Contract Theory: A New Frontier for AGT

Part I: Classic Theory

Paul Dütting – London School of Economics

Inbal Talgam-Cohen – Technion

ACM EC’19 Tutorial

June 2019

Plan

- Part I (Inbal): Classic Theory
  - Model
  - Optimal Contracts
  - Key Results
- Break (5-10 mins)
- Part II (Paul): Modern Approaches
  - Robustness
  - Approximation
  - Computational Complexity

*We thank Tim Roughgarden for feedback on an early version and Gabriel Carroll for helpful conversations; any mistakes are our own

1. What is a Contract?

An Old Idea

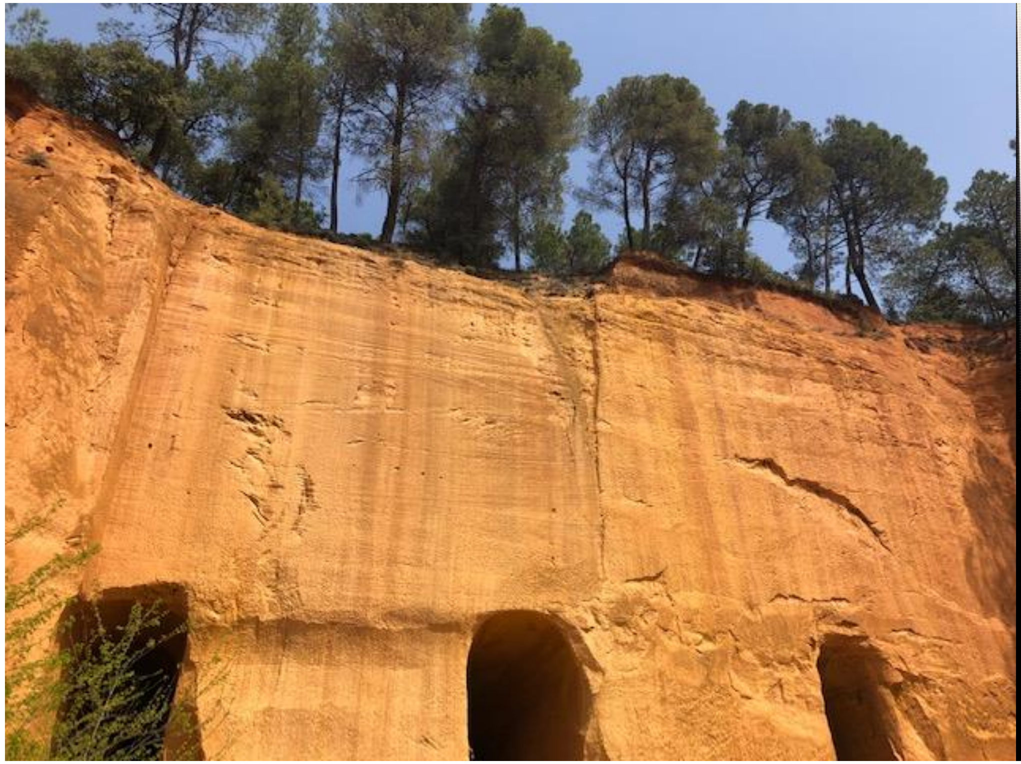

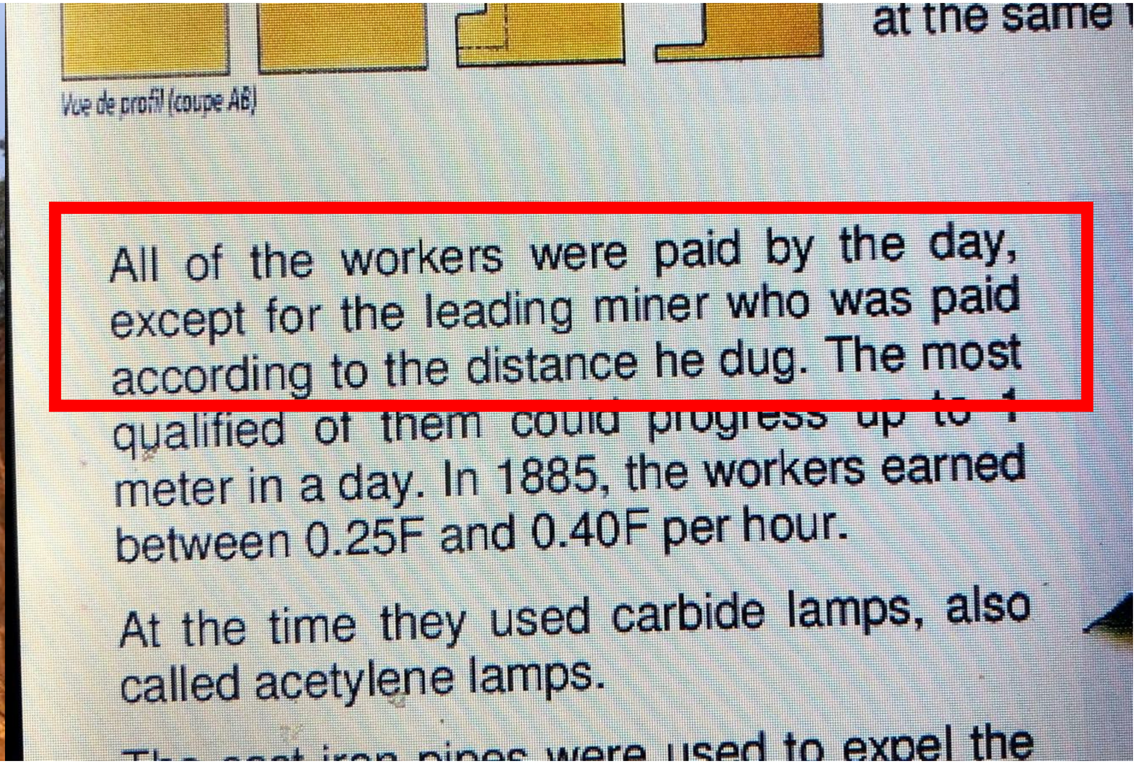

at the same

Vue de profil (coupe AB)

All of the workers were paid by the day, except for the leading miner who was paid according to the distance he dug. The most qualified of them could progress up to 1 meter in a day. In 1885, the workers earned between 0.25F and 0.40F per hour.

At the time they used carbide lamps, also called acetylene lamps.

Les Mines de Bruoux, dug circa 1885

Purpose of Contracts

- Contracts align interests to enable exploiting gains from cooperation
- “What are the common wages of labour, depends everywhere upon the contract usually made between those two parties, whose interests are not the same.” [Adam Smith 1776]

Classic Contract Theory

"Modern economies are held together by innumerable contracts"

[2016 Nobel Prize Announcement]

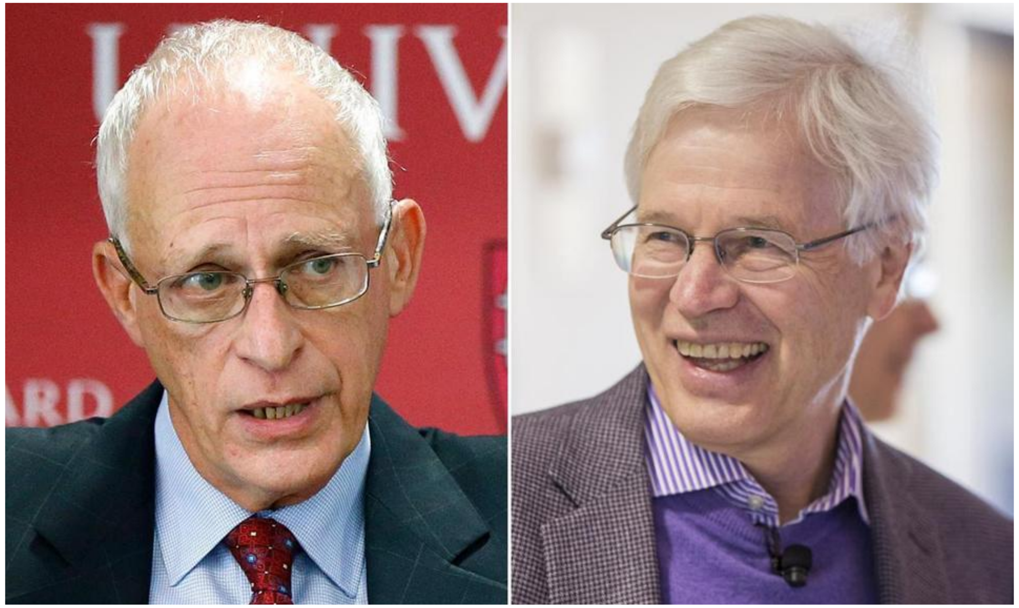

Laureates Oliver Hart and Bengt Holmström

Classic Applications

- Employment contracts
- Venture capital (VC) investment contracts
- Insurance contracts
- Freelance (e.g. book) contracts
- Government procurement contracts
- …

→ Contracts are indeed everywhere

New Applications

Classic applications are moving online and/or increasing in complexity

- Crowdsourcing platforms
- Platforms for hiring freelancers
- Online marketing and affiliation
- Complex supply chains
- Pay-for-performance medicare

→ Algorithmic approach becoming more relevant

Basic Contract Setting [Holmström'79]

- 2 players: principal and agent
- Familiar ingredients: private information and incentives
- Let's see an example…

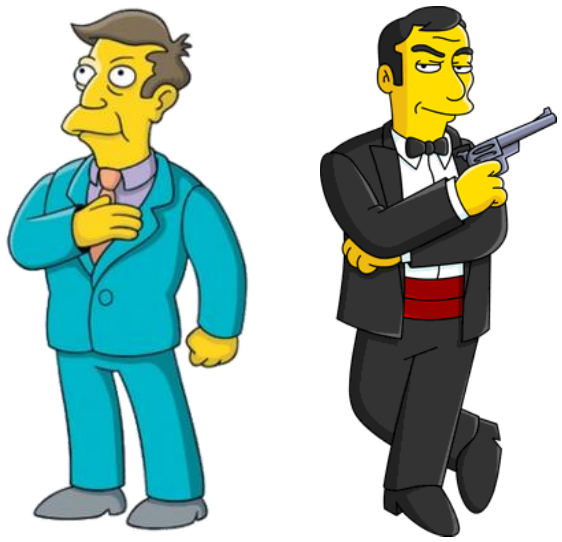

Example

- Website owner (principal) hires marketing agent to attract visitors

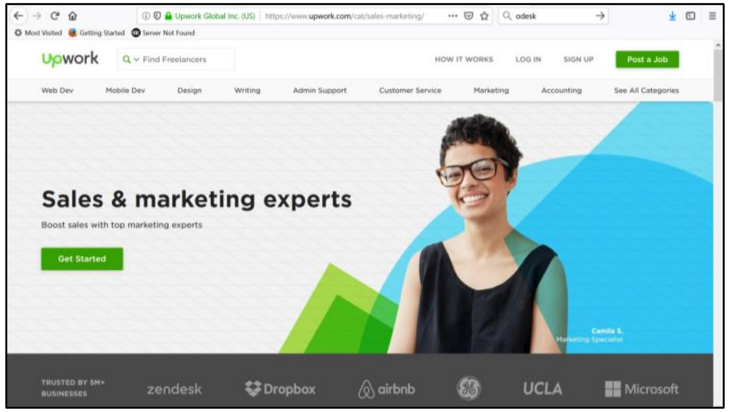

- <u>Two defining features</u>:
  1. Agent's actions are hidden - "moral hazard"
  2. Principal never charges (only pays) agent - "limited liability"

Moral Hazard

"Well then, says I, what's the use of you learning to do right when it's troublesome to do right and ain't no trouble to do wrong, and the wages is just the same?"

Mark Twain, *Adventures of Huckleberry Finn*

Limited Liability

Typical example: an entrepreneur and a VC

- The entrepreneur builds the company
- The VC diversifies the risks and has deep pockets

Timing

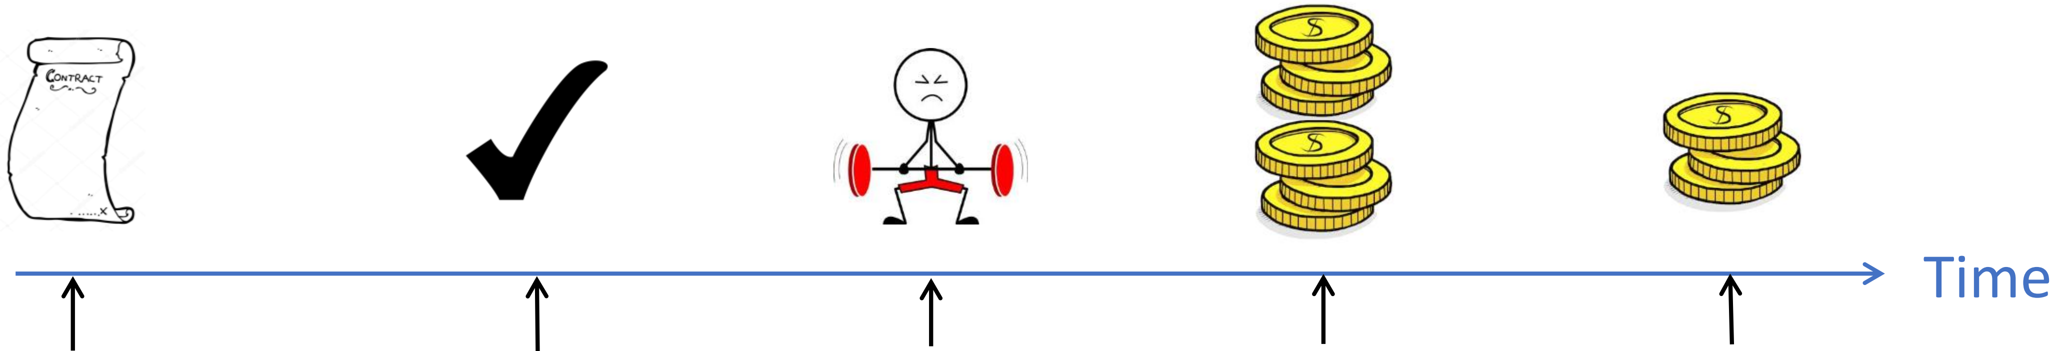

Principal offers agent a contract (parties have symmetric info)

Agent accepts (or refuses)

Agent takes costly, hidden action

Action's outcome rewards the principal

Principal pays agent according to contract

2. Connection to AGT

Relation to Other Incentive Problems [Salanie]

	Uninformed player has the initiative	Informed player has the initiative
Private information is hidden type	Mechanism design (screening)	Signaling (persuasion)
Private information is hidden action	Contract design	-

New Frontier

- Economics and computation – lively interaction over past 2 decades
- Especially true for mechanism design and signaling

Can we recreate the success stories of AGT in the context of contracts?

- Are insights from CS useful for contracts? Is contract theory useful for AGT applications? In Part II: A preliminary YES to both

Already Building Momentum

- Pioneering works:
  - Combinatorial agency [Babaioff Feldman and Nisan’12,…]
  - Contract complexity [Babaioff and Winter’14,…]
  - Incentivizing exploration [Frazier Kempe Kleinberg and Kleinberg’14]
  - Robustness [Carroll’15,…]
  - Adaptive design [Ho Slivkins and Vaughan’16,...]
- Recent works:
  - Delegated search [Kleinberg and Kleinberg’18,…]
  - Information acquisition [Azar and Micali’18,…]
  - Succinct models [Dütting Roughgarden and T.-C.’19b,…]
- EC’19 papers:
  - [Kleinberg and Raghavan’19, Lavi and Shamash’19, Dütting Roughgarden and T.-C.’19a]

The Algorithmic Lens

- Offers a language to discuss complexity
- Has popularized the use of approximation guarantees when optimal solutions are inappropriate
- Puts forth alternatives to average-case / Bayesian analysis that emphasize robust solutions to economic design problems

More on this in Part II

But first, let's cover the basics

3. Formal Model

Contract Setting

- Parameters n, m
- Agent has actions $a_1, \dots, a_n$
  - with costs $0 = c_1 \leq \cdots \leq c_n$ (can always choose action with 0 cost)
- Principal has rewards $0 \leq r_1 \leq \cdots \leq r_m$

- Action a_i induces distribution F_i over rewards (“technology”)
  - with expectation R_i
  - Assumption: $R_1 \leq \cdots \leq R_n$

- Contract = vector of transfers $\vec{t} = (t_1, \dots, t_m) \geq 0$

Recall two defining features

Example

Contract:	$t_1 = 0$	$t_2 = 1$	$t_3 = 2$	$t_4 = 5$

	No visitor $r_1 = 0$	General visitor $r_2 = 3$	Targeted visitor $r_3 = 7$	Both visitors $r_4 = 10$	
Low effort $c_1 = 0$	0.72	0.18	0.08	0.02	$R_1 = 1.3$
Medium effort $c_2 = 1$	0.12	0.48	0.08	0.32	$R_2 = 5.2$
High effort $c_3 = 2$	0	0.4	0	0.6	$R_3 = 7.2$

Expected Utilities

Contract setting:
- n actions $\{a_i\}$, costs $\{c_i\}$
- m rewards $\{r_j\}$
- $n \times m$ matrix F of distributions with expectations $\{R_i\}$

Fix action a_i.

Agent
- $\mathbb{E}$[utility] = expected transfer $\sum_{j\in[m]} F_{i,j}t_j$ minus cost c_i

Payoff $\neq$ payment/transfer

Principal
- $\mathbb{E}$[payoff] = expected reward R_i minus expected transfer $\sum_j F_{i,j}t_j$

Utilities sum up to $R_i - c_i$, action a_i's expected welfare

Example: Agent's Perspective

	Contract:	$t_1 = 0$	$t_2 = 1$	$t_3 = 2$	$t_4 = 5$
		No visitor $r_1 = 0$	General visitor $r_2 = 3$	Targeted visitor $r_3 = 7$	Both visitors $r_4 = 10$
0.44	Low effort $c_1 = 0$	0.72	0.18	0.08	0.02
1.24	Medium effort $c_2 = 1$	0.12	0.48	0.08	0.32
1.4	High effort $c_3 = 2$	0	0.4	0	0.6

Expected transfers: (0.44, 2.24, 3.4) for (low, medium, high)

Example: Agent's Perspective

Contract:	$t_1 = 0$	$t_2 = 1$	$t_3 = 2$	$t_4 = 5$

	No visitor $r_1 = 0$	General visitor $r_2 = 3$	Targeted visitor $r_3 = 7$	Both visitors $r_4 = 10$
Low effort $c_1 = 0$	0.72	0.18	0.08	0.02
Medium effort $c_2 = 1$	0.12	0.48	0.08	0.32
High effort $c_3 = 2$	0	0.4	0	0.6

Expected transfers: (0.44, 2.24, 3.4) for (low, medium, high)

Example: Principal's Perspective

Contract:	$t_1 = 0$	$t_2 = 1$	$t_3 = 2$	$t_4 = 5$	
	No visitor $r_1 = 0$	General visitor $r_2 = 3$	Targeted visitor $r_3 = 7$	Both visitors $r_4 = 10$	
Low effort $c_1 = 0$	0.72	0.18	0.08	0.02	$R_1 = 1.3$
Medium effort $c_2 = 1$	0.12	0.48	0.08	0.32	$R_2 = 5.2$
High effort $c_3 = 2$	0	0.4	0	0.6	$R_3 = 7.2$

R_3 - expected transfer $= 7.2 - 3.4 = 3.8$

A Remark on Risk Averseness

- Recall 2nd defining feature: agent has limited liability ($\vec{t} \geq 0$) [Innes'90]

- Popular alternative to risk-averseness
  - Utility from transfer t_j is $u(t_j)$ where u strictly concave

- Both assumptions justify why the agent enters the contract
  - Rather than "buying the project" and being her own boss

A Remark on Tie-Breaking

- Standard assumption: If the agent is indifferent among actions, he chooses the one that maximizes the principal's expected payoff

4. Computing Optimal Contracts

Contract Design

Goal: Design contract that maximizes principal's payoff

Optimization s.t. incentive compatibility (IC) constraints:

- Maximize $\mathbb{E}$[payoff] from action a_i
- Subject to a_i maximizing $\mathbb{E}$[utility] for agent

Related Problems: Implementability of action a_i; min pay for action a_i

Can all be solved using LPs!

First-Best Benchmark

- First-best = solution ignoring IC constraints

- What principal could extract if actions weren't hidden
  - I.e., if could pick action and pay its cost

$$\text{First-best} = \max_i \{R_i - c_i\}$$

- OPT $\neq$ first-best due to IC constraints

Implementability Problem

Given: Contract setting; action a_i

Determine: Is a_i implementable (exists contract $\vec{t}$ for which a_i is IC)

LP duality gives a simple characterization!

Proposition: Action a_i is implementable (up to tie-breaking) $\Leftrightarrow$

no convex combination of the other actions has same distribution over rewards at lower cost

Implementability LP

a_i implementable $\Longleftrightarrow$ LP feasible

m variables $\{t_j\}$ (transfers); $n-1$ IC constraints

$$\text{minimize } 0$$

$$\text{s.t. } \underbrace{\sum_j F_{i,j} t_j - c_i}_{\text{Agent's expected utility from } a_i \text{ given contract } \vec{t}} \geq \sum_j F_{i',j} t_j - c_{i'} \quad \forall i' \neq i \quad \text{(IC)}$$

$$t_j \geq 0 \quad \text{(LL)}$$

Dual* for Action a_i

Primal infeasible $\Leftrightarrow$ $\exists$ feasible dual solution with objective > 0

$n-1$ variables $\{\lambda_{i'}\}$ (weights); m constraints

$$\text{maximize } c_i - \underbrace{\sum_{i' \neq i} \lambda_{i'} c_{i'}}_{\text{Combined cost}}$$

$$\text{s.t. } \underbrace{\sum_{i' \neq i} \lambda_{i'} F_{i',j}}_{\text{Convex combination of actions}} \leq F_{i,j} \quad \forall j \in [m]$$

$$\lambda_{i'} \geq 0; \quad \sum_{i' \neq i} \lambda_{i'} = 1$$

Min Pay Problem

- Find minimum total transfer of a contract implementing action a_i
- Same LP with updated objective:

$$\text{minimize} \sum_j F_{i,j} t_j$$

$$\text{s.t.} \sum_j F_{i,j} t_j - c_i \geq \sum_j F_{i',j} t_j - c_{i'} \quad \forall i' \neq i \quad \text{(IC)}$$

$$t_j \geq 0 \quad \text{(LL)}$$

Optimal Contract Problem

Key observation:

- Can compute optimal contract by solving n LPs, one per action

Run-time per LP:

- Polynomial in $n - 1$ (constraints), m (variables)

Corollary:

- ∃ optimal contract with $\leq n - 1$ nonzero transfers

Criticism of LP-Based Approach

"More normative than positive":

Part II

1. Requires perfect knowledge of distribution matrix F

Part II

2. What if polytime in n, m is too slow?
   - Recall example: m is exponential in number of visitor types to website

Now

3. The contract that comes out of the LP may seem arbitrary

5. Structure of Optimal Contracts

Optimal Contract for 2 Actions, 2 Rewards

Let $\pi > p$

	"Failure" reward r_1	"Success" reward $r_2 > r_1$
"Shirking" cost $= 0$	$1 - p$	p
"Working" cost $= c > 0$	$1 - \pi$	π

Optimal Contract for $n = m = 2$

	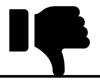	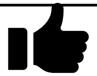
	$1 - p$	p
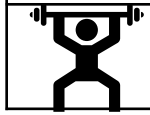	$1 - \pi$	π

Q: What does the optimal contract look like?

- Principal can always extract $R_1 = (1 - p)r_1 + pr_2$

→ Question interesting when optimal contract incentivizes work

- In this case:

first-best = $R_2 - c = (1 - \pi)r_1 + \pi r_2 - c$

Optimal Contract for $n = m = 2$

	$1-p$	p
	$1-\pi$	π

- Notation: Contract pays t_1 for failure, t_2 for success
- IC constraint for working is:

$$t_1(1-\pi) + t_2\pi - c \geq t_1(1-p) + t_2 p$$
$$\Leftrightarrow \quad (\pi - p)(t_2 - t_1) \geq c \quad (*)$$

- $(*)$ binds at the optimal contract

→ Optimal contract is: $t_1 = 0; \quad t_2 = \frac{c}{\pi - p}$

→ Principal extracts: $R_2 - c\frac{\pi}{\pi - p}$

Compare to first-best = $R_2 - c$

Optimal Contract for $n = m = 2$

	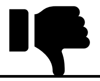	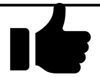
	$1 - p$	p
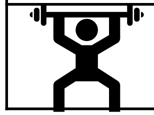	$1 - \pi$	π

Q: Structural properties of the optimal contract $t_1 = 0;\ t_2 = \frac{c}{\pi - p}$?

- Monotonicity property = transfer increases w/ reward
  - Generalizes to any n as long as $m = 2$
- As π, p draw closer, harder to distinguish work from shirk, so t_2 grows

Optimal Contract for 2 Actions, m Rewards

- $n = 2, m > 2$

	 Very Poor	 Poor	 Average	 Good	 Excellent
	$F_{1,1}$	$F_{1,2}$	$F_{1,3}$	$F_{1,4}$	$F_{1,5}$
	$F_{2,1}$	$F_{2,2}$	$F_{2,3}$	$F_{2,4}$	$F_{2,5}$

- Recall: There's an optimal contract with $n - 1 = 1$ nonzero transfers

Optimal Contract for $n = 2, m > 2$

Q: Which reward r_j gets the nonzero transfer t_j in optimal contract?

- Binding IC constraint for working is $t_j F_{2,j} - c = t_j F_{1,j}$

→ Optimal contract is $t_j = \frac{c}{F_{2,j} - F_{1,j}}$; principal extracts $R_2 - c\frac{F_{2,j}}{F_{2,j} - F_{1,j}}$

Optimal Contract for $n = 2, m > 2$

- Principal extracts $R_2 - c\frac{1}{\left(1-\frac{F_{1,j}}{F_{2,j}}\right)}$

→ To maximize over all j, choose j^* that minimizes $\frac{F_{1,j}}{F_{2,j}}$

- $\frac{F_{1,j}}{F_{2,j}}$ is called the likelihood ratio of actions a_1, a_2
  - Numerator (denominator) is likelihood of shirk (work) given reward r_j

Takeaway: Optimal contract pays for reward with min likelihood ratio

Optimal Contract for $n = 2, m > 2$

Recap

Q: Which reward r_j gets the nonzero transfer t_j in optimal contract?

A: Pay for r_j with min likelihood ratio $\frac{F_{1,j}}{F_{2,j}}$

Statistical inference intuition (holds for general n):

Principal is inferring agent's action from the reward

→ Pays more for rewards from which can infer agent is working

An Extreme Example

- Assume reward r_{j^*} has nonzero probability ϵ only if agent works
  - I.e. if r_{j^*} occurs, "gives away" agent's action
- Optimal contract has single nonzero transfer $t_{j^*} = \frac{c}{\epsilon}$
- The good: Principal extracts first-best = $R_2 - c$
- The bad: Contract non-monotone
  - (Recall: monotone = transfer increases with reward)

Example with $n > 2$

Optimal contract incentivizes action a_3

Contract:	$t_1 = 0$	$t_2 = 0$	$t_3 \approx .15$	$t_4 \approx 3.9$	$t_5 \approx 2$	$t_6 = 0$
	$r_1 = 1$	$r_2 = 1.1$	$r_3 = 4.9$	$r_4 = 5$	$r_5 = 5.1$	$r_6 = 5.2$
$c_1 = 0$	3/8	3/8	2/8	0	0	0
$c_2 = 1$	0	3/8	3/8	2/8	0	0
$c_3 = 2$	0	0	3/8	3/8	2/8	0
$c_4 = 2.2$	0	0	0	3/8	3/8	2/8

Recap

Role of rewards in the model is two-fold:

1. Represent surplus to be shared
2. Signal to principal the agent’s action

- The optimal contract is shaped by (2)
- Can be mismatched with (1)

6. Results on Monotonicity and Informativeness

Regularity Conditions [Mirrlees'99]

Q: Natural conditions for optimal contract monotonicity?

For $n = 2$ actions, mth reward must have min likelihood ratio

Definition: A contract setting satisfies MLRP (monotone likelihood ratio property) if

$$\forall \text{ actions } a_i, a_{i'}, i < i': \quad \frac{F_{i,j}}{F_{i',j}} \text{ decreasing in } j$$

Intuition: The higher the reward, the more likely the higher-cost action

Note: MLRP implies FOSD ($F_{i'}$ first-order stochastically dominates F_i)

Regularity Conditions [Mirrlees'99]

- MLRP insufficient for monotonicity with $n > 2$ (recall example)
- Sufficient with "CDF Property" or if actions have increasing welfare
  - "CDFP really has no clear economic interpretation, and its validity is much more doubtful than that of MLRP" [Salanie'05]

	$r_1 = 1$	$r_1 = 1.1$	$r_1 = 4.9$	$r_1 = 5$	$r_1 = 5.1$	$r_1 = 5.2$
$c_1 = 0$	3/8	3/8	2/8	0	0	0
$c_2 = 1$	0	3/8	3/8	2/8	0	0
$c_3 = 2$	0	0	3/8	3/8	2/8	0
$c_4 = 2.2$	0	0	0	3/8	3/8	2/8

Informativeness [Grossman-Hart'83]

Fix n actions, $m \times m$ stochastic matrix Π

Consider 2 contract settings $(F, r), (F', r')$ s.t. $\forall$ action a_i:

- $R_i' = R_i$
- F_i' obtained from F_i as follows: Draw reward-index j' by drawing j from F_i, then drawing from jth column of Π

→ Settings have same expected rewards but F' is a coarsening of F

Proposition: Min pay for action a_i is higher in coarser setting

Informativeness [Holmstrom'79]

Suppose principal can observe additional signals indicating action, e.g., a report from agent's direct supervisor

Statistical model connection:

- Action = underlying parameter
- Reward + report = observed data

Given reward, does report give further info on action? If so – use it!

<u>Sufficient statistic theorem</u>: The principal should condition transfers on a sufficient statistic for all available signals

Recap

- Classic lit has made headway in making sense of optimal contracts
- E.g. through statistical inference connections

Limitations:

- Conditions like actions having increasing welfare are too strong
- “Coarsening” relation is a very partial order on contract settings

A Way Forward: Simple Contracts

- Linear contracts: Determined by parameter $\alpha \in [0,1]$
  - For reward r_j the principal pays the agent αr_j
  - Generalization to affine: $\alpha r_j + \alpha_0$

- Agent's expected utility from action a_i is $\alpha R_i - c_i$
- Principal's expected payoff is $(1-\alpha)R_i$

Notice: No dependence on details of distribution!

7. Model Extensions & Summary

Extensions

1. Continuum of actions: Studied in particular with 2 rewards [Mirrlees'99]
2. Continuum of rewards: Functional analysis [Page'87]
3. Multiple agents: Teamwork, free-riding [Holmstrom'82]
4. Multiple principals: Agent's success in a project benefits 2 principals [Bernheim-Whinston'86]
5. Multitasking: Actions can be substitutes or complements for agent [Holmstrom-Milgrom'91]
6. Adverse selection: Agents also have hidden types [E.g., Chiappori et al.'94]

Dynamics

1. Multiple time periods, agent takes action at each period

- In this model [Holmstrom and Milgrom'87] give first robustness explanation for real-life contracts taking a simple, often linear form

2. Renegotiation after action is taken

- May prevent implementing costly actions

Incomplete Contracts

Famous example from 1920s [Klein et al.'78]:

- Contract between GM and car-part manufacturer
- GM committed; manufacturer kept costs high ("held up" GM)

Problem caused by incomplete contract setting:

- Players can make specific investments
- Not all appear in contract due to transaction costs [Coase'37]

→Leads to underinvestment; here renegotiation can be socially useful

Recap of Part I

Contracts incentivize someone to do something for us although we get the rewards and they incur the cost

- A model with familiar components of private info, incentives that "fit together" in fundamentally different way than auctions

- Two defining features: (1) Hidden actions (2) Limited agent liability

Recap of Part I: Main Results

- Implementability, Min Pay and Optimal Contract all solvable with LPs in poly(n, m) runtime if distributions known

Optimal Contract	2 rewards	m rewards
2 actions	Monotonicity & Pay for reward with min likelihood ratio	Pay for reward with min likelihood ratio
n actions	Monotonicity	Strong assumptions needed

Resources

1. Jean-Jacques Laffont and David Martimort, “The Theory of Incentives: The Principal-Agent Model”, Princeton U. Press 2002
2. Patrick Bolton and Mathias Dewatripont, “Contract Theory”, MIT Press 2005
3. Bernard Salanie, “The Economics of Contracts: A Primer”, MIT Press 2005 (see in particular Chapter 5)

*See Appendices of [Dütting Roughgarden and T.-C.’19a] for more details on many of the basics covered in this tutorial

*For tutorial bibliography see tutorial website

Questions?

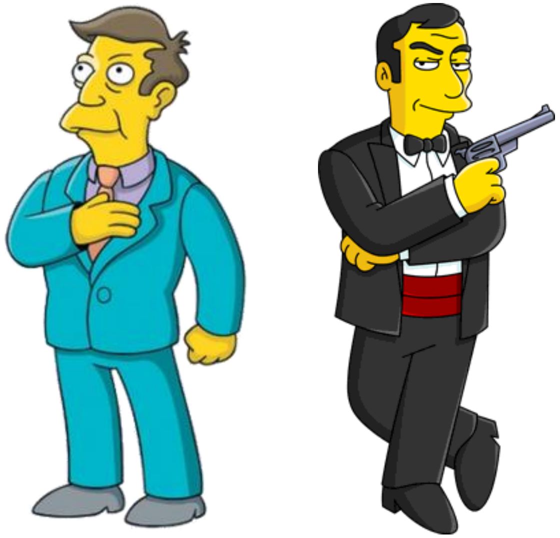

- After the break: Algorithmic aspects