\ \ object)\}$$

  subject to the constraints that the buyers have to willingly participate and play an equilibrium
- The surprising thing: this is actually doable.

4 Mechanism Design

- The framework we'll use to answer this question is referred to as **mechanism design**
- This basically puts us in the role of a seller, or government, or whatever, who asks, suppose I have complete freedom to design a game – an auction format, or a voting rule, or whatever – and people will play it; how much can I accomplish?
- The only constraints we put on the mechanism designer are the following:
  1. You can't force people to play – they have to be willing to play. Which generally means their expected payoff from the game can't be less than 0.
  2. You need to assume people will play a Bayesian Nash equilibrium within whatever game you define – that is, you can't trick people into doing things that aren't in their own interest.
- We also assume the mechanism designer has **full commitment power** – once he defines the rules of the game, the players have complete confidence that he'll honor those rules. (This is important – you'd bid differently in an auction if you thought that, even if you won, the seller might demand a higher price or mess with you some other way.)
- Broadly speaking, mechanism design takes the *environment* as given – the *players*, their *type distributions*, and their *preferences over the different possible outcomes* – and designs a game for the players to play in order to select one of the outcomes. Outcomes can be different legislative proposals, different allocations of one or more objects, etc.
- We'll be focusing on the auction problem – designing a mechanism to sell a single object, and trying to maximize expected revenue
- So the set of possible outcomes X consists of who (if anyone) gets the object, and how much each person pays
- A *mechanism* is a strategy set $S_1, S_2, \ldots, S_N$ for each player, and a mapping from strategy profiles to outcomes: $\omega : S_1 \times \cdots \times S_N \rightarrow X$
- Actually, given a set of strategies played, the outcome selected by the mechanism can be deterministic or stochastic, so really $\omega : S_1 \times \cdots \times S_N \rightarrow \Delta(X)$
- For example, the first-price auction we considered last week is one mechanism: each player's strategy set is the positive reals, the mechanism allocates the good to whoever plays the highest strategy and demands a payment from that player equal to his strategy, and there are no payments to or from anyone else

- We'll assume the players play a Bayesian Nash equilibrium within the mechanism, but in general, if there are multiple equilibria, the mechanism designer is assumed to be able to pick which equilibrium gets played

- A *performance* is a mapping of types to outcomes; we say that a given performance is *implemented* by a mechanism if that mechanism has an equilibrium where the players' equilibrium strategies lead to that mapping of types to outcomes

  - For example, in a symmetric IPV world, the first-price auction is a mechanism that implements the efficient allocation (the highest type being allocated the object), along with a particular transfer from that player
  - In general, in mechanism design, we don't worry about there being multiple equilibria, just that the one that we want is indeed one of the equilibria
  - Implicitly, we're kind of assuming that, in addition to setting up the game, the mechanism designer gets to select the equilibrium played if there is more than one
  - By using this approach, we're also kind of assuming that "all that matters" to the players is what outcome is reached, not the exact process by which it is chosen

5 Direct Revelation Mechanisms and the Revelation Principle

- Much of the time, we are able to restrict our attention to a particular class of mechanisms called *direct revelation mechanisms*

- Informally, a direct revelation mechanism consists of the mechanism designer specifying a mapping from types directly to outcomes, and asking each player (in private) to tell him their type

- Formally, this is just a mechanism where the strategy set for each player matches his type space, and given the mapping of types to outcomes, we expect there to be an equilibrium where every player reveals his type truthfully

- We'll consider only mechanisms that promise possible outcomes at every strategy profile (or reported type profile) – that is, if there is no combination of actions/types at which it promises the same object to multiple players, or anything stupid like that.

Lemma 1. *(The Revelation Principle.) Given any equilibrium of any auction mechanism, there exists an equivalent direct revelation mechanism in which truthful revelation is an equilibrium and in which the seller and all the bidders get the same expected utilities.*

- Proof is basically this: take, for example, the first-price auction

- Instead of running the first-price auction, the mechanism designer approaches each of the players and says, "tell me your type, and then I'll calculate the equilibrium bids of you and the other bidders in the BNE of the first-price auction, and if yours is highest, I'll give you the object and charge you that much"

- So any deviation (to reporting another type) in the direct-revelation mechanism, would be a deviation (to that type's equilibrium bid) in the first-price auction

- By the same logic, a direct-revelation mechanism can implement any equilibrium of any feasible auction.

- For this reason, we will focus only on direct-revelation mechanisms (since they're easier to analyze).

- Again, we do not require truthful revelation to be the only equilibrium, just that it be an equilibrium.

6 Myerson's Optimal Auctions

- Jump back to our IPV setting
  - N bidders, each with independent type $t_i \sim F_i$ with support $[a_i, b_i]$
  - We don't need symmetry, but we do need independence.
  - Seller values the object at t_0
- An *outcome* is a choice of who (if anyone) gets the object, and transfers to/from each player
- So we can summarize any direct-revelation mechanism in this world as a set of mappings

$$p_i : T_1 \times T_2 \times \cdots \times T_N \to [0, 1]$$

and

$$x_i : T_1 \times T_2 \times \cdots \times T_N \to \Re$$

where at a profile of reported types $t = (t_1, t_2, \ldots, t_n)$, player i gets the object with probability $p_i(t)$ and pays in expectation $x_i(t)$

- (When the allocation is nondeterministic, it doesn't matter whether each player pays only when he gets the object or not; since the buyers are risk-neutral, all that matters is their expected payment at each profile.)
- Define $U_i(p, x, t_i)$ to be player i's equilibrium expected payoff under mechanism (p, x), assuming he has type t_i and everyone reports their type truthfully
  - This is $E_{t_{-i}} \{t_i p_i(t_i, t_{-i}) - x_i(t_i, t_{-i})\}$ – for each opponent profile t_{-i}, player i wins with probability $p(t)$ a prize worth t_i and pays $x_i(t)$

Definition 1. *A direct revelation mechanism is FEASIBLE if it satisfies...*

1. *Plausibility: for any profile of reported types t, $\sum_{j \in N} p_j(t) \leq 1$, and $p_i(t) \geq 0$ for all $i \in N$.*
2. *Individual rationality: $U_i(p, x, t_i) \geq 0$ for all $i \in N$, all $t_i \in [a_i, b_i]$.*
3. *Incentive compatibility:*

$$U_i(p, x, t_i) \quad \geq \quad E_{t_{-i}} \left\{ t_i p_i(t_{-i}, t'_i) - x_i(t_{-i}, t'_i) \right\}$$

for all $t'_i \in [a_i, b_i]$

- The first condition is just that the mechanism never promises the impossible, like giving the object to two people at once
- The second condition is that nobody would choose to opt out of the game
- The third is that given the game, everyone truthfully revealing their type is a BNE
  - That is, if everyone else reports their type truthfully, I can't gain by lying about my type – it's a best-response for me to report truthfully too
- Note that, we assume the bidders already know their types before deciding whether or not to participate
  - That is, individual rationality has to hold for each type of bidder, not just in expectation over the bidders' types
  - (If bidders had to commit to play before learning their types, mechanism design would actually be much easier – basically, pick any mechanism that allocates the object efficiently, and then calculate each bidder's expected payoff within that mechanism, and charge them that much as an entry fee. You get efficient outcomes and full surplus extraction by the seller. But that isn't realistic. We assume bidders can opt out after learning their types, so we can't charge an entry fee that low-type bidders would refuse to pay.)
- Define $Q_i(p, t_i) = E_{t_{-i}} p_i(t_i, t_{-i})$, that is, given a mechanism with allocation rule p, $Q_i(t_i)$ is the probability of bidder i getting the object given type t_i
- Lemma 2 of Myerson:

Lemma 2. *A mechanism (p, x) is feasible if and only if:*

4. *Q_i is weakly increasing in t_i*
5. *$U_i(p, x, t_i) = U_i(p, x, a_i) + \int_{a_i}^{t_i} Q_i(p, s) ds$ for all i, all t_i*
6. *$U_i(p, x, a_i) \geq 0$ for all i, and*
7. *$\sum_j p_j(t) \leq 1$, $p_i \geq 0$*

To prove this, we need to show $\{1,2,3\} \leftrightarrow \{4,5,6,7\}$.

First, $\{1,2,3\} \rightarrow \{4,5,6,7\}$

- First, 1 and 7 are the same condition, so $1 \rightarrow 7$
- Second, 2 says $U_i(p,x,t_i) \geq 0$ for all t_i, so it must hold for $t_i = a_i$, which is 6, so $2 \rightarrow 6$
- So what's left to show is that 3 implies 4 and 5.
- Let's show (3) $\rightarrow$ (4), that is, IC implies Q_i increasing
  - Pick $t_i' > t_i$
  - Let $x_i = E_{t_{-i}} x_i(t_i, t_{-i})$ and $x_i' = E_{t_{-i}} x_i(t_i', t_{-i})$
  - Incentive compatibility (3) requires

$$\begin{aligned} t_i Q_i(t_i) - x_i &\geq t_i Q_i(t_i') - x_i' \\ t_i' Q_i(t_i') - x_i' &\geq t_i' Q_i(t_i) - x_i \end{aligned}$$

  - Rearranging gives

$$t_i'(Q_i(t_i') - Q_i(t_i)) \quad \geq \quad x_i' - x_i \quad \geq \quad t_i(Q_i(t_i') - Q_i(t_i))$$

    and therefore

$$(t_i' - t_i)(Q_i(t_i') - Q_i(t_i)) \quad \geq \quad 0$$

    so since $t_i' > t_i$, $Q_i(t_i') \geq Q_i(t_i)$
  - Meaning Q_i is weakly increasing, which is condition 4
- Finally, we need to show that $3 \rightarrow 5$
  - This is basically just the envelope theorem
  - By construction, everyone is truthfully announcing their types in equilibrium, so fix everyone elses' strategy (truthful reporting) and consider bidder i's problem
  - He has type t_i, and can pick a type t_i' to report, which will give him payoff $t_i Q_i(t_i') - x_i(t_i')$
  - So his expected payoff is $U_i(t_i) = \max_{t_i'} t_i Q_i(t_i') - x_i(t_i')$
  - So the envelope theorem says that $U_i'(t_i) = \frac{\partial}{\partial t_i} \left(t_i Q_i(t_i') - x_i(t_i') \right)|_{t_i' = t_i} = Q_i(t_i)$
  - So integrating, $U_i(t_i) - U_i(a_i) = \int_{a_i}^{t_i} U_i'(s) ds = \int_{a_i}^{t_i} Q_i(s) ds$
  - (If you don't like relying on the envelope theorem, we can also prove it directly – see next page)

Parenthetical: If We Wanted To Prove This Without the Envelope Theorem

IC implies

$$U_i(t_i') - U_i(t_i) \quad = \quad t_i'Q_i(t_i') - x_i' - (t_iQ_i(t_i) - x_i) \quad \geq \quad t_i'Q_i(t_i) - x_i - (t_iQ_i(t_i) - x_i) \quad = \quad (t_i' - t_i)Q_i(t_i)$$

and likewise

$$U_i(t_i') - U_i(t_i) \quad = \quad t_i'Q_i(t_i') - x_i' - (t_iQ_i(t_i) - x_i) \quad \leq \quad t_i'Q_i(t_i') - x_i' - (t_iQ_i(t_i') - x_i') \quad = \quad (t_i' - t_i)Q_i(t_i')$$

So for any $t_i' > t_i$,

$$(t_i' - t_i)Q_i(t_i') \quad \geq \quad U_i(t_i') - U_i(t_i) \quad \geq \quad (t_i' - t_i)Q_i(t_i)$$

- So now, we calculate $U_i(t_i) - U_i(a_1)$
- We can break the interval $[a_i, t_i]$ up into M smaller intervals, each of size $\epsilon = \frac{t_i - a_i}{M}$, and write

$$U_i(t_i) - U_i(a_i) \quad = \quad (U_i(a_i + \epsilon) - U_i(a_i)) \quad + \quad (U_i(a_i + 2\epsilon) - U_i(a_i + \epsilon)) \quad + \quad \cdots \quad + \quad (U_i(t_i) - U_i(t_i - \epsilon))$$

Using the result we just showed,

$$\begin{array}{ccccccccc} U_i(t_i) - U_i(a_i) & = & (U_i(a_i + \epsilon) - U_i(a_i)) & + & (U_i(a_i + 2\epsilon) - U_i(a_i + \epsilon)) & + & \cdots & + & (U_i(t_i) - U_i(t_i - \epsilon)) \\ & & \geq & & \geq & & & & \geq \\ & & \epsilon Q_i(a_i) & & \epsilon Q_i(a_i + \epsilon) & & & & \epsilon Q_i(t_i - \epsilon) \end{array}$$

and also

$$\begin{array}{ccccccccc} U_i(t_i) - U_i(a_i) & = & (U_i(a_i + \epsilon) - U_i(a_i)) & + & (U_i(a_i + 2\epsilon) - U_i(a_i + \epsilon)) & + & \cdots & + & (U_i(t_i) - U_i(t_i - \epsilon)) \\ & & \leq & & \leq & & & & \leq \\ & & \epsilon Q_i(a_i + \epsilon) & & \epsilon Q_i(a_i + 2\epsilon) & & & & \epsilon Q_i(t_i) \end{array}$$

- If we make M big, so each interval (and ϵ) get small, we have an upper and lower bound on $U_i(t_i) - U_i(a_i)$ that go to the same limit – which is the Riemann integral of Q_i over $[a_i, t_i]$
- So
$$U_i(t_i) - U_i(a_i) \quad = \quad \int_{a_i}^{t_i} Q_i(s)ds$$
which is exactly property 5, so we're done with the first direction.
- (Note that Q_i is Riemann-integrable because it's bounded (always within $[0, 1]$) and, since it's monotonically increasing, it must be continuous almost everywhere (if it had uncountably many discontinuities, it couldn't be bounded).)
- (Or if we want to see it "more directly," the upper and lower bound have the same limit because the two sums differ by exactly $\epsilon Q_i(t_i) - \epsilon Q_i(a_i)$, but $Q_i(t_i)$ and $Q_i(a_i)$ are both bounded within $[0, 1]$, so the two sums differ by at most ϵ, which goes to 0.)

Back to Work – still need to show $\{4,5,6,7\} \rightarrow \{1,2,3\}$

- Again, 1 and 7 are the same
- 5 implies U_i is increasing; along with $U_i(a_i) \geq 0$, this implies IR holds everywhere (2)
- All that's left is to show that 4 and 5 together imply 3 – that is, that monotonicity and the envelope condition together imply incentive compatibility
- Given a true type t_i, we can write the gain from deviating and reporting a type t_i' as

$$\begin{aligned} gain &= t_i Q_i(t_i') - x_i' - U_i(t_i) \\ &= t_i' Q_i(t_i') - (t_i' - t_i) Q_i(t_i') - x_i' - U_i(t_i) \\ &= U_i(t_i') - U_i(t_i) - (t_i' - t_i) Q_i(t_i') \end{aligned}$$

- For the case where $t_i' > t_i$, using 5, this is

$$gain = \int_{t_i}^{t_i'} Q_i(s)ds - (t_i' - t_i)Q_i(t_i') = \int_{t_i}^{t_i'} \left[Q_i(s) - Q_i(t_i')\right] ds$$

- But since (4) Q_i is weakly increasing, the integrand is nonpositive, so the gain is nonpositive
- For the case where $t_i' < t_i$,

$$gain = -\int_{t_i'}^{t_i} Q_i(s)ds + (t_i - t_i')Q_i(t_i') = \int_{t_i'}^{t_i} \left[Q_i(t_i') - Q_i(s)\right] ds$$

  which is once again nonpositive because Q_i is nondecreasing

- So together, 4 and 5 imply 3 – and we're done.

That proves the lemma.

7 Formulating the Seller's Problem

- We'll, we've now established that:

    1. we can accomplish something with any mechanism if and only if we can accomplish with a direct mechanism
    2. we can accomplish something with a direct mechanism if and only if we can accomplish it with a mechanism satisfying conditions (4), (5), (6), and (7)

- So our task is now, very literally, to solve the seller's problem: to maximize expected revenue (or really, expected seller profit) over all possible mechanisms (p, x), subject to the constraints (4), (5), (6), and (7).

- We can write the seller's expected payoff as

$$E_t \left\{ \sum_{i \in N} x_i(t) + t_0 \left(1 - \sum_{i \in N} p_i(t) \right) \right\}$$

  and so our problem is to maximize this, over all direct-revelation mechanisms, subject to (4), (5), (6), and (7)

- And that's where we'll start next week.