Mechanism Design

- Setup:
- n agents $i = 1, \dots, n$
- agent i has type $\theta_i \in \Theta_i$ which is i 's private information
- $\theta = (\theta_1, \dots, \theta_n) \in \Theta = \prod_i \Theta_i$
- We denote $\theta_{-i} = (\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_n)$
- $\theta = (\theta_i, \theta_{-i})$
- $y \in Y$ is a decision to be taken by the principal P
- E.g.: $y = (x, t)$, where x is the allocation (who gets the good in an auction; how much of a public good is built; etc) and t is the transfer (how much people pay/are paid)

Mechanism

- A mechanism $\Gamma = \{M, y\}$ specifies a message space M and a decision rule $y(m)$
- Each agent sends a message $m_i(\theta_i)$ to P from message space M_i, and then P chooses action $y(m_1, \dots, m_n)$
- P has commitment power

Preferences

- Agent i has utility $u_i(y, \theta)$
- P has utility $v(y, \theta)$
- (Note: i's utility can depend on other players' types, but in some examples it will only depend on her own type, $u_i(y, \theta_i)$)

Beliefs

- $p(\theta)$ is a common prior belief
- Players have posteriors given their type $p(\theta_{-i} \mid \theta_i)$ derived from their prior

Timing

(1) P chooses a mechanism (M, y) and commits to it

(2) Agents play the "game" induced by the mechanism, with equilibrium $m^*(\theta) = \left(m_1^*(\theta_1), \dots, m_n^*(\theta_n)\right)$

(3) Outcome $\tilde{y}(\theta) = y(m^*(\theta))$

For now we will be agnostic about the equilibrium concept used to determine m^*

Questions

- Which allocations $\tilde{y}(\theta)$ can be implemented? (Depending on the solution concept)
- Which $\tilde{y}(\theta)$ among the implementable ones is optimal for P ?
- E.g.: in our screening problem, $\tilde{y} = (x(\theta), t(\theta))$ and we could implement any non-decreasing schedule $x(\theta)$ (but with restrictions on $t(\theta)$)

Two Solution Concepts

- DSE (Dominant Strategy equilibrium): i has a best strategy independently of the other agents' types (even if I knew their types)
- BNE (Bayesian Nash equilibrium)

Revelation Principle

BNE version: suppose Γ has BNE $m^*(\theta)$ with outcome $\tilde{y}(\theta) = y(m^*(\theta))$.

Then there exists a direct revelation mechanism Γ^d with $M = \Theta$ and $y^d(\theta) = \tilde{y}(\theta)$, such that $m_i^d(\theta_i) = \theta_i$ is BNE-implementable.

- In a direct mechanism, P just asks agents to reveal their type, and chooses some allocation accordingly
- It is incentive-compatible for agents to tell their true type
- The revelation principle says that decision rule $\tilde{y}(\theta)$ is implementable with some mechanism (M, y) iff truth-telling is a BNE of mechanism $(\Theta, \tilde{y})$
- This greatly reduces the space of mechanisms we need to study
- We already saw the revelation principle in our screening problem:
- A solution was initially framed as a payment schedule $t(x)$, which would induce some equilibrium production $x(\theta)$ by the agent

- But we reframed it as directly choosing $(x(\theta), t(\theta))$ for each θ, subject to IC and IR conditions
- Note: P's commitment power matters
- If P did not have commitment power it would be hard to get agents to reveal θ since it might allow for more deviations ex post by P
- The TSA has rules to punish people detected to have drugs
- In the direct mechanism version, you would always tell the truth, and you would not get punished if you had some amount that they would not have detected anyway
- But they don't have the commitment power to do this: if you say "yes, I have five grams of cocaine" you will go to jail

Proof

- "If" direction is obvious: if truth telling is a BNE of mechanism $(\Theta, \tilde{y})$, then this mechanism implements allocation $\tilde{y}(\theta)$
- "Only if": start with general (M, y)
- If m^* is a BNE, then $m_i^*(\theta_i) \in \text{argmax}_{m_i}\ E_i[u_i(y(m_i, m_{-i}^*(\theta_{-i}), \theta) \mid \theta_i]$
- That is:

$$E_i[u_i(y(m_i^*(\theta_i), m_{-i}^*(\theta_{-i}), \theta) \mid \theta_i] \geq E_i[u_i(y(m_i', m_{-i}^*(\theta_{-i}), \theta) \mid \theta_i]$$

- In particular, one may take $m_i' = m_i^*(\tilde{\theta}_i)$:

$$E_i\left[u_i(y(m_i^*(\theta_i), m_{-i}^*(\theta_{-i}), \theta) \mid \theta_i\right] \geq E_i\left[u_i\left(y\left(m_i^*(\tilde{\theta}_i), m_{-i}^*(\theta_{-i}), \theta\right) \mid \theta_i\right]\right.$$

- for any $\tilde{\theta}_i$: no point in mimicking any other type $\tilde{\theta}_i$
- Hence:

$$E_i[u_i(\tilde{y}(\theta_i, \theta_{-i}), \theta) \mid \theta_i] \geq E_i\left[u_i\left(\tilde{y}\left(\tilde{\theta}_i, \theta_{-i}\right), \theta\right) \mid \theta_i\right] \text{ for all } \tilde{\theta}_i$$

- Then $\theta_i \in \text{argmax}_{\tilde{\theta}_i}\ E_i\left[u_i\left(\tilde{y}\left(\tilde{\theta}_i, \theta_{-i}\right), \theta\right) \mid \theta_i\right]$, so truth-telling is an equilibrium of $(\Theta, \tilde{y})$

DSE

- Same theorem holds for the DSE solution concept
- Here, m^* is a DSE if

$$m_i^*(\theta_i) \in \text{argmax}_{m_i}\ u_i(y(m_i, m_{-i}), \theta)$$

- for any m_{-i}
- Notes: DSE implies BNE
- Revelation principle is a "testing device"
- Commitment is again critical
- More general mechanisms may be useful for unique implementation

Vickrey-Clarke-Groves (VCG) mechanism

- VCG is a DSE implementation of any decision rule
- The catch: it is not necessarily budget-balanced
- $y(x, t)$ allocation
- $t = (t_1, \dots, t_n)$ transfers
- E.g.: x is a public good, or $x = (x_1, \dots, x_n)$ is an allocation of private goods
- $u_i(y, \theta) = u_i(x, \theta_i) + t_i$: quasilinear preferences
- First-best allocation:

$$x^*(\theta) \in \text{argmax} \sum_i u_i(x_i, \theta_i)\,, \forall \theta$$

- Question: can $x^*(\theta)$ be implemented?
- Yes
- Counterintuitive: it seems like in real life it is very hard to get people to reveal preferences for a public good and build it whenever optimal
- DSE means that

$$\theta_i \in \text{argmax}_{m_i}[u_i(x^*(m_i, m_{-i}), \theta_i) + t_i(m_i, m_{-i})]\ \forall \theta_i, m_{-i}$$

- Note: DSE requires that declaring your true type is optimal even if other people are lying and sending whatever messages m_{-i}

- By definition of x^*,

$$\theta_i \in \text{argmax}_{m_i} \left[u_i(x^*(m_i, m_{-i}), \theta_i) + \sum_{j \neq i} u_j\left(x^*(m_i, m_{-i}), m_j\right) \right] \quad \forall \theta_i, m_{-i}$$

- Since sending $m_i = \theta_i$ implements the socially optimal x^* (assuming other players' types are given by m_j)

- Idea: we can just set the transfers for player i equal to all the remaining terms!

$$t_i^{VCG}(m_i, m_{-i}) = \sum_{j \neq i} u_j\left(x^*(m_i, m_{-i}), m_j\right) + h_i(m_{-i})$$

- Then i 's incentives are always to implement $x^*(\theta_i, m_{-i})$, so he has a weakly dominant strategy to announce $m_i = \theta_i$
- That is, we turned the individual agent's problem into Pareto's problem.
- h_i is any function that depends on m_{-i} and hence does not affect i 's incentives
- May be useful if we want transfers to add up to 0

Uniqueness

- Not only does VCG implement x^*
- But it is also essentially the unique mechanism that does this

If Θ_i is "smoothly connected" $\forall i$, then $\{t_i^{VCG}\}$ uniquely implements $x^*(\theta)$ (up to "constants" $h_i(m_{-i})$).

- Smoothly connected means that, for any $\theta_i, \theta_i' \in \Theta_i$, there is a curve c: $[0,1] \to \Theta_i$ s.t. $c(0) = \theta_i, c(1) = \theta_i', c$ is C^2 and $u \circ c$ is C^2

Example 1

- Suppose $x = 1$ or 0: build or not build
- Building has social cost K (for simplicity $K = 0$)
- θ_i is i 's willingness to pay
- $x^*(\theta) = 1$ if $\sum_i \theta_i \geq K$ and 0 otherwise

- Then what are the VCG transfers?
- $t_i^{VCG}(m) = 0$ if i 's WTP is not pivotal
- $t_i^{VCG}(m) = \sum_{j \neq i} \theta_j \leq 0$ if i is pivotal for $x = 1$
- $t_i^{VCG}(m) = -\sum_{j \neq i} \theta_j$ if i is pivotal for $x = 0$
- Idea: i always pays for the externality of his message
- Our example above is called a pivot scheme
- It implies a particular choice of h_i :

$$h_i(m_{-i}) = -\max_x u_{j \neq i}(x, m_j)$$

- In particular this choice of h_i guarantees that $\sum_i \ t_i(m_i, m_{-i}) \leq 0$ for all m (the principal never has to pay money on net)

Example 2

- Second price auction
- n buyers, each i has value θ_i, submits bid b_i (simultaneous bids)
- Highest bid gets the good, highest bidder pays second highest bid
- Check: this is a pivot scheme
- This seems too easy; what is the catch?
- To get the right decision, the mechanism generates very steep incentives
- In reality, this makes it hard to satisfy the IR of all participants, if they have any
- If we choose h_i as in the pivot scheme, agents may get very negative payoffs in some states, so their IR may be violated (especially if they know their θ_i before agreeing to the mechanism, in which case they would have a limited liability constraint)
- If we increase transfers to agents so their IRs are satisfied, the principal may have to pay a lot in some states

Balancing the Budget

- Suppose the principal does not want to pay or be paid money for setting up the mechanism

- In other words, we want $\sum_i t_i(m) = 0$ for all m
- When can we do this?
- Let $S(\theta) = \sum_i u_i(x^*(\theta), \theta_i)$
- Suppose to begin that we take $h_i(m_{-i}) \equiv 0 \ \ \forall i, m_{-i}$
- Then $\sum_i t_i^{VCG} = (n-1)S(\theta)$: massive deficit
- Taking h_i as in the pivot scheme gives $\sum_i t_i^{VCG}(m) \leq 0$ (budget surplus), but not $\equiv 0$
- Solution: we can take h_i such that $\sum_i t_i(m) = 0 \ \forall m$ iff we can write

$$S(m) = f_{i=1}^{n} f_i(m_{-i})$$

- for some functions f_i
- If this is true, we can set $h_i(m_{-i}) = -(n-1)f_i(m_{-i})$
- Then $\sum_i t_i(m) = (n-1)S(m) - (n-1) \ {}_i f_i(m_{-i}) \equiv 0$
- This condition is also sufficient: if $\sum_i t_i(m) \equiv 0$, then $(n-1)S(m) + \sum_i h_i(m_{-i}) \equiv 0$, so we can use $f_i = -\frac{h_i}{n-1}$
- How hard is this condition to satisfy?
- In our public good example, $S(\theta) = \sum_i \theta_i$ or $S(\theta) = 0$
- This S satisfies the condition: can take $f_i = -\frac{\sum_{j \neq i} \theta_j}{n-1}$
- Another case where it is satisfied is if you add an agent $n+1$ who does not care about the outcome, so we can set $S(m) = -f_{n+1}$
- But it's hard in general

BNE Implementation

- With BNE implementation, we want to satisfy

$$E_{\theta_{-i}}[u_i(x^*(\theta), \theta_i) + t_i(\theta)] \geq E_{\theta_{-i}}[u_i(x^*(m_i, \theta_{-i}), \theta_i) + t_i(m_i, \theta_{-i})]$$

- If we assume types are independent, the RHS can be written as

$$\bar{u}_i(m_i, \theta_i) + \bar{t}_i(m_i)$$

- where $\bar{u}_i$ is expected utility from the allocation and $\bar{t}_i$ is the expected transfer
- These are not conditioned on θ_{-i} because we are taking expectations (and if types are independent, θ_i is uninformative about θ_{-i})
- In this case it is easier to balance the budget because BNE implementation requires fewer constraints on the t_i
- If we choose t_i^{VCG} then x^* is DSE-implementable so in particular it is BNE-implementable, but we can then tweak the transfers further without breaking BNE implementation

Summary

- We covered how to generally implement the optimal allocation with the VCG mechanism
- Intuition: use transfers so that each i 's incentives are identical to the social planner's
- Have to pay i for the externality that his decision generates on everyone else
- Caveat: this mechanism runs a massive budget deficit
- Can fix it by just lowering all the transfers so the planner runs a surplus (e.g. pivot scheme)
- But getting the budget to be always balanced can only be done if the surplus function $S(\theta)$ satisfies a separability property
- We then moved on to BNE implementation
- The Bayesian IC condition is now:

$$E_{\theta_{-i}}[u_i(x^*(\theta), \theta_i) + t_i(\theta)] \geq E_{\theta_{-i}}[u_i(x^*(m_i, \theta_{-i}), \theta_i) + t_i(m_i, \theta_{-i})]$$

- Assuming independent types, this can be rewritten as

$$\bar{u}_i(\theta_i, \theta_i) + \bar{t}_i(\theta_i) \geq \bar{u}_i(m_i, \theta_i) + \bar{t}_i(m_i)$$

- Note: BNE implementation has many fewer IC constraints
- With DSE implementation, need constraints $\text{IC}_{\theta_i, m_i, m_{-i}}$ for all θ_i, m_i, m_{-i}
- $\text{IC}_{\theta_i, m_i, m_{-i}}$ says that type θ_i prefers to send a truthful message rather than reporting m_i, when other players send m_{-i}
- With BNE, i does not know m_{-i} and just cares about the effect of his message under the expected m_{-i}

- So only has conditions $\text{IC}_{\theta_i, m_i}$
- This allows us to pick non-VCG transfers and still implement the same allocation

Budget Balancing with BNE

- Can we use this new freedom to still implement x^* while balancing the budget?
- Yes
- Pick transfers

$$t_i^{AGV}(m) = \bar{t}_i^{VCG}(m_i) - \frac{1}{N-1} - \underset{j \neq i}{VCG}(m_j)$$

- Then

$$t_i^{AGV}(m) = 0 \ \forall m$$

- From i 's point of view, $\bar{t}_i^{AGV}(m_i) = \bar{t}_i^{VCG}(m_i) +$ constant, so it generates the same incentives as VCG: the extra terms cannot be influenced by i's message
- Note: this works for BNE implementation because t_i^{AGV} gives the right incentive for the expected m_{-i}
- If we wanted DSE implementation, t_i would have to make $m_i = \theta_i$ IC for every m_{-i} possible
- So t_i would have to condition on $m = (m_i, m_{-i})$ jointly
- This would make it impossible to funnel other t_j into a function $h_i(m_{-i})$, which is what we are doing now

Caveats

- However, BNE implementation has its own set of problems, so not necessarily more realistic than DSE implementation
- This only works under independent types
- Types may well be correlated in reality
- This also requires that the players have common knowledge of the distribution of everyone's type
- DSE implementation does not rely on this
- Finally, mechanisms which BNE implement an allocation may also have other equilibria

Envelope Theorem

- We will use the envelope theorem to study implementation in the continuous case
- Let $\theta \in [0,1]$ state of the world
- X arbitrary choice set
- Agent with utility $u(x, \theta)$
- Maximized utility $U(\theta) \equiv \sup_{x \in X} u(x, \theta)$
- Optimal choice $X^*(\theta) = \text{argmax}_{x \in X}\ u(x, \theta)$
- $x^*(\theta) \in X^*(\theta)$ is a selection

Theorem (Envelope Theorem in Integral Form)

Assume:

- $u(x, \theta)$ is differentiable in θ for all $x \in X$
- There is $B < \infty$ such that $|u(x, \theta)| \leq B$ for all x, θ
- $X^*(\theta) \neq \emptyset$ for all θ

Then

$$U(\theta) = U(0) + \int_0^{\theta} u_\theta\big(x^*(\theta), \tilde{\theta}\big) d\tilde{\theta}$$

and

$$U'(\theta) = u_\theta\big(x^*(\theta), \tilde{\theta}\big)$$

exists for all $\theta \in [0,1]$.

- Note the statement is completely agnostic about the set X and the behavior of u with respect to x
- No assumption that X is an interval, or connected, or even made up of real numbers
- No assumption that u is differentiable or even continuous with respect to X

Continuous BNE Implementation

- Let $E\big(t_i(m_i, \theta_{-i})\big) = \bar{t}_i(m_i)$
- Let $E(u_i(x^*(m_i, \theta_{-i}), \theta_i) = \bar{u}_i(m_i, \theta_i)$
- Then

$$U_i(\theta_i) = U_i(0) + \int_0^{\theta_i} \frac{\partial \bar{u}_i}{\partial \theta_i}(\tilde{\theta}_i, \tilde{\theta}_i) d\tilde{\theta}_i$$

- In other words, $U_i(\theta_i)$ is completely pinned down by the allocation
- Hence, any two schemes $t_i(m), \hat{t}_i(m)$ which implement the allocation must satisfy $\bar{t}_i(m_i) = \bar{\hat{t}}_i(m_i) + \text{constant}$
- In other words, $E\big(t_i(m_i, \theta_{-i})\big) = E\left(t_i^{VCG}(m_i, \theta_{-i})\right) + \text{constant}$
- This means that there is essentially no way to improve on VCG, even if you just want BNE implementation
- (Besides the fact that with VCG you can tweak the actual t_i, so long as you maintain the resulting $\bar{t}_i$, and this may be useful for budget balancing)
- This dashed the hopes of computer scientists that hoped to come up with better implementation mechanisms

Revenue Equivalence Theorem

- The Revenue Equivalence Theorem is a consequence of this analysis
- It says that, if two mechanisms implement the same allocation, and the payoff of each i's lowest type is the same under both mechanisms, then the expected payoff of every type of every player is the same under both mechanisms
- And the mechanism's designer expected surplus is also the same
- In other words, if both mechanisms have the same x^*, and the same $U_i(0)$ for every i, then they have the same $U_i(\theta_i)$ for every i, θ_i, and the same expected surplus $-E(\ _i t_i)$

RET Example

- The RET has important consequences for auctions
- Compare a first and second price auction with symmetric buyers, with continuous independent types θ_i distributed on an interval
- Bids will be different (in first price auction, buyers underbid to increase their profit)
- But both will end up giving the good to the highest bidder, which is the buyer with highest value: same $x(\theta)$
- Lowest type never wins, so payoff 0 in both cases
- RET: both auctions must generate the same expected revenue! (both for the auctioneer and for every type of every player)

Summary

- Reminder: we had seen how to BNE-implement the optimal allocation x^*
- We constructed t_i^{AGV}, which balanced the budget and BNE-implemented x^*
- In particular, t_i^{AGV} was the same as VCG in expectation, in other words $\bar{t}_i^{AGV}(m_i) = \bar{t}_i^{VCG}(m_i)$
- But $t_i^{AGV}(m) = t_i^{VCG}(m)$

Myerson-Satterthwaite Theorem

- Can we achieve efficient bilateral trade between two agents with private information about their values?
- M-S theorem: no!
- How come?

Setup

- 2 agents B, S
- One good
- $x_S + x_B = 1$: $x_S = 1$ means sell, $x_S = 0$ means don't sell
- Payoffs: $u_S = t_S - x_S\theta_S, u_B = x_B\theta_B + t_B$
- $\theta_i \sim F_i$ independent, with full support
- Assume the supports overlap, so exchanging may or may not be efficient

Requirements:

(1) Efficiency: $x_S(\theta) = 1$ iff $\theta_B \geq \theta_S$ (the mechanism should result in trade whenever it is welfare-improving)

(2) Ex ante budget balance: $E(t_S(\theta) + t_B(\theta)) \leq 0$ (the principal does not lose money on average)

(3) (Interim) Individual rationality: $EU_i(\theta_i) \geq 0$ under the mechanism, for every i and θ_i

M-S: no mechanism can satisfy all requirements

- Why doesn't our BNE implementation theory contradict the M-S theorem?
- Note that requirements 2 and 3 differ from our usual assumptions
- 2 is actually quite weak: in BNE implementation, we can balance the budget exactly for every θ; here, we just require expected balance (or surplus)

- But 3 is strong and we never had that condition before
- In BNE implementation, we never required that each agent get some minimum expected utility
- Here we have a stronger condition: agents must not want to pull out after knowing their type

M-S Theorem Proof Sketch

- Start with pivot scheme
- This is a VCG mechanism, so implements the efficient allocation ($x_S = 1$ iff $\theta_B \geq \theta_S$)
- It gives the transfers: $t_S = 0$ iff $x_S = 0, t_S = \theta_B$ iff $x_S = 1$
- $t_B = 0$ iff $x_S = 0, t_B = -\theta_S$ iff $x_S = 1$
- Pivot scheme runs a budget deficit: $E(t_B(\theta) + t_S(\theta)) = E\big(\max(\theta_B - \theta_S, 0)\big) > 0$
- We could change it-how?
- One thing we can do is decrease transfers by a fixed amount: set $\tilde{t}_B(\theta) = t_B(\theta) - C_B$, or $\tilde{t}_S(\theta) - C_S$ for some $C_B, C_S > 0$
- This does not affect incentives, but is impossible because of the IRs
- Already with the pivot scheme, $U_S(1) = 0$ and $U_B(0) = 0$, so setting $C_B > 0$ or $C_S > 0$ would violate IR for some types
- Can we change the transfers in some θ-dependent way?
- Yes-if we just want BNE implementation, we can change the t_i in any way that preserves $\bar{t}_i(m_i)$
- However, any change to the t_i which preserves $\bar{t}_i(m_i)$ for each m_i, will also preserve $E_{m_i}\big(\bar{t}_i(m_i)\big) = E_m(t_i(m))$
- Hence such changes will not affect $E(t_S(\theta) + t_B(\theta))$!
- And so any such change cannot fix the expected budget deficit
- Does that really finish the proof? Yes
- Because we are leveraging another powerful result we already know: that any mechanism implementing x^* must have the same $\bar{t}_i$ as VCG, up to a constant

MHT

Consider the following related team production problem:

- N agents
- $x = f(e_1, \dots, e_n) = e_1 + \cdots + e_n$ is total production (a function of agents' efforts)
- $s_i(x)$ payment to agent i
- $u_i = s_i(x) - c_i(e_i)$

Note: no uncertainty or private information

Can you satisfy:

(1) Efficiency

(2) Nash Equilibrium

(3) Budget Balance: $\sum_i s_i(x) = x$ for all x

Answer: no, under some conditions

Proof Sketch

- The efficient allocation satisfies $\frac{\partial f}{\partial e_i} = \frac{\partial c_i}{\partial e_i}$ for all i
- Nash equilibrium requires that e_i solve

$$\max_{e_i} \{s_i(f(e_1, \dots, e_n)) - c_i(e_i)\}$$

- So $\frac{\partial s_i}{\partial x}\frac{\partial f}{\partial e_i} = \frac{\partial c_i}{\partial e_i}$
- Using the efficiency condition, we get $\frac{\partial s_i}{\partial x} = 1$ for all i
- Here is the contradiction: we must have $s_i'(x^*) = 1$ for all i
- But $\sum_i s_i(x) = x$ for all x requires that $\sum_i s_i'(x^*) = 1$ instead
- In other words, I need much stronger incentives than I can provide
- However, you can solve the contradiction if you allow

$$s_i(x) \leq x$$

for all x instead (budget surplus)

- Then you can take $s_i(x) = x - \frac{N-1}{N}x^*$ for x up to x^*, and $s_i(x) = \frac{X}{N}$ thereafter

- Idea: incentives are weaker for $x > x^*$, but that's fine because we are trying to implement a fixed x^* (no uncertainty)
- For $x < x^*$ I create steep incentives by using a steep punishment
- If anyone screws up, everyone pays for it (team punishment)
- This does not result in low utility for the agents (IR problems) because the punishment only happens off the equilibrium path
- But, when types are random, everything can happen on the equilibrium path
- Note: incentive problems can be created by informational externalities even if there isn't joint production
- In this example, the production function is additive (no interaction between agents)
- But still hard to incentivize simply because the principal doesn't observe individual outputs

The Market for Lemons

- A lemon is a used car that is not very good
- Idea: S owns a used car and wants to sell it
- S knows whether the car is a lemon or a peach, but B can't tell
- Suppose $v \sim U[5000,10000]$ where v is the value of the car to the seller
- B's value is $v + \Delta$, where $\Delta = 1000$
- So $v_B \sim U[6000,11000]$, but unlike our previous models, here the values are correlated
- Even though supports overlap, trade is always efficient
- But can they trade?
- Suppose S offers to sell for 7500
- B can infer that, if 7500 is the market price, then sellers with value above 7500 would never actually sell (they would rather keep the car)
- And sellers with value below 7500 would sell
- So the offer must come from the latter group, which has mean value 6250

- Hence $E(v_B \mid \text{offer}) = 7250 < 7500$, and B would not buy
- What is the equilibrium price?
- It must be v such that $\frac{p+5000}{2} + 1000 = p$, so $p = 7000$
- Hence 60% out of the efficient trades do not happen
- In general $p = 5000 + 2\triangle$: the smaller B's extra value, the lower the equilibrium price
- For small Δ, most of the market unravels
- This market unraveling problem creates incentives for people to signal
- The seller may let you take the car to a third party mechanic, or do a test drive, or give you a guarantee
- But without such signals, the information problem has big consequences