Simple vs. Optimal Contracts

Paul Dütting – London School of Economics (Math)
Tim Roughgarden – Columbia (CS)
Inbal Talgam-Cohen – Technion (CS)

ACM EC @FCRC
Phoenix, June 2019

Contract Theory

- Contracts align interests to enable exploiting gains from cooperation
- "Modern economies are held together by innumerable contracts"

[2016 Nobel Prize Announcement]

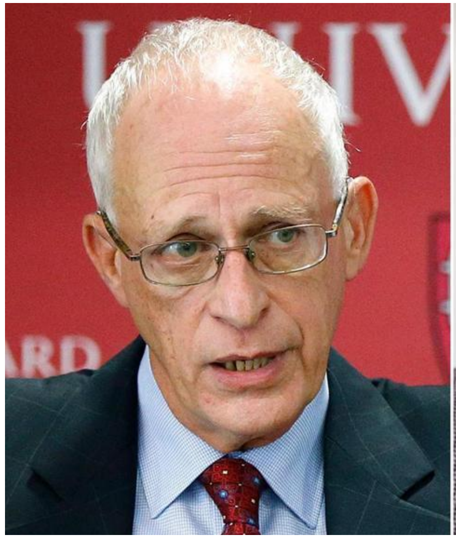

Oliver Hart   Bengt Holmström

Classic Applications

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

→ Contracts are indeed everywhere

Modern 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

→ Of interest to AGT; algorithmic approach becoming more relevant

The Algorithmic Lens

Research agenda: What can we learn about contract design through the algorithmic lens?

1. Robust alternatives to average-case / Bayesian analysis

2. Approximation guarantees when optimal solutions inappropriate

3. (Complexity issues – in different work)

For more information please see our EC’19 tutorial website

A Building Momentum

- Some 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]
- Some 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]

Basic Contract Setting: An Example

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

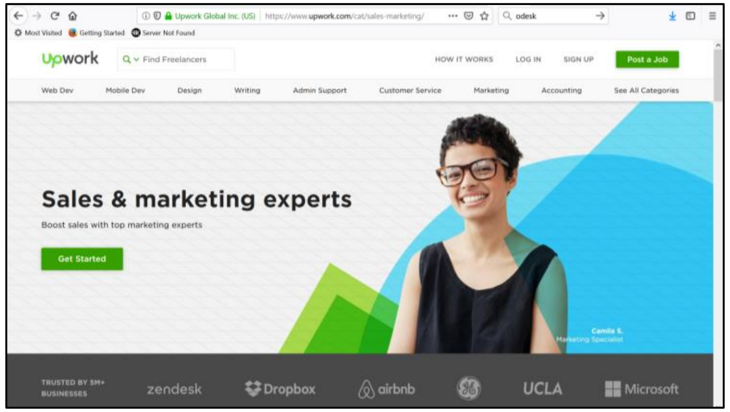

- Two defining features:
  1. Agent's actions are hidden - "moral hazard"
  2. Principal never charges (only pays) agent - "limited liability"

Relation to other Incentive Problems

- Mechanism design
  - Agents have hidden types
- Signaling (Bayesian persuasion)
  - Principal more informed

- Contracts [Holstrom'79]
  - No hidden types
  - Principal less informed

Our Results

- In the model of [Holmstrom'79]:

1. New robustness (max-min) justification for simple, linear contracts
    - "Standing on the shoulders" of [Carroll'15]

2. Approximation guarantees for linear contracts
    - Linear is far from optimal only in pathological cases
    - Approximation is tight even for monotone contracts

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	$R_3 = 7.2$

Contract Design Problem

An optimization problem with incentive compatibility (IC) constraints

Maximize principal's $\mathbb{E}$[payoff] from action a_i

subject to action a_i maximizing $\mathbb{E}$[utility] for agent

- $\mathbb{E}$[payoff] = expected reward R_i minus expected payment $\sum_j F_{i,j} t_j$
- $\mathbb{E}$[utility] = expected payment $\sum_j F_{i,j} t_j$ minus cost c_i

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
	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 $= 7.2 - 3.4 = 3.8$

LP-Based Solution

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

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

Expected transfer to agent for action a_i

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

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

Agent's expected utility from a_i given contract $\vec{t}$

- Caveats: (1) imperfect distribution knowledge (2) impractical contract

Result 1: Robust Optimality

Linear Contracts

- Determined by parameter $\alpha \in [0,1]$:
  - Given reward r_j, principal transfers αr_j to agent
  - 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 $(1 - \alpha) R_i$

No dependence on details of distribution!

- Really popular in practice

Robustness

- “It is probably the great robustness of linear rules […] that accounts for their popularity” [Milgrom-Holmström’87]

- Breakthrough formulation of [Carroll’15]: Linear contracts are optimal in the worst-case over unknown extra actions available to agent

- Alternative formulations?
- Standard CS formulation of uncertainty when input is stochastic: assume only first moments of the distribution are known [Scarf’58]

Example

	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	$R_3 = 7.2$

Example

	No visitor $r_1 = 0$	General visitor $r_2 = 3$	Targeted visitor $r_3 = 7$	Both visitors $r_4 = 10$	
Low effort $c_1 = 0$	?	?	?	?	$R_1 = 1.3$
Medium effort $c_2 = 1$	?	?	?	?	$R_2 = 5.2$
High effort $c_3 = 2$	?	?	?	?	$R_3 = 7.2$

New Robustness Result

Theorem:

- Given a contract setting with unknown distributions but known expectations,
- a linear contract is optimal in the worst-case over all compatible distributions

→ Same conclusion as [Carroll'15], under very different hypothesis!

Intuition: If you don't know enough to design a contract depending on anything but the expected rewards, optimize wrt what you know

Proof Overview: Max-Min Visualization

- Fix a contract setting with known expected rewards

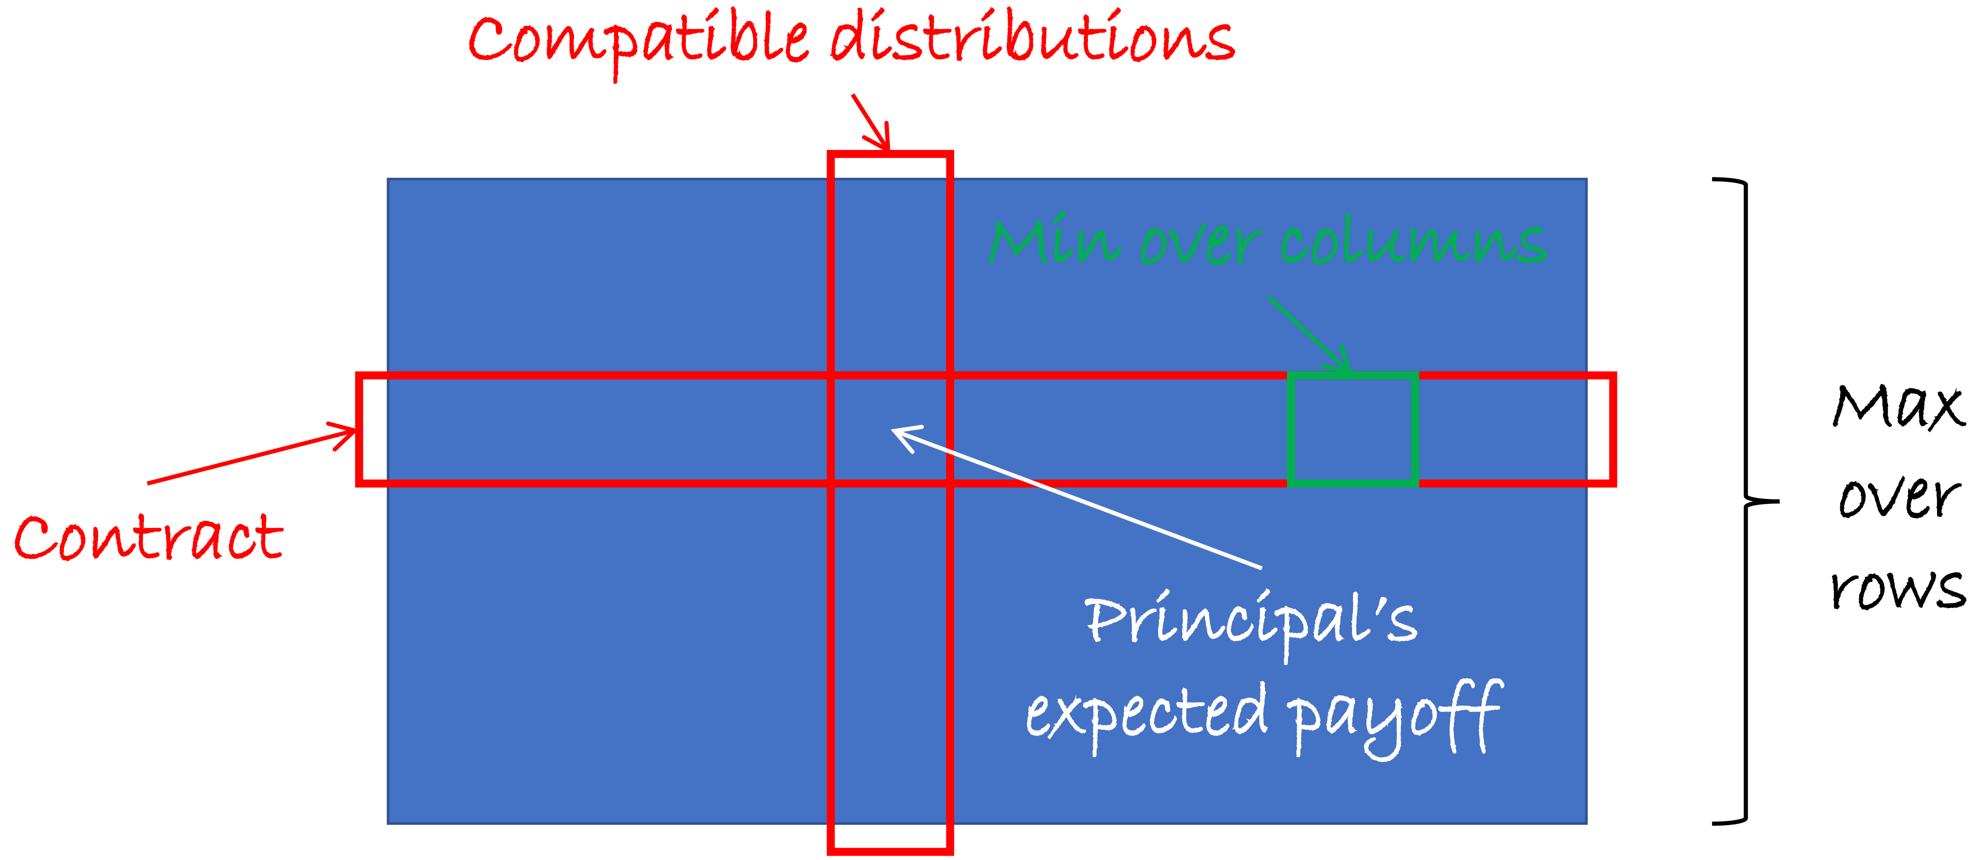

Proof Overview: Linear Contracts are Robust

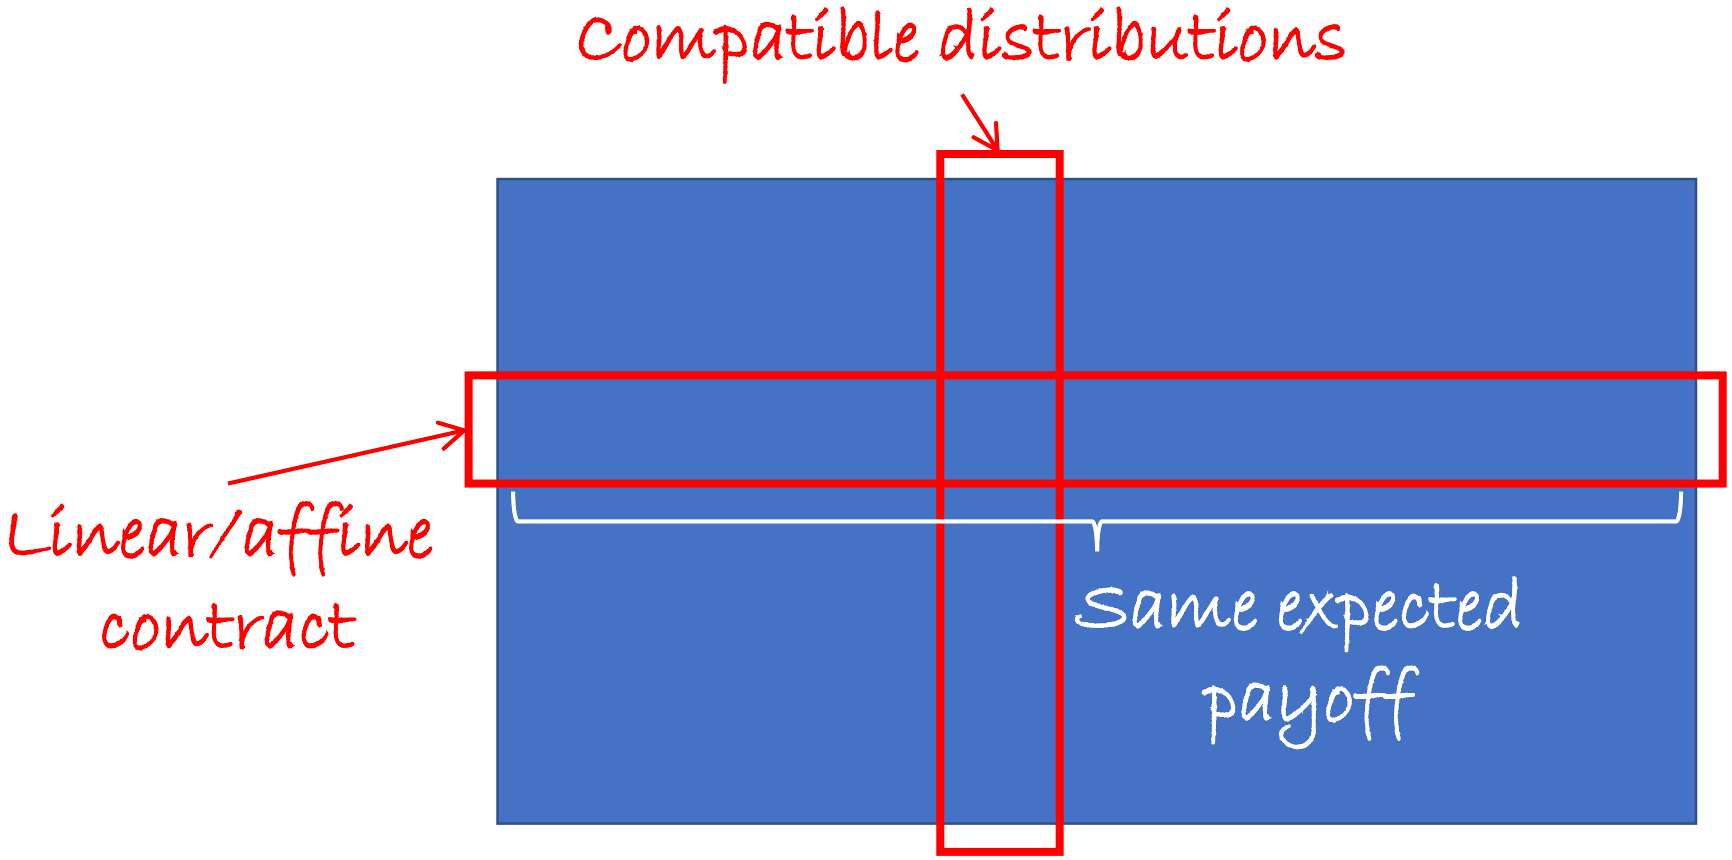

Proof Overview: Key Lemma

Lemma: For every contract t there exist compatible distributions and an affine contract with $\alpha_0 \geq 0$ and better expected payoff

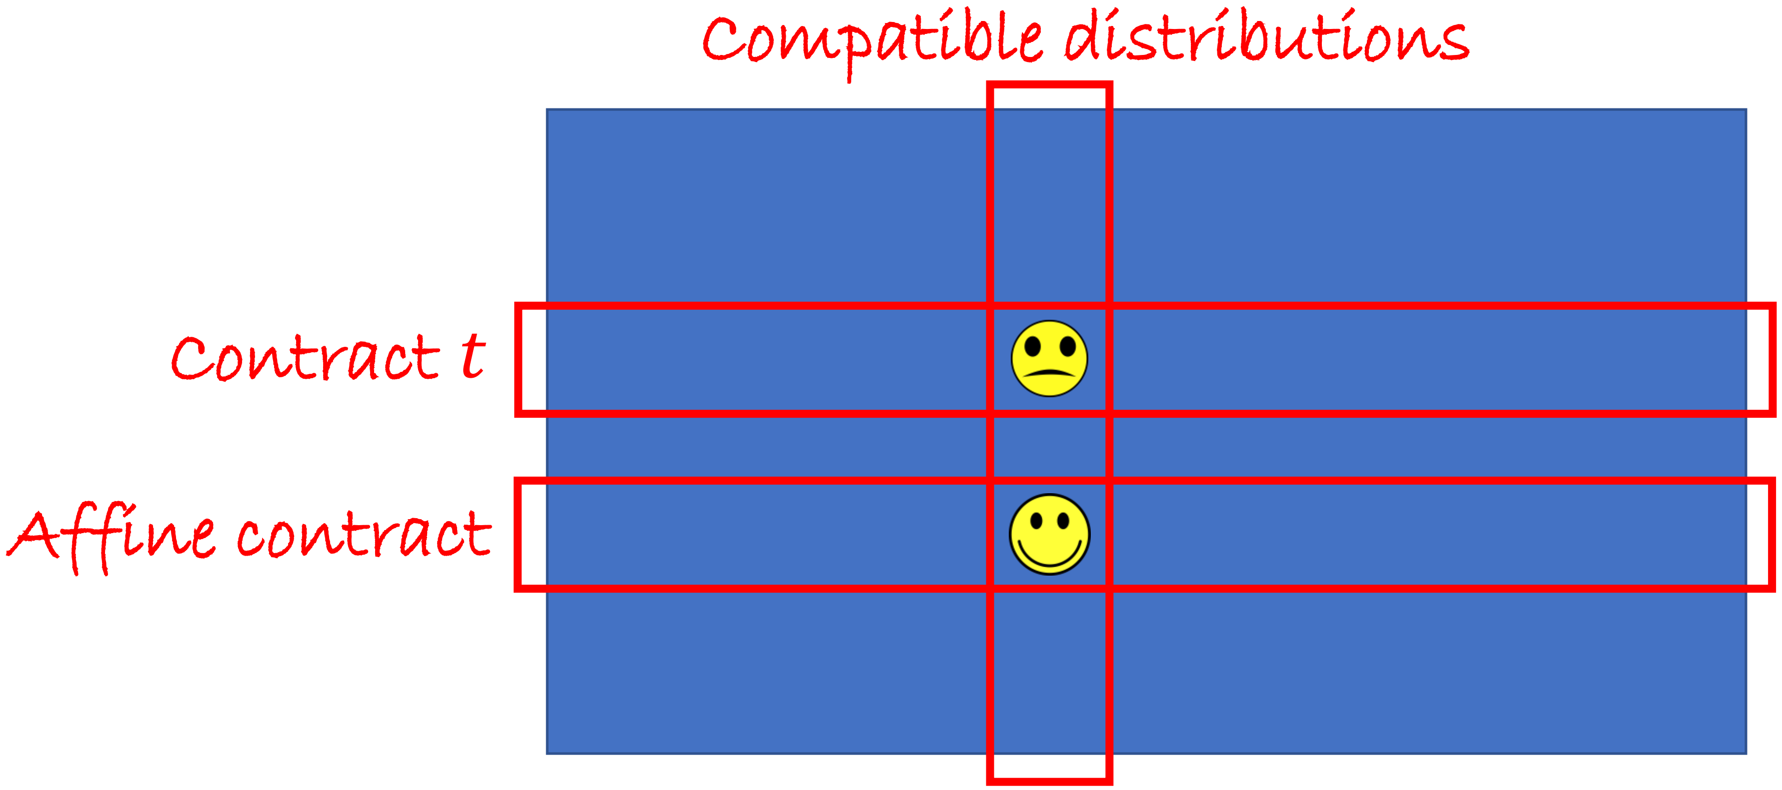

Key Lemma Suffices

→ For every contract t there exists an affine contract with $\alpha_0 \geq 0$ and better worst-case expected payoff

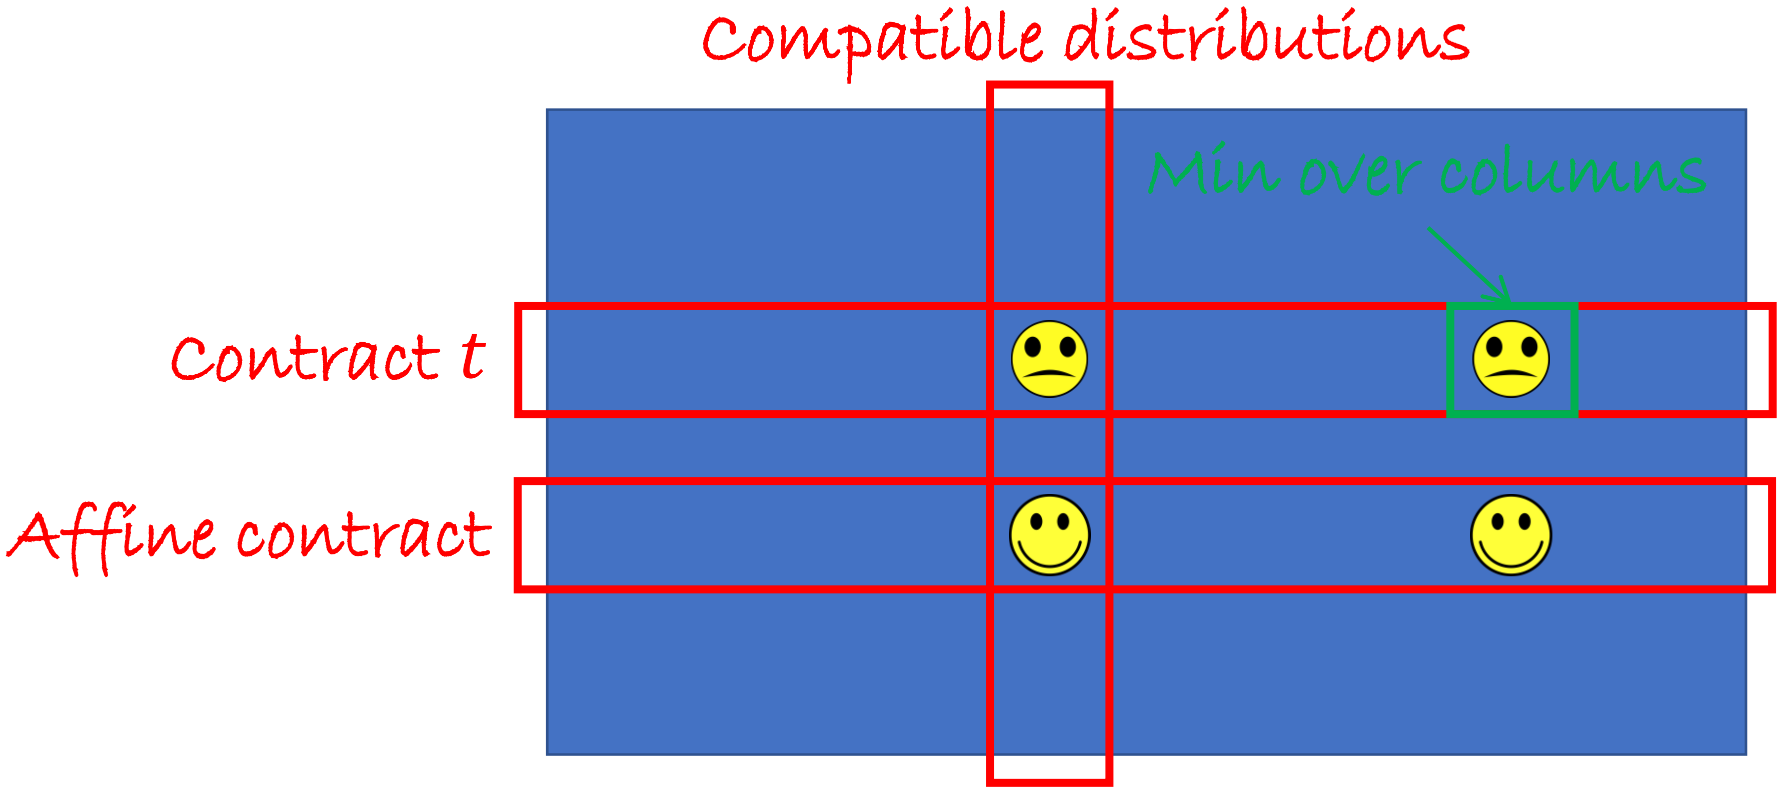

Key Lemma Suffices

In an affine contract, setting $\alpha_0 = 0$ increases expected payoff
→ Optimal linear contract has best worst-case expected payoff

QED

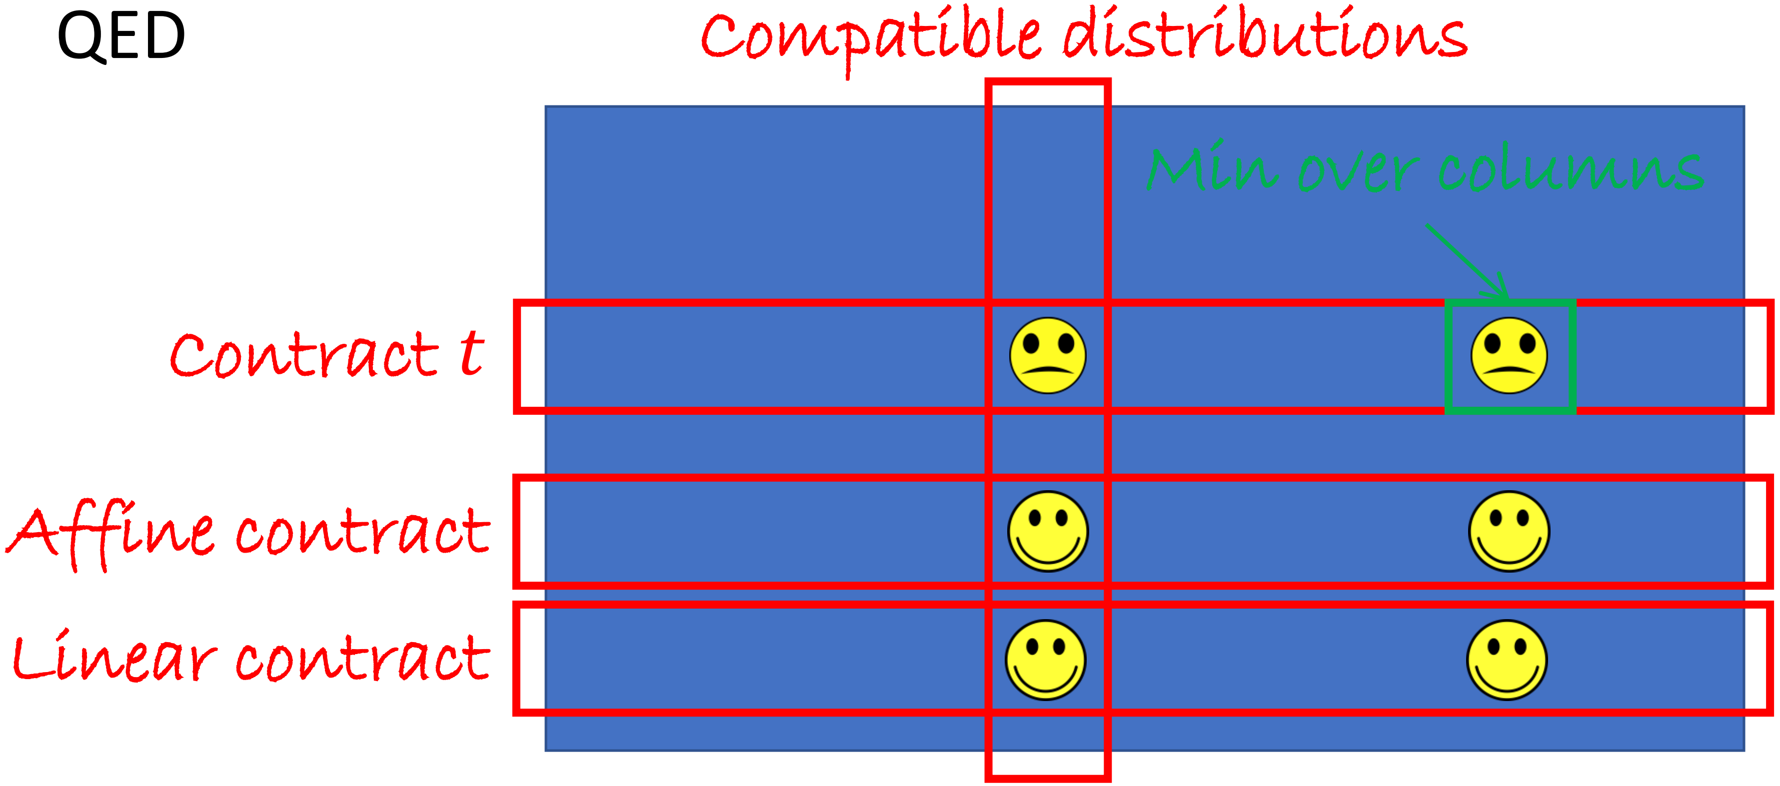

Result 2: Approximation

Approximation

What fraction of the optimal payoff is achievable by a simple contract?

- Result (informal): Linear contracts achieve constant approximation except in pathological settings with simultaneously:
  - many actions;
  - big spread of expected rewards;
  - big spread of costs

Example of Pathological Setting

- Let $\epsilon \to 0$

$$(R_1, R_2, R_3, \dots) = \left(1, \frac{1}{\epsilon}, \frac{1}{\epsilon^2}, \dots\right)$$

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

Formally

<u>Theorem</u>:

ρ = 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, $\rho = \Theta(\log C)$

Bounds are tight even for best monotone contract

Summary

- Contract theory as an interesting new frontier for AGT

- Algorithmic approach can provide new insights, such as:
  - Optimize the contract to available moment information
  - Expect linear contracts to perform well except in pathological cases

- Opportunities for new success stories