

# Contract Theory: A New Frontier for AGT

## **Part II: Modern Approaches**

Paul Dütting – London School of Economics

Inbal Talgam-Cohen – Technion

ACM EC'19 Tutorial

June 2019

# Overview

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

# 1. Robustness

# Motivation

The classic principal-agent model [Holmström 1979, Grossmann and Hart 1983] suggests optimal contracts that

- Are rather **complex** and intransparent
- Exhibit **undesirable properties** (e.g., non-monotonicity)
- Do **not resemble** contracts used in **practice** (which tend to be simple, often linear)



Linear contract:  $t(r) = \alpha \cdot r, \alpha \in [0,1]$

# Milgrom-Holmström [1987]

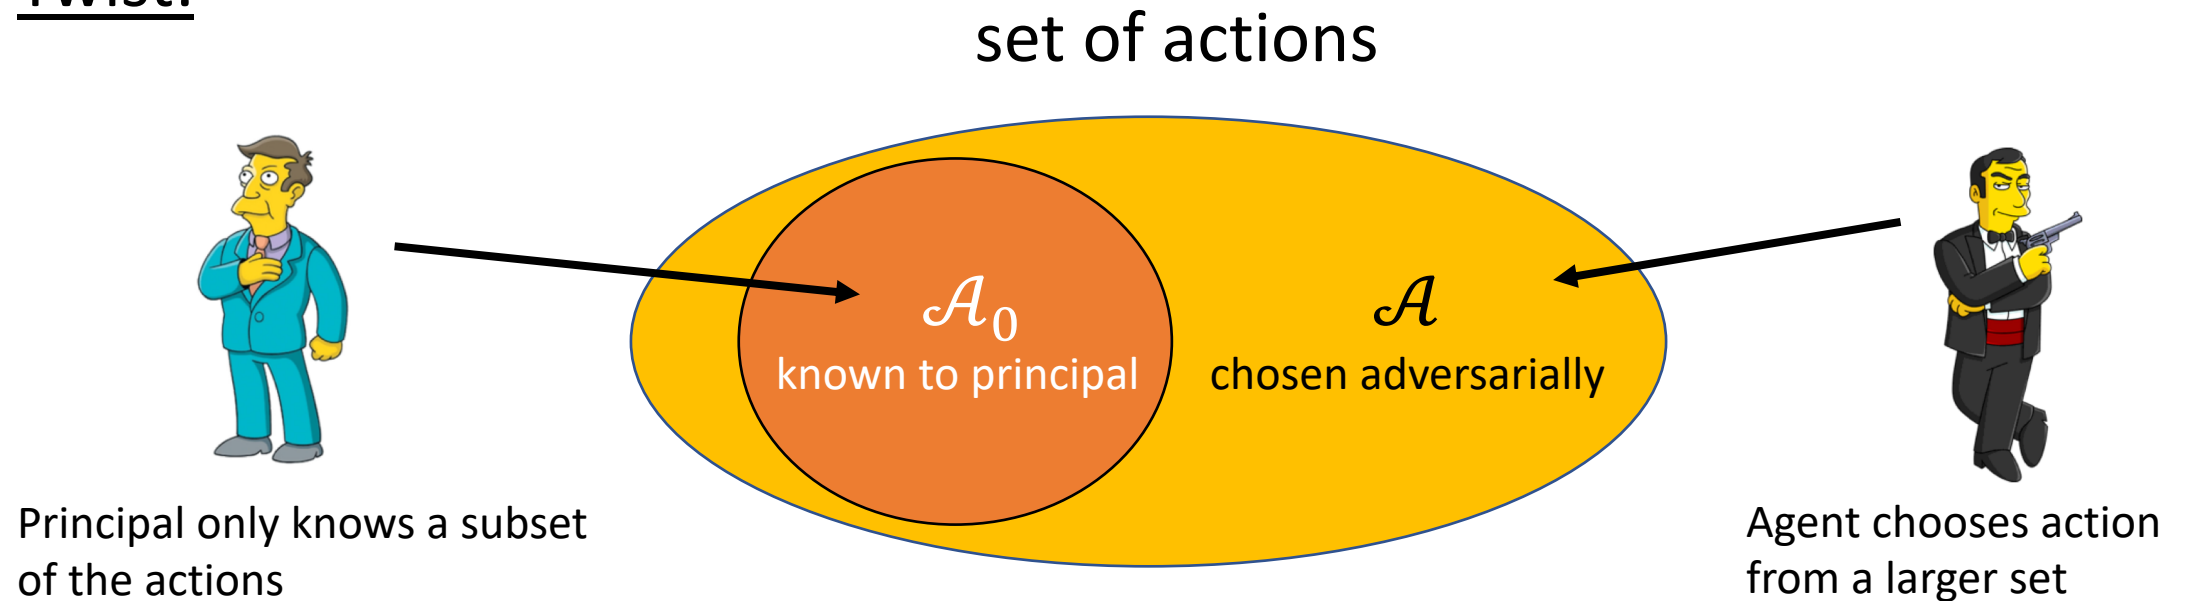
“It is probably the great **robustness** of **linear rules** based on aggregates that accounts for their popularity.

That point is not made as effectively as we would like by our model; we suspect that it **cannot be made** effectively **in any traditional Bayesian model.**”

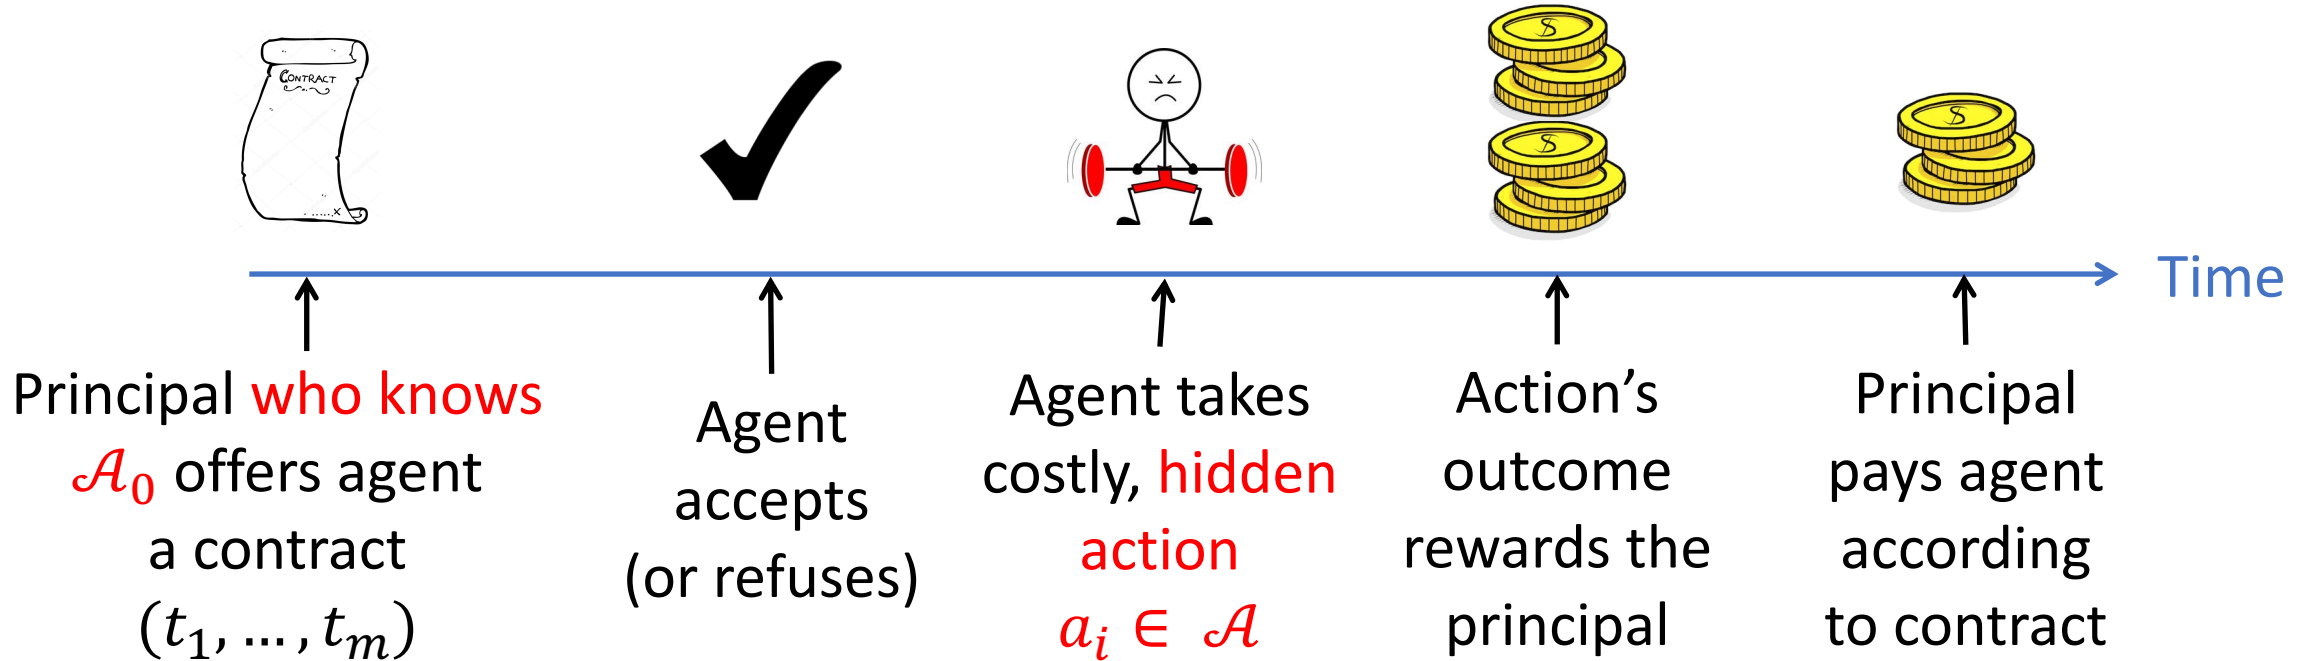
# Carroll's Model [2015]

Recall: Action  $a_i$  is specified by distribution  $F_{i,j}$  over rewards  $r_j$ , and a cost  $c_i$

Twist:



# Timing



# The Agent's Perspective

- The agent chooses action  $a^*$  from  $\mathcal{A}$  that maximizes expected payment minus cost

$$a^* \in \operatorname{argmax}_{a=(F,c) \in \mathcal{A}} (\mathbb{E}_{r \sim F} [t(r)] - c)$$

$\Rightarrow$  agent utility  $V_A(t|\mathcal{A})$

- Note: The agent can guarantee himself a certain expected utility by **only** maximizing over  $\mathcal{A}_0$

 “reserve agent utility”  $V_A(t|\mathcal{A}_0)$



# The Principal's Perspective

- Denote the set of actions that maximize the agent's utility for a given contract  $t$  and set of actions  $\mathcal{A}$  by

$$A^*(t|\mathcal{A}) = \operatorname{argmax}_{a=(F,c)\in\mathcal{A}} (\mathbb{E}_{x\sim F} [t(r)] - c)$$

- Then the principal solves the following **max-min problem**

$$\sup_t \inf_{\mathcal{A} \supseteq \mathcal{A}_0} \max_{a=(F,c)\in A^*(t|\mathcal{A})} \mathbb{E}_{r\sim F} [r - t(r)]$$

$V_P$

principal payoff  $V_P(t|\mathcal{A})$

# Reserve Principal Payoff?

- With a **linear contract**  $t(r) = \alpha \cdot r$ , for any action  $a = (F, c)$ :

$$\mathbb{E}_{r \sim F}[t(r)] = \alpha \cdot \mathbb{E}_{r \sim F}[r]$$

$$\mathbb{E}_{r \sim F}[r - t(r)] = (1 - \alpha) \cdot \mathbb{E}_{r \sim F}[r]$$



welfare pie

- So for every linear contract  $t(r) = \alpha \cdot r$  and incentivized action  $a = (F, c)$ :

$$V_P \geq \frac{1 - \alpha}{\alpha} \cdot \mathbb{E}_{r \sim F}[t(r)] \geq \frac{1 - \alpha}{\alpha} \cdot (\mathbb{E}_{r \sim F}[t(r)] - c)$$

$$\Rightarrow V_P \geq \frac{1 - \alpha}{\alpha} \cdot V_A(t | \mathcal{A}_0)$$

Maximizing the RHS gives max-  
min optimal contract



# Max-Min Robustness


## Theorem [Carroll'15]

For all partially specified principal agent-settings with rewards  $r_1, \dots, r_m$  and known action set  $\mathcal{A}_0$  there exists a **linear contract** that maximizes  $V_P$ .

# Key Steps in Proof


1. Argue that for any (not necessarily monotone) contract  $t$  there is an **affine contract**  $t'$  with the same or better worst-case guarantee (see next few slides)
2. Show that for any such affine contract  $t'$  there is an even better **linear contract**  $t''$  (see Carroll's paper for details)

# Why Affine is Enough




- Fix an arbitrary contract  $t$  (black dots)
- For any action  $a = (F, c)$  the agent may take, consider the point  $(\mathbb{E}_F[r], \mathbb{E}_F[t(r)])$
- This point lies in the **convex hull** of  $\{(r_j, t(r_j)) : 1 \leq j \leq m\}$  (gray area)

# Why Affine is Enough



- Moreover, the agent will only take actions that give him payoff at least  $V_A(t|\mathcal{A}_0)$  (dark gray area)
- Point  $Q$  is the point where expected payoff to the principal  $\mathbb{E}[r - t(r)]$  is smallest (bottom left of dark gray area)

# Why Affine is Enough



- Support line  $t'$  to the convex hull at  $Q$  is an **affine contract**, whose worst-case payoff to the principal is **no worse** than that of contract  $t$

# Discussion

- Obviously: Not the only way in which one can formalize model uncertainty
- Standard approach in **computer science** in cases where input is stochastic:
  - Assume **details** of the distributions are **unknown**
  - But **first moments** (or first few moments) **are known**

[E.g., Scarf'58, ..., Azar-Daskalakis-Micali-Weinberg'13, Bandi-Bertsimas'14]



# New Notion of Robustness

In an **EC'19 paper** (with Tim Roughgarden) we explore contract design with moment information:

- Fixed set of outcomes  $r_1, \dots, r_m$
- There are  $n$  actions with costs  $c_1, \dots, c_n$
- Details of the distributions  $F_1, \dots, F_n$  are **unknown**
- But their expected rewards  $R_i = \mathbb{E}_{r \sim F_i}[r]$  for  $i = 1, \dots, n$  are **known** (“compatible distributions”)

# New Notion of Robustness

**Theorem** [Dütting, Roughgarden, Talgam-Cohen'19a]

For every contract setting with **known expected rewards**, a **linear contract** maximizes the principal's expected payoff in the **worst-case** over compatible distributions.

So: Carroll's same conclusion, but under a very different hypothesis!

(Come to the EC talk!)

# Open Questions

- Is there a unification of Carroll's and our result?
- Study other models of uncertainty (e.g., **distributions** over outcomes are only **known approximately** [Bergemann-Schlag'11, Cai-Daskalakis '17, Dütting-Kesselheim'19])

# More Generally

A rapidly growing area in economics and computer science:

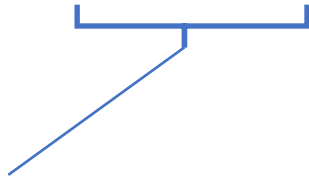
- Contracts [Carroll'15, Dütting-Roughgarden-Talgam-Cohen'19a]
- Revenue maximizing auctions [Bergemann-Schlag'11, Azar-Daskalakis-Micali-Weinberg'13, Bandi-Bertsimas'14, Carroll'17, Cai-Daskalakis'17, Carrasco-et-al.'18, Gravin-Lu'18, Bei-Gravin-Lu-Tang'19]
- Posted pricing and prophet inequalities [Dütting-Kesselheim'19]

## 2. Approximation

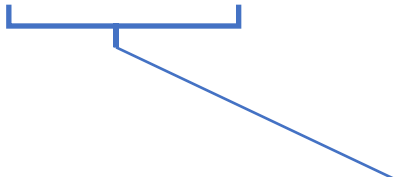
# A Powerful Tool from AGT

- Given a simple microeconomic mechanism, bound the **worst-case performance loss** relative to the optimal mechanism
- For a maximization problem: Find largest  $\beta \in [0,1]$  such that for all instances

$$ALG(I) \geq \beta \cdot OPT(I)$$



Performance of simple  
mechanism on instance



Optimal performance on  
instance


# Example: Linear Contracts

	$r_1 = 1$	$r_2 = 3$	
Action 1	$F_{1,1} = 1$	$F_{1,2} = 0$	$c_1 = 0$
Action 2	$F_{2,1} = 0$	$F_{2,2} = 1$	$c_2 = 4/3$


To find the **optimal contract**:

- The best way to incentivize action  $a_1$  is to pay  $t = (0,0)$  for an expected payoff of 1
- The best way to incentivize action  $a_2$  is to pay  $t = (0, 4/3)$  for an expected payoff of  $3 - 4/3 = 5/3$

$$\Rightarrow OPT = 5/3$$

$$t_2 = \frac{c_2}{F_{2,2} - F_{1,2}}$$


# Example: Linear Contracts




To find the **best linear contract**:

- Draw upper envelope with  $\alpha$  on  $x$ -axis and  $\alpha R - c$  on  $y$ -axis
- Each action corresponds to a line
- For every given  $\alpha$ , highest line corresponds to best (= chosen) action



# Example: Linear Contracts



- Here smallest  $\alpha$  at which action 1 and action 2 are implemented is  $\alpha = 0$  and  $\alpha = 2/3$

$$\Rightarrow ALG = 1 < 5/3$$

(Note: This shows that  $\beta$  can be at most  $3/5$ )

# Approximation Result

**Theorem (informal):** [Dütting, Roughgarden, Talgam-Cohen'19a]

**Linear** contracts achieve good approximation except in pathological settings with simultaneously:

- many **actions**;
- big spread among actions of expected **rewards**;
- big spread among actions of **costs**

# Example of a Pathological Setting

Let  $\epsilon \rightarrow 0$

$$(R_1, R_2, R_3, \dots) = (1, \frac{1}{\epsilon}, \frac{1}{\epsilon^2}, \dots)$$
$$(c_1, c_2, c_3, \dots) = (0, \frac{1}{\epsilon} - 2 + \epsilon, \frac{1}{\epsilon^2} - 3 + 2\epsilon, \dots)$$

# Formally

**Theorem** [Dütting, Roughgarden, Talgam-Cohen'19a]

$\rho$  = worst-case ratio of optimal contract and best linear contract

- with  $n$  actions,  $\rho = n$ ;
- with ratio  $R$  of highest to lowest  $R_i$ ,  $\rho = \Theta(\log R)$ ;
- with ratio  $C$  of highest to lowest  $c_i$ ,  $\rho = \Theta(\log C)$

- Upper bound w.r.t. to first best, lower bound w.r.t. optimal contract
- Lower bounds apply even under MLRP
- Bounds are tight, even for best monotone contract!

# Open Questions

- We only scratched the surface!
- The general question is: For which classes of contracts and under which assumptions on the setting can we get good (constant factor) approximations?
- Cf. "simple vs. optimal mechanisms" literature [Hartline and Roughgarden'09,...]



# 3. Computational Complexity

# Motivation

- If everything is given explicitly and there is only one agent then not interesting computationally
- If there is **more than one agent** or if **some part of the input is given implicitly** things become interesting:
  - E.g. an action could consist of several binary decisions
  - E.g. outcomes could be subsets of a ground set
  - E.g. ...

# Prior Work

- A paper which was way ahead of its time:  
Combinatorial Agency paper of Babaioff-Feldman-Nisan [2006, 2012]  
(and follow-up work)
- Studies a setting with multiple agents, in which each agent can take a binary action



# New Approach

In **ongoing work** (with Tim Roughgarden) we consider the following **succinct single-agent model**:

- There are  $\mu$  items,  $m = 2^\mu$  possible outcomes
- Given action  $a_i$ , each item  $k$  is included in the outcome independently wp  $F_{i,k}$
- The principal's reward is the sum of rewards  $r_k$  for each item  $k$  included in the outcome

# Example from Part I

Additive

	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

Product

E.g.  $\Pr[\text{general} \mid a_3] = 1, \Pr[\text{targeted} \mid a_3] = 0.6$

# Goal

Use **succinct** structure to exponentially **speed-up** finding the optimal contract in comparison to the naïve LP-based method

# Recall: Naïve LP-based Approach

- Based on solving  $n$  instances of the “MIN-PAY” problem
- Given action  $a_i$ , find optimal contract that implements  $a_i$

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

There are polynomial in  $m$ , exponential in  $\mu$  many variables, but only  $n$  constraints – **Ellipsoid** to the rescue?

# The Dual

$$\begin{aligned} & \text{maximize} \sum_{i' \neq i} \lambda_{i'} (c_i - c_{i'}) \\ & \text{s.t.} \sum_{i' \neq i} \lambda_{i'} - 1 \leq \boxed{\frac{\sum_{i' \neq i} \lambda_{i'} F_{i',j}}{F_{i,j}}} \quad \forall j \in [m] \end{aligned}$$

A **separation oracle** boils down to finding an item subset with **minimum likelihood** in the **combination distribution**  $\sum_{i' \neq i} \lambda_{i'} F_{i'}$  relative to  $F_i$

# Computational Hardness

- Solving the separation oracle exactly is **NP-hard**
- In fact computing the optimal expected payoff in succinct contract settings in time polynomial in  $\mu$  turns out to be **NP-hard**

# Approximate IC

A solution from AGT: Relax the IC constraints!

**Definition:** Given a contract  $t$ , action  $a_i$  is  $\delta$ -IC if

$$(1 + \delta) \sum_j F_{i,j} t_j - c_i \geq \sum_j F_{i',j} t_j - c_{i'} \quad \forall i' \neq i$$

In normalized settings, the agent loses  $\leq \delta$  by choosing a  $\delta$ -IC action

[By  $\delta$ -IC contract we mean a contract  $t$  and  $\delta$ -IC action  $a_i$  that pleases the principal]

# Theorem

Let  $OPT$  be the expected payoff of the optimal (IC) contract.

**Theorem** [Dütting, Roughgarden, Talgam-Cohen'19b]

There is an **Ellipsoid-based algorithm** that given a **succinct** contract setting with  $\mu$  items and a parameter  $\delta > 0$ , returns a  $\delta$ -IC contract with expected payoff  $\geq OPT$  in time **polynomial in  $\mu$  and  $1/\delta$** .

(Recall: Running time of naïve method is exponential in  $\mu$ )



# Ellipsoid-Based Algorithm

- Strengthened dual:

$$\begin{aligned} & \text{maximize } \sum_{i' \neq i} \lambda_{i'} (c_i - c_{i'}) \\ & \text{s.t. } (1 + \delta) \left( \sum_{i' \neq i} \lambda_{i'} - 1 \right) \leq \frac{\sum_{i' \neq i} \lambda_{i'} F_{i',j}}{F_{i,j}} \quad \forall j \in [m] \end{aligned}$$

- Run Ellipsoid calling an FPTAS for the separation oracle
- FPTAS runs in time polynomial in  $\mu$  and  $\frac{1}{\delta}$ , and exponential in  $n$

# Additional Results

In the paper [Dütting, Roughgarden, Talgam-Cohen'19b] we also show:

- Hardness of approximation for exactly IC contracts
- Constant factor  $\delta$ -IC contracts
- ....

(Watch out for the paper!)

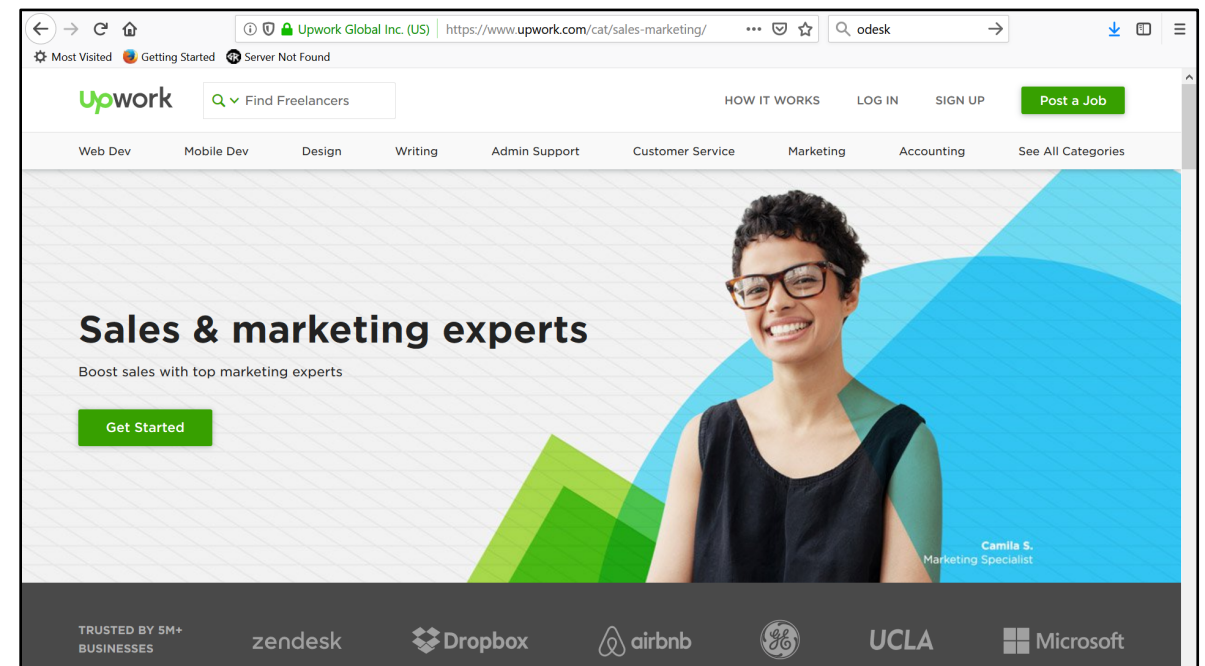
# Open Questions

- Many interesting computational questions
- Approximation probably even more natural than in the mechanism design world
- Mostly unexplored ...!

## 4. Concluding Remarks

# Important Applications

- Freelancing and crowdsourcing platforms
- Start-up funding platforms
- Blockchain and smart contracts
- Venture capital contracts
- Government procurement
- ...



# Growing Momentum

- Combinatorial agency [Babaioff-Feldman-Nisan'12,...]
- Contract complexity [Babaioff and Winter'14,...]
- Incentivizing exploration [Frazier-Kempe-Kleinberg-Kleinberg'14,...]
- Robustness [Carroll'15,...]
- Adaptive design [Ho-Slivkins-Vaughan'16,...]
- Delegated search [Kleinberg and Kleinberg'18,...]
- Information acquisition [Azar and Micali'18,...]
- Robustness [Dütting-Roughgarden-Talgam-Cohen'19a,...]
- Succinct models [Dütting-Roughgarden-Talgam-Cohen'19b,...]
- VCG contracts [Lavi-Shamash'19,...]
- Strategic classification [Kleinberg-Raghavan'19,...]

(At this year's EC)

# Many Open Problems

- There are **lost of interesting open questions** even in the most basic/classic models!
- The **algorithmic perspective** could be a **powerful tool** to complement the classic econ approach



Tutorial website:  
<http://personal.lse.ac.uk/act/index.htm>

Thanks! Questions?

# References

Gabriel Carroll. Robustness and Linear Contracts. *American Economic Review*, 105 (2), 2015, 536-563.

Paul Dütting, Tim Roughgarden, Inbal Talgam-Cohen. Simple versus Optimal Contracts. *Proc. 20th ACM Conference on Economics and Computation*, 2019, 369-387.

Paul Dütting, Tim Roughgarden, Inbal Talgam-Cohen. The Complexity of Contracts. *Working paper*, 2019.