Hotelling Lectures: Evidence in Games and Mechanisms
Part 3: Other Directions

Barton L. Lipman
Boston University

November 2023

Costly Verification

Have assumed evidence is costless to receive and the agent determines whether it is provided.

Alternative: Principal has costly access to evidence.

One–agent mechanisms: Townsend (1979), Gale and Hellwig (1985), Border and Sobel (1987), Mookherjee and Png (1989); Glazer and Rubinstein (2004).

Multi–Agent Mechanisms: BDL (*AER*, 2014)

Consider the simple allocation problem again but with costly verification.

Principal has one good to allocate, no monetary transfers.

Agent i has privately–known type t_i where $v_i(t_i)$ is the value to the principal from giving the good to i.

Temporarily simplify notation to think of types as v_i's.

v_i's independent, continuous distribution F_i on $[\underline{v}, \bar{v}]$.

All agents want the good.

Principal can pay cost c to obtain perfect evidence about i's type. (Easy to generalize to c_i.)

Result (again): Favored–agent mechanism.

Have favored agent i^* and threshold value v^* such that

- If every non–favored agent reports a value below v^*, then i^* gets the object and no verification is carried out.
- Otherwise, agent with highest reported value is checked and gets the object if (as will happen in equilibrium) report is found to be correct.

Characterizing Favored Agent and Threshold

For each i, define v_i^* by

$$\mathrm{E}(v_i) = \mathrm{E}\max\{v_i, v_i^*\} - c.$$

Results:

- If i is favored agent, best threshold is v_i^*.
- Optimal favored agent is any i with largest v_i^*.

Intuition: Suppose i is favored. Compare thresholds ν and v_i^*, where $\nu > v_i^*$.

Let x be highest report of agent other than i.

	$x < v_i^* < \nu$	$v_i^* < x < \nu$	$v_i^* < \nu < x$
v_i^*	$\mathrm{E}(v_i)$	$\mathrm{E}\max\{v_i, x\} - c$	$\mathrm{E}\max\{v_i, x\} - c$
ν	$\mathrm{E}(v_i)$	$\mathrm{E}(v_i)$	$\mathrm{E}\max\{v_i, x\} - c$

$x > v_i^*$ implies

$$\mathrm{E}\max\{v_i, x\} - c > \mathrm{E}\max\{v_i, v_i^*\} - c = \mathrm{E}(v_i).$$

So v_i^* is a better threshold than ν.

Why is best favored agent the one with the highest v_i^*?

Very partial intuition: This i requires best showing among others to not just give it to i.

As before, mechanism satisfies robust IC and is deterministic.

Detail

But commitment is needed: if not committed, principal would not verify ex post.

Jump

Robust incentive compatibility: for every i and v_i, truth is optimal given *any* strategies of others.

Suppose i is not favored.

If $v_i > v^*$:

- Truth: i gets object if all reports by others below v_i.
- Any lie: i cannot get the object.
- So truth weakly dominates any lie.

If $v_i < v^*$: No matter what i does, she can't get object. So truth is an optimal strategy.

Suppose i is favored and $v_i > v^*$.

- If all other reports are below v_i, she gets the good if she reports truthfully and can't improve on that.
- If some other agent is above v_i, she can't get good regardless.
- So truth weakly dominates any lie.

Suppose i is favored and $v_i < v^*$.

- If some other agent is above v^*, she can't get the good regardless.
- If all other agents below v^*, she gets the good regardless.
- Her report is irrelevant, so truth telling is an optimal strategy. Back

Similarity between mechanisms with Dye evidence and costly verification is more general.

Erlanson and Kleiner (*TE*, 2020) consider costly verification in a public good problem and BDL (2019) show their mechanism "looks like" optimal mechanism with Dye evidence.

BDL (2019) shows that for a subclass of problems considered, one can solve for optimal mechanism in Dye evidence and do "change of variables" to get optimal mechanism for costly verification.

For simple allocation problem: $P_i(t) =$ probability principal gives good to i given reports t and not learning i lied.

$Q_i(t) =$ probability principal verifies i given reports t.

So principal's objective function is

$$\mathrm{E}_t \left[\sum_i (P_i(t) v_i(t_i) - Q_i(t) c_i) \right] = \sum_i \mathrm{E}_{t_i} [\hat{p}_i(t_i) v_i(t_i) - q_i(t_i) c_i]$$

where $\hat{p}_i(t_i) = \mathrm{E}_{t_{-i}} P_i(t_i, t_{-i})$ and $q_i(t_i) = \mathrm{E}_{t_{-i}} Q_i(t_i, t_{-i})$.

Principal never verifies unless he'll give good to agent.

Hence incentive compatibility constraint is:

$$\hat{p}_i(t_i) \geq \hat{p}_i(t_i') - q_i(t_i'), \ \forall t_i, t_i'$$

$$\min_{t_i} \hat{p}_i(t_i) \geq \hat{p}_i(t_i') - q_i(t_i'), \ \forall t_i'$$

$$q_i(t_i') \geq \hat{p}_i(t_i') - \min_{t_i} \hat{p}_i(t_i), \ \forall t_i'.$$

$$q_i(t_i') = \hat{p}_i(t_i') - \min_{t_i} \hat{p}_i(t_i), \ \forall t_i'.$$

$$q_i(t_i') = \hat{p}_i(t_i') - \min_{t_i} \hat{p}_i(t_i), \ \forall t_i'.$$

Can show type of i who gets good least often is type with lowest v_i — call this t_i^0.

Substituting into objective function, principal maximizes

$$\mathrm{E}_t \sum_i \left[\hat{p}_i(t_i)(v_i(t_i) - c_i) + \hat{p}_i(t_i^0)c_i \right].$$

Rewrite as

$$\max \mathrm{E}_t \sum_i \hat{p}_i(t_i)\tilde{v}_i(t_i)$$

subject to $\hat{p}_i(t_i^0) \leq \hat{p}_i(t_i)$.

$$\max \mathrm{E}_t \sum_i \hat{p}_i(t_i)\tilde{v}_i(t_i) \tag{1}$$

subject to $\hat{p}_i(t_i^0) \leq \hat{p}_i(t_i)$.

Consider Dye evidence problem where values are $\tilde{v}_i(t_i)$ and t_i^0 is only type without evidence.

Objective function is equation (1).

Incentive constraint is types provide maximal evidence — i.e., no one pretends to be t_i^0. Same constraint as above.

Take the Dye solution, substitute in for $\tilde{v}_i$, rearrange, and you get the favored–agent mechanism for costly verification.

Other papers: Mylovanov and Zapechelnyuk (*AER*, 2017), Ball and Kattwinkel (2019), Li (*JET*, 2020), Li and Libgober (2023).

Game–Theoretic Approaches to Evidence Acquisition

Che and Kartik (JPE, 2009).

Receiver can choose a sender to get evidence from.

Initially, receiver and sender symmetrically informed and may have different priors about state, t.

Sender can spend resources to try to learn and get evidence about t.

$u(a, t) = v(a, t) = -(a - t)^2$.

No conflict of interest in utilities. Both want to set $a = t$.

If sender spends $c(q)$, she has probability q of getting evidence in the form of a noisy signal about t, otherwise nothing.

Dye with noise and endogenous q.

Prior over t for receiver $N(0, \sigma^2)$, for sender is $N(\mu_S, \sigma^2)$.

Main result: Receiver prefers sender with $\mu_S \neq 0$.

Intuition: If sender reveals a signal about t, receiver picks her best a given her prior and this signal.

If sender doesn't show proof, receiver picks a based entirely on her prior.

If $\mu_S \neq 0$, sender doesn't always reveal signal.

But sender also works harder to get one: the conflict is smaller given information than given priors.

This creates the key tradeoff: If priors are further apart, sender will be informed more often but will distort disclosure decision more often.

Optimal tradeoff has best μ_S different from 0.

Chade and Pram (2023).

Paper not available yet, so my information limited.

Policy issues:

Should SAT/college entrance exams be required?

A less severe restriction: If an applicant has taken the test, should they be required to disclose the results?

Model: Population of students and of schools.

Schools vary in quality; this is observed.

Students vary in quality; this is not observed.

If quality were known, matching would be assortative: better students would go to better schools.

If quality is unknown, do same with *expected* quality.

Students' match utility increasing in the quality of the school they go to.

Students' types are noisy signals of their quality. Types are privately known, no evidence about them available.

Students can pay to take a test which provides a better signal of quality and result is evidence.

Three Cases:

Case 1. Test is voluntary, disclosure is voluntary.

In equilibrium, students with high enough types will take the test.

Students who take the test and get high enough results will disclose.

Students with low results pool with those who didn't take the test.

A different way to endogenize q from Dye.

One difference: Not having evidence is itself a bad signal.

Case 2. Test is mandatory, disclosure is mandatory.

I think result same is test is mandatory and disclosure voluntary because of unraveling.

Comparison to previous case: Low types worse off since they have to pay for a test they'll do poorly on.

Can have higher types better off, but "unexpected" complexity in which higher types are better off and which are worse off.

Intuition?: Suppose under voluntary, I take the test.

1. Suppose my score is in range where I don't disclose.

So under voluntary, I'm tied with a mass of people at the bottom, but disclosing would put me below them.

Now I have to disclose, but everyone else does too.

The mass of types I was tied with will have a distribution of scores with a mean above me.

Whether I'm better off or worse off depends on how high my score was and the asymmetry of the distribution of the new scores.

2. Suppose my score is in the range where I disclose.

So I'm above this mass of people who don't disclose, beating all of them.

If this group is forced to take and reveal the test, at least some will beat me, so I'm worse off.

So there are situations in which I'm better off under mandatory and situations where I'm worse off.

As my type increases, probability of these situations change in complicated ways, creating ambiguous signs.

Case 3. Test is voluntary, disclosure is mandatory.

This is the case where fewest people take the test.

The lowest types prefer this case to either of the others.

Reason: When they don't provide a test result, school knows they didn't take the test — can rule out possibility they took it but did poorly.

Other papers: DeMarzo, Kremer, and Skrzypacz (*AER*, 2019), Shiskin (2022).

Mechanism Design with Evidence Acquisition (BDL, 2023)

Simple allocation problem: Principal has one unit of a good to allocate to one of N agents.

Value to principal of giving good to agent i is v_i, value to the agent is 1.

Change: We start with no one knowing any v_i's.

Agent i can learn v_i at cost $c \in (0, 1)$. Learning provides proof.

One option: Principal asks all agents to get proof and gives good to the best.

Assume distribution of v_i same across agents.

Suppose $c > 1/N$. Then principal can't get all agents to respond.

So what is the optimal mechanism?

Optimal mechanism: Principal picks an agent at random and asks them to learn v_i.

If $v_i \geq v^*$, principal stops and gives good to that agent.

Otherwise, principal picks second agent at random and continues.

If all agents have $v_i < v^*$, principal will get information from all and then will give to agent with highest v_i.

Key: When agent is asked, they aren't told how many others were already asked.

Intuition: Suppose principal does reveal order. If you know you're last, you know the others all have $v_i < v^*$.

Hence you get good if you're above the threshold *or* with probability $1/N$ if you're below.

But if you know you're first, you get the good if you're above the threshold *or* with probability $1/N$ if *everyone* is below.

So if first has sufficient incentive, we have slack on the rest and could improve.

What happens when costs and/or distributions of values differ?

Three changes:

First: Compare *virtual values*, not values. $v_i + \lambda_i$.

Intuition: Need to give an advantage to agents who are harder to incentivize.

λ_i will be Lagrange multiplier on incentive constraint for i.

Second: Probability distribution over order in which agents asked for evidence not uniform.

Two reasons:

1. Put less pressure on agents who are hard to incentivize by asking them later.
2. May want to get to "better" agents earlier.

Third: May lower standards over time.

Start with "tier 1" agents and proceed as before with a certain threshold.

If none of them are above it, *lower* the threshold.

If some are above the lower threshold, give to the best.

Otherwise, continue to next tier and ask agents for evidence, comparing to the lower threshold.

Intuition: Extreme version of second point.

Model

- N agents and one principal.
- Principal has one unit of a good to give to an agent.
- Value to principal of giving good to i is v_i; value to i of receiving it is 1.
- v_i unknown ex ante. $v_i \sim [0, 1]$ with density $f_i > 0$, cdf F_i. Independent.

- Cost to i of learning v_i is $c_i \in (0, 1)$.
- Learning v_i gives evidence allowing i to prove v_i to principal.

Mechanisms

Consider multistage mechanisms where at each stage, as function of history so far, the principal does one of the following:

- stops process and keeps the good
- stops process and gives good to some agent
- asks some agent to obtain and provide evidence.

By Revelation Principle, mechanism gives agents incentives to obey — to get evidence and report it.

So best to use most severe punishment possible if an agent disobeys: not giving her the good.

Also not hard to see that wlog, principal does not give agent any information about history when asking for evidence.

Otherwise, we can improve agent incentives by pooling.

New property:

No free lunch: Agent is never given the good without first being asked for evidence.

If mechanism violates this on a positive measure set of histories, principal can strictly improve by changing to mechanism satisfying this.

$D =$ set of (randomizations over) mechanisms satisfying no–free–lunch.

$P_i(v \mid d) =$ probability mechanism allocates good to i given profile of types v and mechanism d given obedience.

$e_i(d) =$ ex ante probability agent i is asked for evidence in mechanism d given obedience.

The Maximization Problem

$$\max_{d \in D} \mathrm{E}_v \left[\sum_i P_i(v \mid d) v_i \right]$$

subject to

$$\textbf{(IC)} \quad \mathrm{E}_v P_i(v \mid d) - c_i e_i(d) \geq 0 \quad \forall i.$$

The Lagrangian:

$$\max_{d \in D} \left[\mathrm{E}_v \sum_i P_i(v_i \mid d) v_i + \sum_i \lambda_i \left(\mathrm{E}_v P_i(v \mid d) - c_i e_i(d) \right) \right]$$

$$\max_{d \in D} \left[\mathrm{E}_v \sum_i \left(P_i(v_i \mid d)(v_i + \lambda_i) - \lambda_i c_i e_i(d) \right) \right].$$

Digression: Weitzman's "Pandora's box" problem.

N boxes. "Prize" in box i has distribution $\hat{F}_i$. Cost of opening box i is $\hat{c}_i$.

Searcher cannot take a box without opening it.

Weitzman characterizes optimal search procedures.

$$\max_{d \in D} \left[\mathrm{E}_v \sum_i \left(P_i(v_i \mid d)(v_i + \lambda_i) - \lambda_i c_i e_i(d) \right) \right].$$

For fixed λ_i's, this is Weitzman's Pandora box problem.

agents $\rightarrow$ boxes

"ask i for evidence" $\rightarrow$ "open box i." e_i = probability of opening box i, $\lambda_i c_i$ = cost of opening box i.

"give good to i" $\rightarrow$ "take box i." $P_i(v)$ = probability of taking box i, $v_i + \lambda_i$ = prize in box i.

Then this is objective function in Weitzman.

What about set of options?

Claim: Set of search procedures in Weitzman essentially same as incentive compatible mechanisms satisfying no–free–lunch.

- Agents can refuse to get evidence, but boxes can't refuse to be opened. Incentive constraints make this irrelevant.
- Weitzman assumes box cannot be taken without first being opened; corresponds to our "no free lunch."

Solution of Weitzman problem for fixed λ_i's:

For each box i, compute an index, denoted $\hat{v}_i + \lambda_i$, defined by

$$\hat{v}_i + \lambda_i = \mathrm{E}_{v_i+\lambda_i} \max\{v_i + \lambda_i, \hat{v}_i + \lambda_i\} - \lambda_i c_i.$$

Procedure:

1. Open a box with the highest index. In case of ties, any randomization over these is optimal.
2. If prize in that box exceeds the highest index of the other boxes, stop and take that box. (Ties here will be measure zero.)
3. Otherwise, open a box with the highest index of remaining boxes, where any randomization over ties is optimal. Etc.

Translating to our problem:

1. Partition agents into tiers with highest tier consisting of those with highest index, second tier with second–highest, etc.
2. Randomly choose agent in first tier to ask for evidence. If virtual value $v_i + \lambda_i$ is above common index for tier 1, stop and give to i.
3. Otherwise, pick another agent in first tier and continue.
4. If no agent in first tier has virtual value above common index, the comparison is now to the next highest index, the common index for tier 2.
5. If some first tier agent has virtual value above this, give her good. Otherwise, ask tier 2 agents, etc.

This is mechanism described earlier *except* randomizations are arbitrary so far and can be history–dependent.

Unlike Weitzman, the randomization over tied indexes is critical since it matters for incentive compatibility.

We have to identify λ_i's and randomizations satisfying incentive compatibility and

$$\lambda_i[\mathrm{E}_v P_i(v) - c_i e_i] = 0, \quad \forall i.$$

Turns out that we can reduce problem to solving for λ_i's and e_i's subject to a feasibility constraint on e_i's.

Feasibility constraint ensures that there are randomizations such that these are the e_i's.

Show that history dependent randomizations do not enlarge set of feasible e_i's so we don't need to consider them.

Jump

Tempting intuition: For "generic" c_i's and F_i's, indices never tie, so randomization "almost always" irrelevant.

Recall that index is $\hat{v}_i + \lambda_i$ where $\hat{v}_i$ is a certain function of λ_i.

Key: λ_i is endogenous. Ties and therefore tiers with more than one agent are *not* nongeneric.

"Similar enough" agents must be in the same tier, so ties not "measure zero."

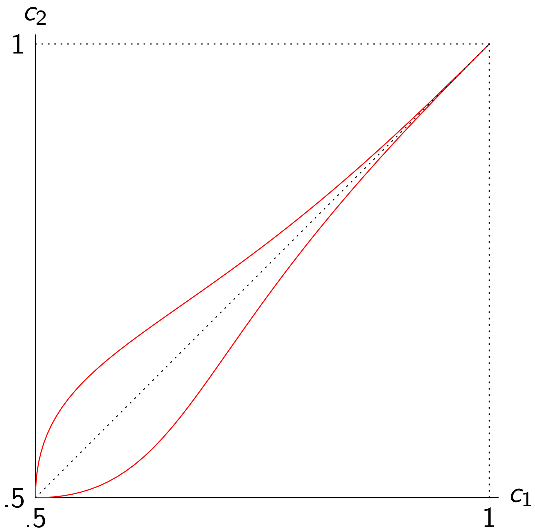
c_2
1
.5
.5
1
c_1

Optimal Stochastic Order of Agents

What determines distribution over order in which agents are asked for evidence?

Definition

Say that i is *stronger* than j if $c_i \leq c_j$ and F_i FOSD F_j.

One or both of these comparisons can hold with equality.

Theorem

If i is stronger than j, then $e_i \geq e_j$.

Implications:

1. i's index weakly larger than j's.
2. If they are in different tiers, then i is in a higher tier than j.
3. If the optimal order is deterministic, then i is asked for evidence before j.

Further implications:

- If $c_i = c_j$ and $F_i = F_j$, then $e_i = e_j$.

So symmetric agents must be treated symmetrically at the optimum.

▸ Jump

Illustration. $v_i \sim U[0,1]$. Drawn for $c_2 = .7$.

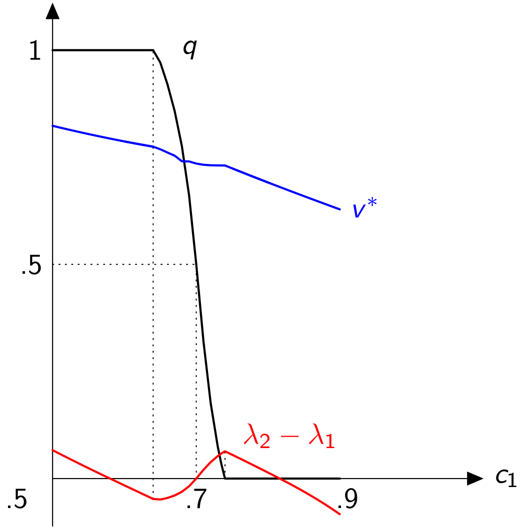

Other Interesting Papers

- Ali, Haghpanah, Lin, and Siegel, "How to Sell Hard Information," *QJE*, 2022.
- Dasgupta, Krasikov, and Lamba, "Hard Information Design," 2022.
- Gratton, Holden, and Kolotilin, "When to Drop a Bombshell," *RES*, 2018.
- Koessler and Skreta, "Selling with Evidence," *TE*, 2019.
- Perez–Richet and Skreta, "Test Design under Falsification," *Econometrica*, 2022.
- Rappoport, "Evidence and Skepticism in Verifiable Disclosure Games," 2022.
- Zhou, "Optimal Disclosure Windows," 2023.

Conclusion

Evidence can be fruitfully introduced in a variety of economic models.

Obtained insights in games and mechanism–design problems:

- What evidence environments result in “good” outcomes.
- What evidence will be provided in equilibrium.
- How this affects/distorts economic decisions.
- When are the solutions to mechanism-design problems robust and do not require commitment.
- How verification costs affect the optimal mechanism.

Some open questions:

- Other types of distortions stemming from the ability to manipulate evidence.
- Using the no–value–to–commitment to characterize optimal mechanisms with other evidence structures.
- Connection between other evidence structures and other forms of costly verification?
- Interaction of evidence/costly verification with other tools such as transfers.