Contract Design Under Uncertainty

Tal Alon | PhD Student @ Technion, Israel Institute of Technology

Joint work with Paul Dütting, Yingkai Li, and Inbal Talgam-Cohen

October 3, 2023

MSRI

Modern Algorithms in Society

- Interactions with **self-interested** individuals
- In AGT, we take participants' **incentives** into account

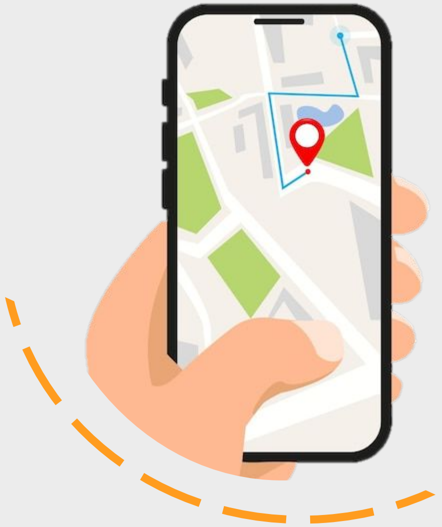

Social Media Marketing

- **Paid Ads** - in the social platform feed
- **Influencer Marketing** - a brand hires popular users

Paid Ads

- Algorithmic **auction** - advertisers bid → allocation and payments

Input $\xrightarrow{\text{hidden type}}$ **Algorithm** → Output

✔ Computational Complexity

✔ Achieving optimality

✔ **Incentivize truthfulness**

Influencer Marketing

- The algorithm determines a **contract** to incentivize effort

Input → **Algorithm** → Output

✔ Computational Complexity
✔ Achieving optimality
✔ **Incentivize effort**

The Principal-Agent Problem [GH83]

- **Moral hazard** - the agent's actions cannot be observed
- **Objective:** a contract maximizing expected **rewards minus payment** $\mathbb{E}[r - t]$

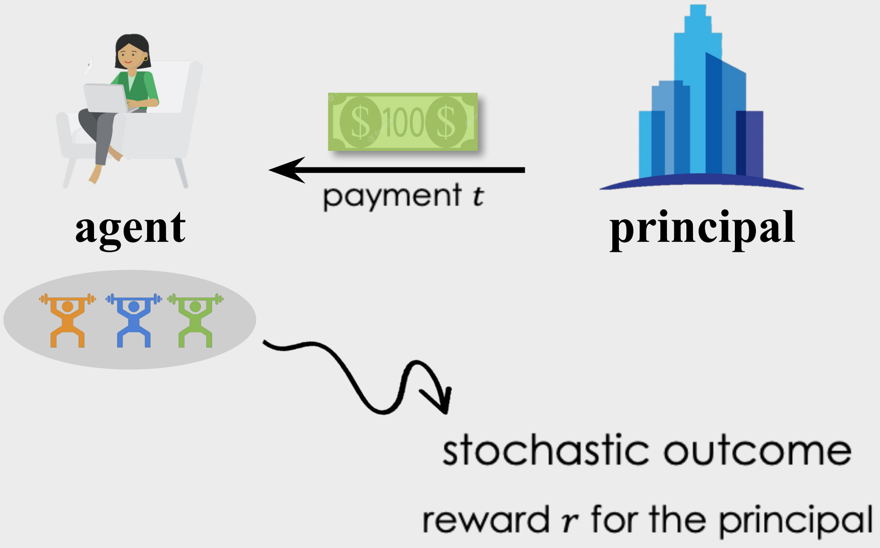

Research Agenda

Generalizations of the classic model:

- **Personalization** for participants from diverse population (with **types**)
- **Multilateral** contracts involving multiple principals or agents
- The need for **simple** contracts

Research Agenda

Generalizations of the classic model:

- **Personalization** for participants from diverse population (with **types**)
- **Multilateral** contracts involving multiple principals or agents
- The need for **simple** contracts

Research Agenda

Input $\xrightarrow{\text{hidden type}}$ **Algorithm** $\xrightarrow{\text{hidden action}}$ Output

- ✔ Computational Complexity
- ✔ Achieving optimality
- ✔ **Incentivize effort**
- ✔ **Incentivize truthfulness**

Outline

- **Single-parameter** model of types [ADT.C. EC'21]
  - Motivated by single-parameter auction design
- **Characterization** of the design space [ADT.C. EC'21]
- Counter-intuitive and **undesirable** properties of **optimal contracts** [ADLT.C. EC'23]
- **Linear contracts** (aka commission-based) are near-optimal [ADLT.C. EC'23]

Recent works on contracts with types. Myerson)1982(, Guruganesh et al.)2021), Castiglioni et al.)2021), Gottlieb and Moreira)2022(, Casto-Pire et al.)2022(.

The Model

- The classic principal-agent model by Grossman-Hart (1983)

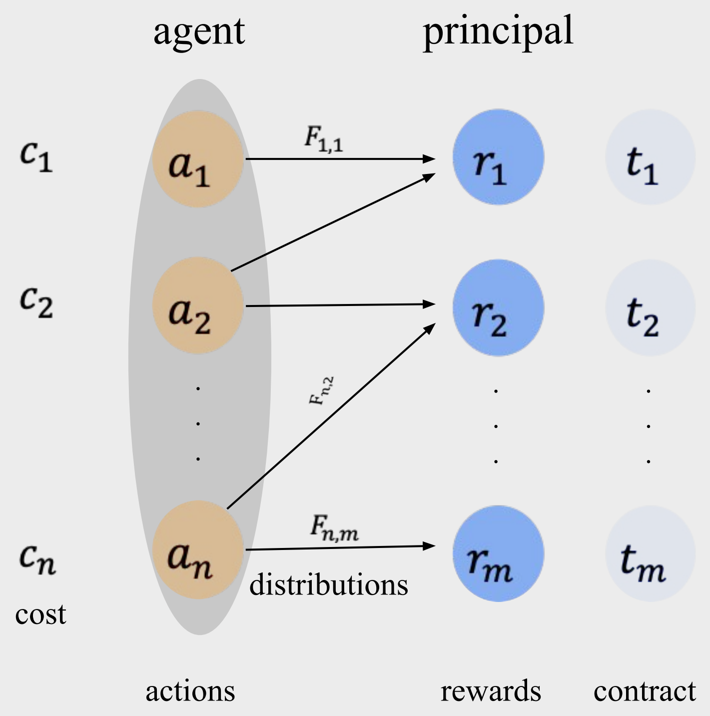

Notation. $T_i = \mathbb{E}_{j \sim F_i}[t_j]$, $R_i = \mathbb{E}_{j \sim F_i}[r_j]$

- Agent's **action** i^* maximizes $T_i - c_i$
- Principal's **revenue** $R_{i^*} - T_{i^*}$

The Model

- Type is drawn from G with density g supported on $C = [\underline{c}, \bar{c}]$

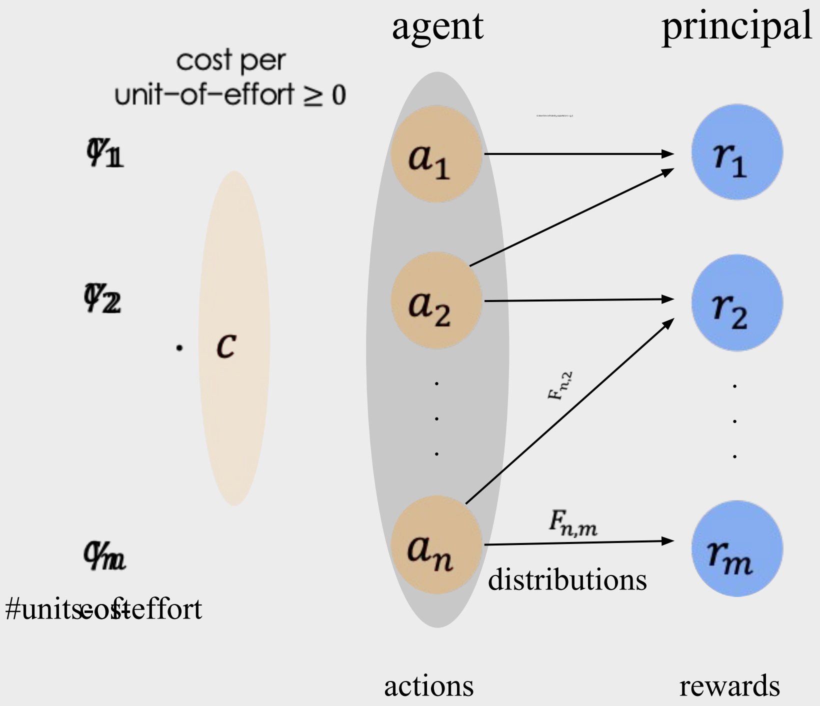

The Model

- A **contract** (x, t) consists of two mappings:
  - **Payment rule** $t: C \to \mathbb{R}^m$ from types to a **paymet** scheme
  - **Allocation rule** $x: C \to [n]$ from types to an **action** recommendation

Notation. $T_i^c = \mathbb{E}_{j \sim F_i}[t(c)_j]$

- **Agent** c report $\hat{c}$ and action $i^*(c)$ maximize $T_{i^*(c)}^{\hat{c}} - \gamma_{i^*(c)} c$
- A contract (x, t) is **incentive compatible (IC)** if $c = \hat{c}$ and $x(c) = i^*(c)$
- **Principal's contract** maximizes $\mathbb{E}_{c \sim G}[R_{x(c)} - T_{x(c)}^c]$ s.t. IC
- **Welfare** sum of utilities $\mathbb{E}_{c \sim G}[R_{i^*(c)} - \gamma_{i^*(c)} c]$

Design Space Characterization

Definition. Allocation rule x is **implementable** if exists payment rule t s.t. (x, t) is IC

Q. What do implementable allocation rules look like?

pure hidden-action
characterization

pure hidden-type
characterization

Design Space Characterization

Definition. Allocation rule x is **implementable** if exists payment rule t s.t. (x, t) is IC

Q. What do implementable allocation rules look like?

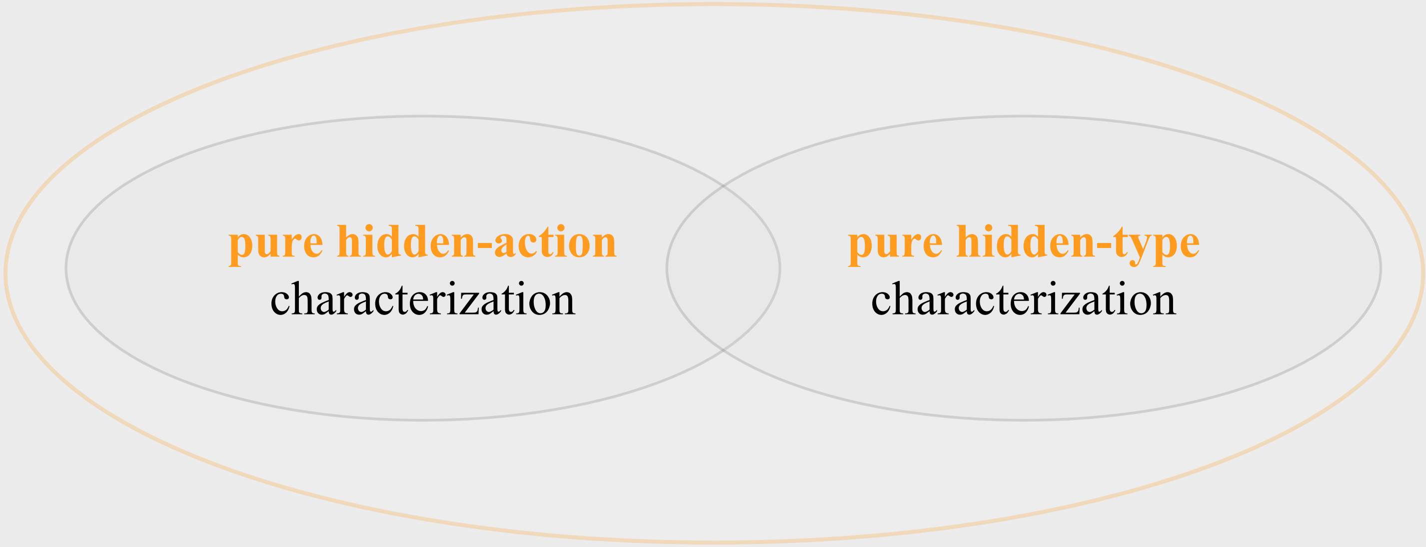

Design Space Characterization

Pure hidden action

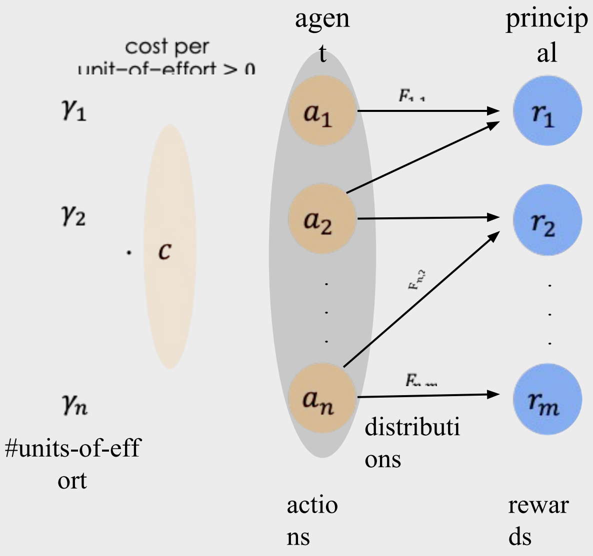

Design Space Characterization

Pure hidden action

- x is **implementable** $\Leftrightarrow$ exists no **deviation scheme** λ_k s.t. (1) dominant distribution $\sum_k \lambda_k F_k \geq F_x$
  (2) strictly lower cost $\sum_k \lambda_k c_k < c_x$

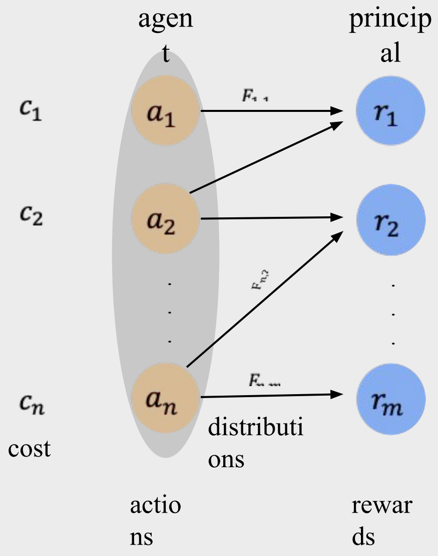

Design Space Characterization

Pure hidden action

- x is **implementable** $\Leftrightarrow$ exists no **deviation scheme** λ_k s.t. (1) dominant distribution $\sum_k \lambda_k F_k \geq F_x$

  (2) strictly lower cost $\sum_k \lambda_k c_k < c_x$

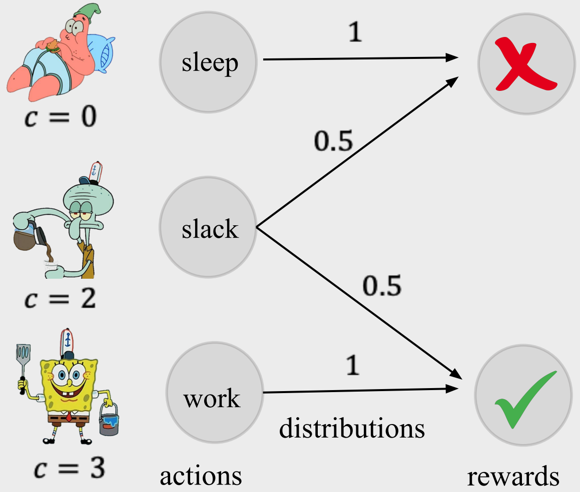

Design Space Characterization

Pure hidden action

- x is **implementable** $\Leftrightarrow$ exists no **deviation scheme** λ_k s.t. (1) dominant distribution $\sum_k \lambda_k F_k \geq F_x$

  (2) strictly lower cost $\sum_k \lambda_k c_k < c_x$

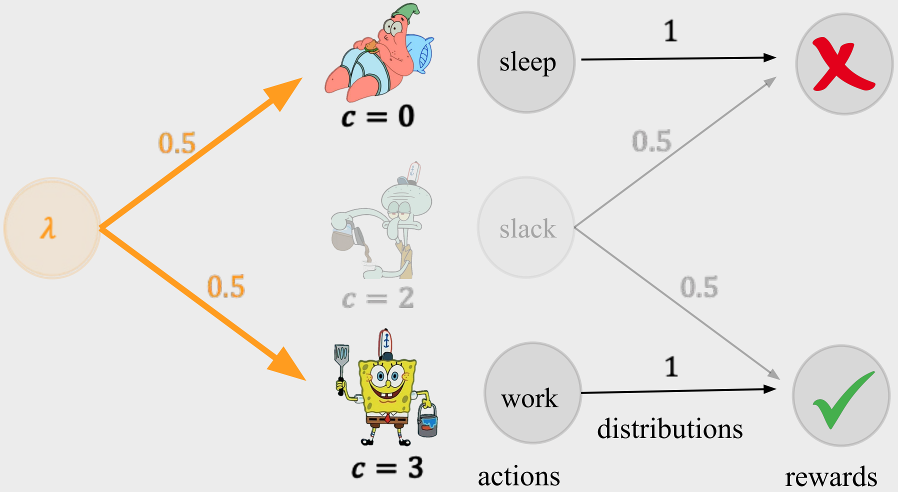

 is not implementable

Design Space Characterization

Pure hidden action

- x is **implementable** $\Leftrightarrow$ exists no **deviation scheme** λ_k s.t. (1) dominant distribution $\sum_k \lambda_k F_k \geq F_x$
  (2) strictly lower cost $\sum_k \lambda_k c_k < c_x$

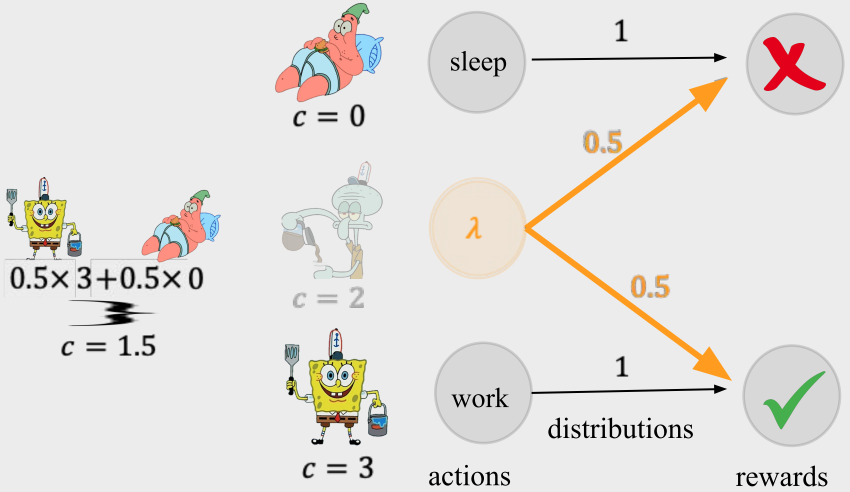

is not implementable

- LP duality approach

Design Space Characterization

Pure hidden type

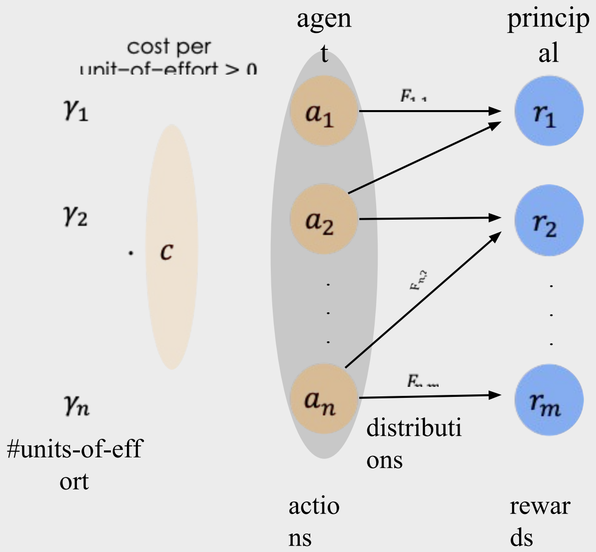

Design Space Characterization

Pure hidden type

- x is **implementable** $\Leftrightarrow$ x is **monotone** [Myerson 1981]

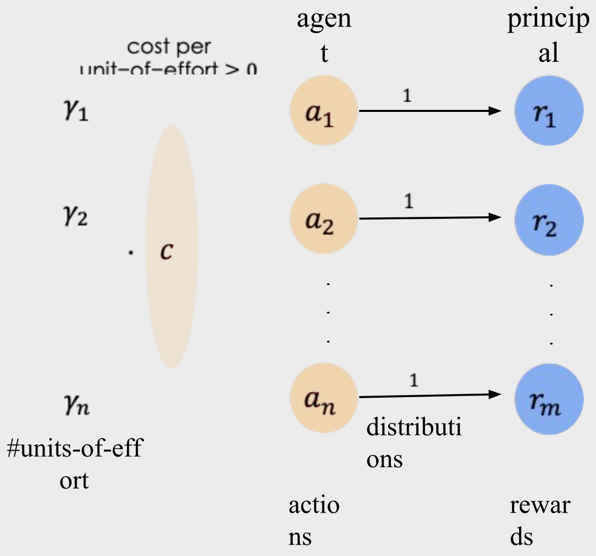

Design Space Characterization

Pure hidden type

- x is **implementable** $\Leftrightarrow$ x is **monotone** [Myerson 1981]

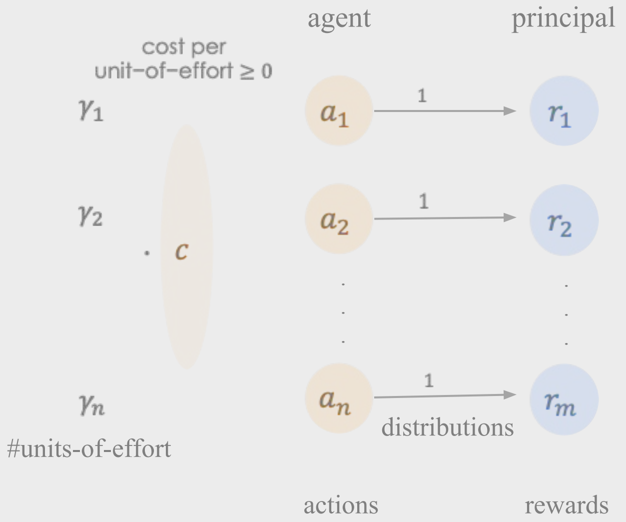

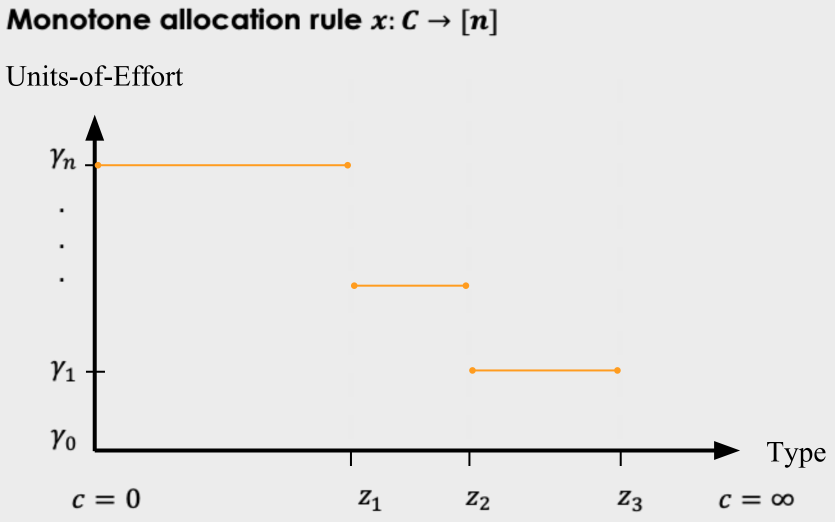

Design Space Characterization

Hidden type and hidden action

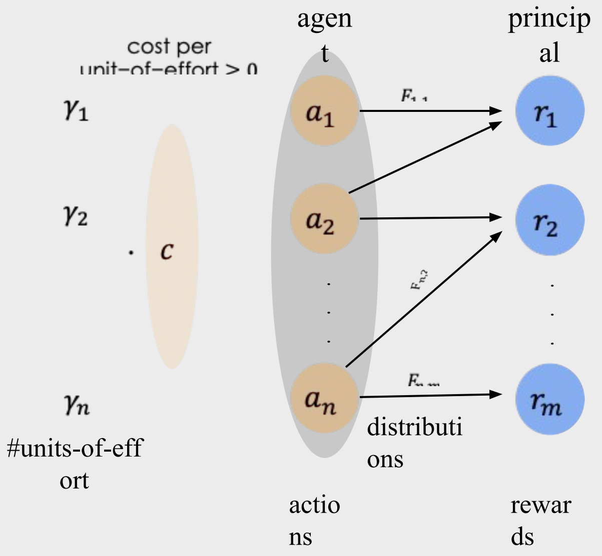

Design Space Characterization

Hidden type and hidden action

Proposition [ADT EC'21]. If x is implementable, it is **monotone**

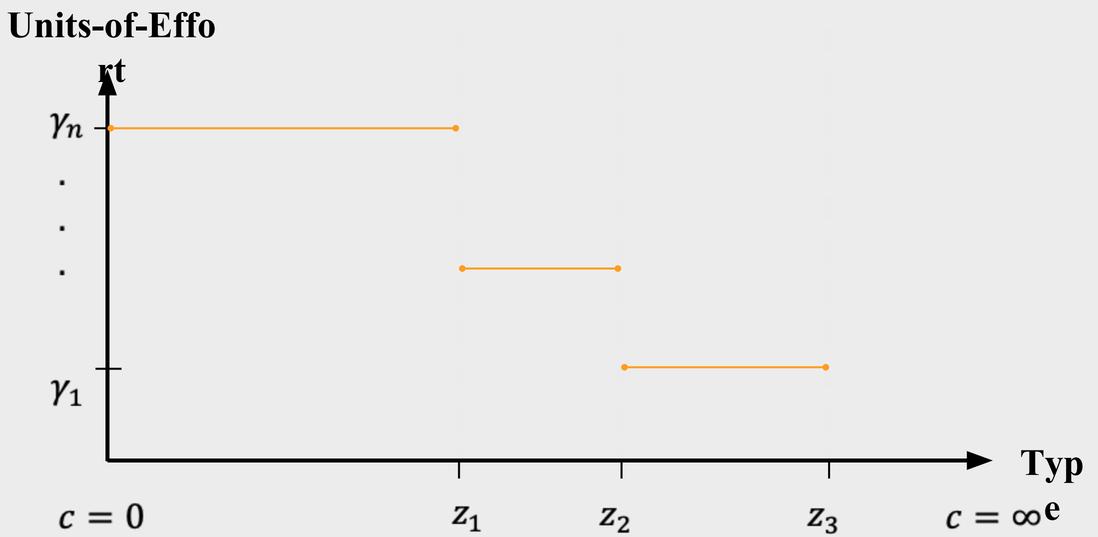

Monotone piecewise-constant allocation rule $x: C \to [n]$

Design Space Characterization

Theorem [ADT EC'21]. x **implementable** $\Leftrightarrow$ exists no **deviation scheme** $\lambda_{(z,k)}$ s.t.

(1) dominant sum of distributions $\sum_{z,k} \lambda_{(z,k)} F_k \geq \sum_z F_{x(z)}$

(2) strictly lower joint cost $\sum_{z,k} \lambda_{(z,k)} \gamma_k z < \sum_z \gamma_{x(z)} z$

Corollary [ADT EC'21]. Optimal contract is **polytime computable** for const #actions

Hardness for constant #actions in the multi-parameter model [Guruganesh-Schneider-Wang'21]

Example

Theorem [ADT EC'21]. x **implementable** $\Leftrightarrow$ exists no **deviation scheme** $\lambda_{(z,k)}$ s.t.

(1) dominant sum of distributions $\sum_{z,k} \lambda_{(z,k)} F_k \geq \sum_z F_{x(z)}$

(2) strictly lower joint cost $\sum_{z,k} \lambda_{(z,k)} \gamma_k z < \sum_z \gamma_{x(z)} z$

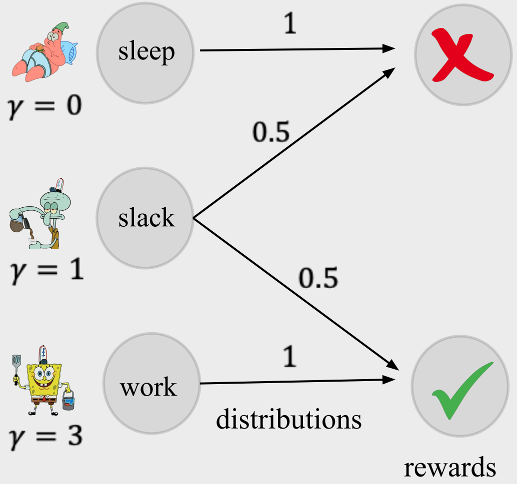

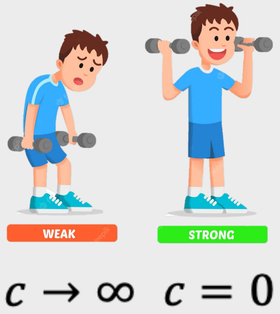

Example

Theorem [ADT EC'21]. x **implementable** $\Leftrightarrow$ exists no **deviation scheme** $\lambda_{(z,k)}$ s.t.

(1) dominant sum of distributions $\sum_{z,k} \lambda_{(z,k)} F_k \geq \sum_z F_{x(z)}$

(2) strictly lower joint cost $\sum_{z,k} \lambda_{(z,k)} \gamma_k z < \sum_z \gamma_{x(z)} z$

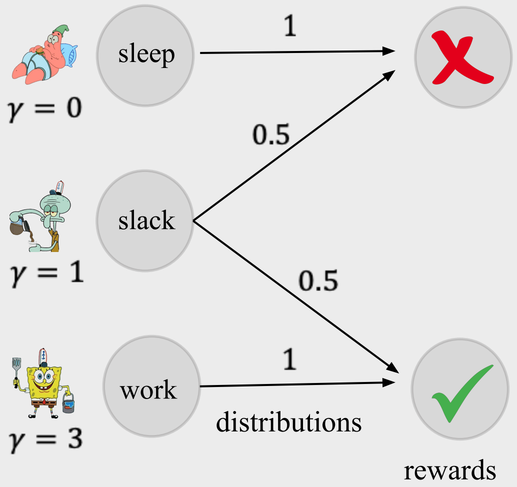

allocation rule: 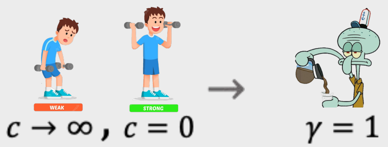

(1) distributions sum (1,1)

(2) joint cost ∞

Example

Theorem [ADT EC'21]. x **implementable** $\Leftrightarrow$ exists no **deviation scheme** $\lambda_{(z,k)}$ s.t.

(1) dominant sum of distributions $\sum_{z,k} \lambda_{(z,k)} F_k \geq \sum_z F_{x(z)}$

(2) strictly lower joint cost $\sum_{z,k} \lambda_{(z,k)} \gamma_k z < \sum_z \gamma_{x(z)} z$

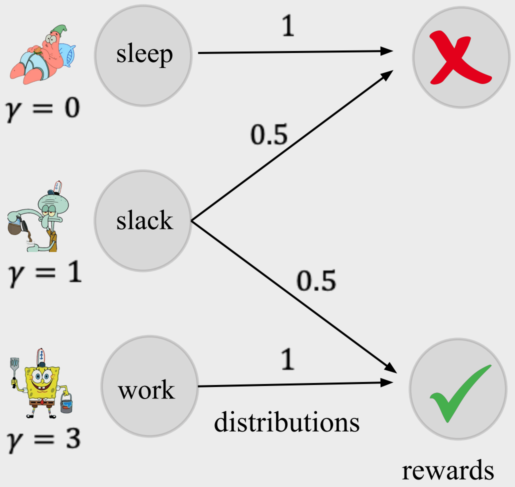

allocation rule:

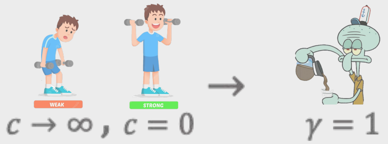

(1) distributions sum (1,1)

(2) joint cost ∞

deviation:

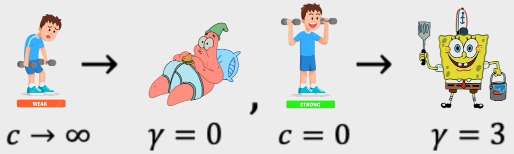

(1) distributions sum (1,1)

(2) joint cost 0

Optimal Contracts and Their Issues

- Informational requirements, extensive analysis ,etc.
- Unintuitive, e.g., non-monotonicity in rewards [DRT EC'19]

Theorem [ADLT EC'23]. In the single dimensional typed model

- Large **menu-size complexity**
- Revenue **non-monotonicity** w.r.t type distribution

Simple Contracts

- In a **linear contract**, the principal offers a fixed share $\alpha \in [0,1]$ of the rewards

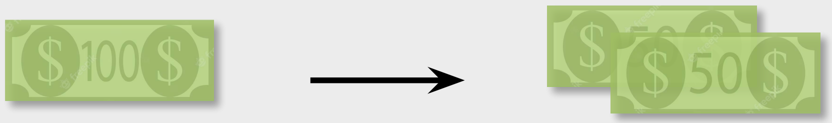

100$

50$ to agent
50$ to principal

“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.” [Milgrom and Holmstrom 1987]

Simple Contracts

- In a **linear contract**, the principal offers a fixed share $\alpha \in [0,1]$ of the rewards

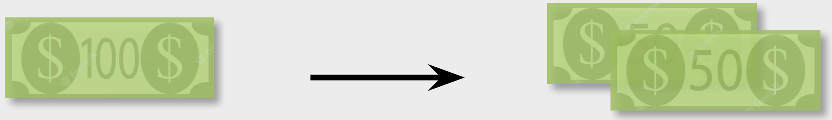

100$

50$ to agent
50$ to principal

- **Robustness of linear contracts.** Carroll)2015(, Dütting et al.)2019(, Yu and Kong)2020(, Dai and Toikka)2022(, Walton and Carroll)2022(
- **Approximation of linear contracts.** Dütting et al.)2019(, Castiglioni et al.)2021(, Guruganesh et al.)2021(

Near-Optimality of Linear Contracts

- $\theta(n)$ separation for **point-mass** distributions [DRT EC'19]
  - Boundary case
- **Approximately optimal** with **sufficient uncertainty**
  - The small-tail assumption

principal-agent instances

Not small-tail	Small-tail
Point-mass like	**Sufficient uncertainty**
[Dütting et al. 2019]	**Linear contracts are (near)-optimal!**

Near-Optimality of Linear Contracts

Theorem [ADLT EC'23]. Revenue benchmark:

- 3-approximation for normal $\mathcal{N}(\mu, \sigma^2)$ truncated at $c = 0$ with $\sigma \geq 5\eta/2\sqrt{2}$
- 2-approximation for uniform $U[0, \bar{c}]$
  - Optimal when $i^*(r, \bar{c}) = 0$
- 2-approximation for decreasing densities (e.g., exponential)

- Constant approximation w.r.t optimal welfare benchmark [ADLT EC'23]

The Small-Tail Assumption

Definition [ADLT EC'23]. Let $\kappa \in [\underline{c}, \bar{c}]$, $\eta \in [0,1]$.

An instance is **(κ, η)-small-tail** if $\mathrm{Wel}_{[\kappa,\bar{c}]} \geq \eta \mathrm{Wel}_{[\underline{c},\bar{c}]}$.

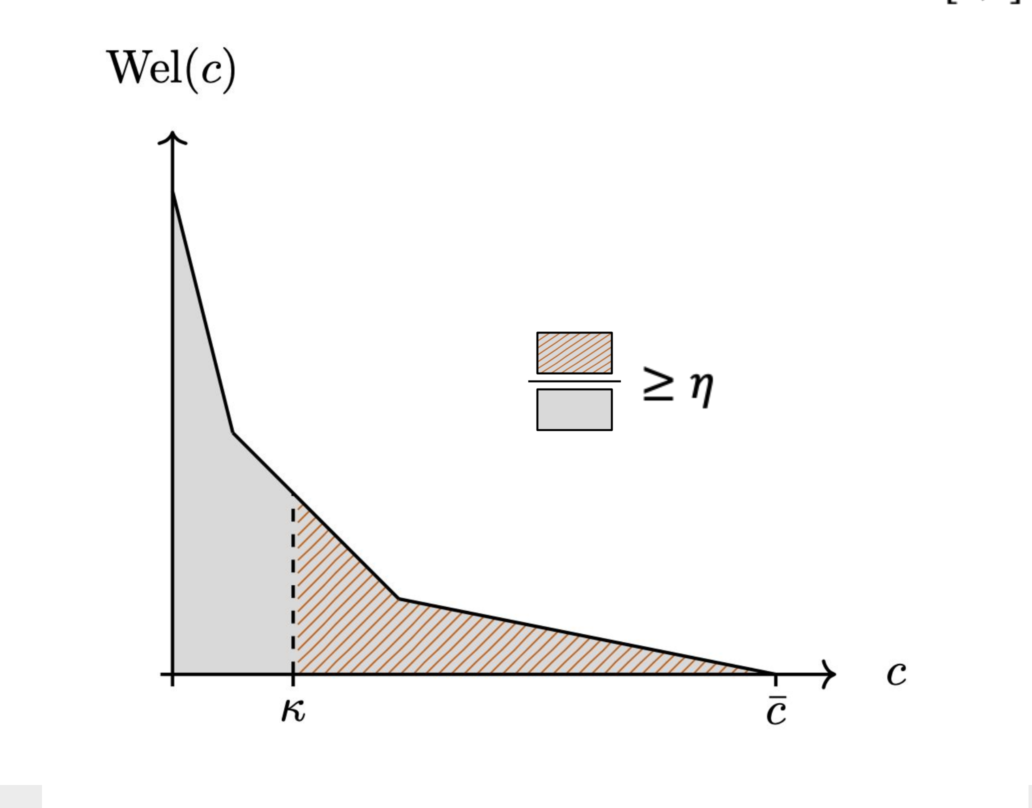

The Small-Tail Assumption

Depends on the entire principal-agent setting

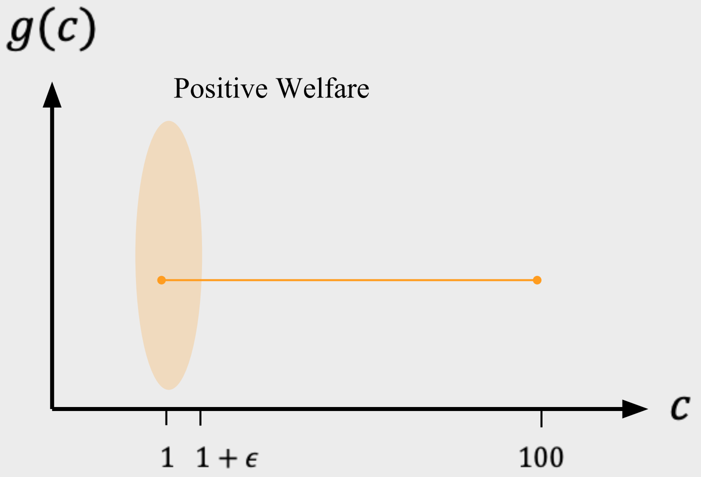

High Uncertainty

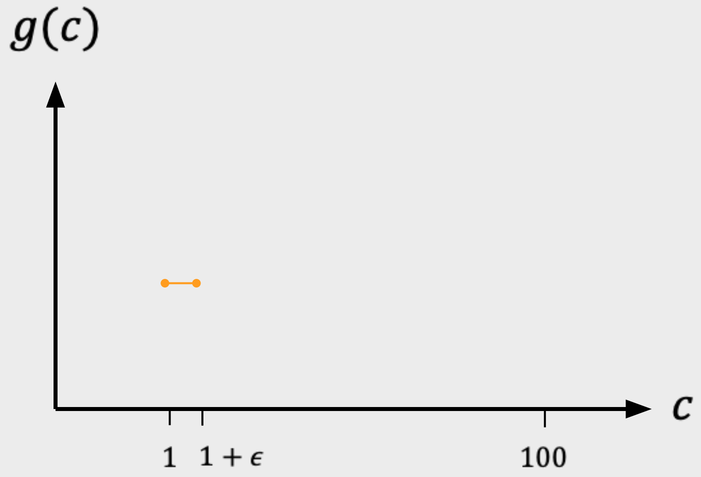

Low Uncertainty

Dütting et al.)2019(

Universal Approximation Guarantee

Theorem [ADLT EC'23]. Let $q \in (0,1)$ and $G(c_q) = q$. If for $\alpha, \eta \in (0,1)$ the settings is $(\frac{c_q}{\alpha}, \eta)$**-small-tail** then linear contract α is at least $(1-\alpha)\eta q$ of the **optimal welfare**

Slowly-Increasing Distributions

- Applies to any CDF and captures its **rate of increase**
- Parametric **approximation** of linear contracts
- Any distribution is slowly-increasing for **some parameters**

Slowly-Increasing Distributions

Definition [ADLT EC'23]. Let $\alpha, \beta \in (0,1)$, and $\kappa \in [\underline{c}, \bar{c}]$.

A distribution G is (α, β, κ)-**slowly-increasing** if $G(c) \le \frac{1}{\beta} G(\alpha c)\ \forall \kappa \le c$

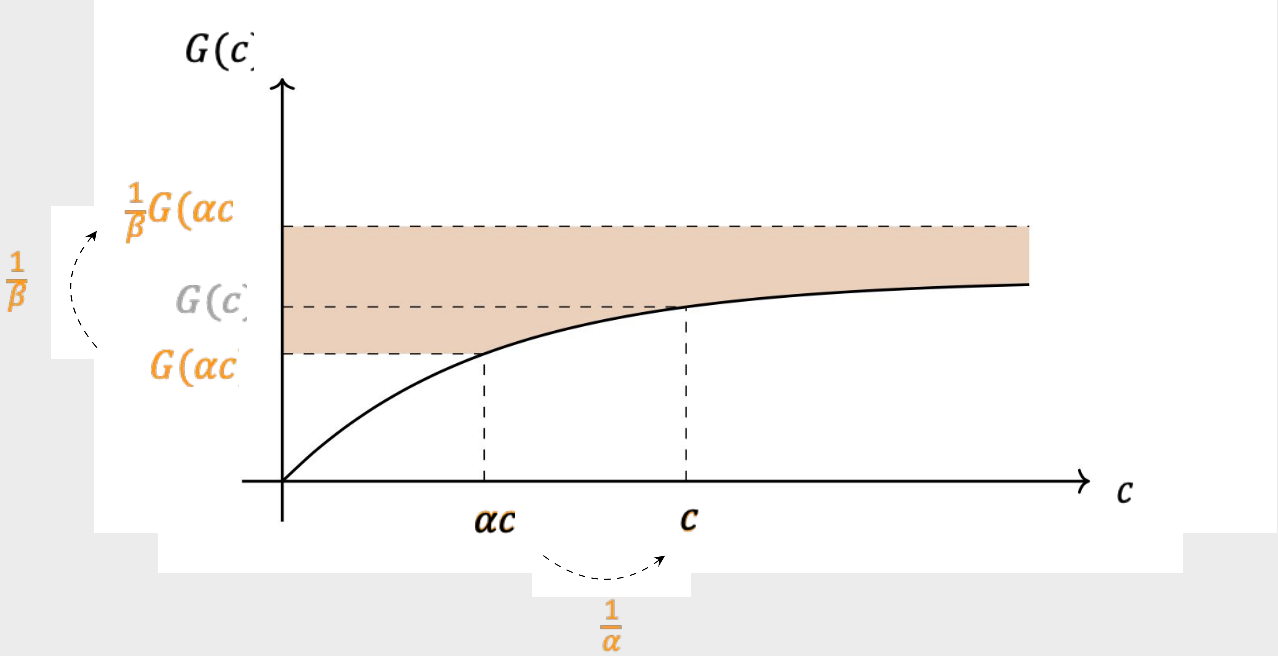

Approximation for Slowly Increasing

Theorem [ADLT EC'23]. Let $\alpha, \beta, \eta \in (0,1)$, and $\kappa \in \left[\underline{c}/_{\alpha}, \bar{c}\right]$.

Under **(α, β, κ)-slowly-increasing** and **(κ, η)-small-tail**

linear contract α is $(1-\alpha)\beta\eta$ of the **optimal welfare**

Proof Idea for Slowly Increasing

Step 1. Revenue of linear contract α is at least $1-\alpha$ of its welfare

revenue of linear contract		welfare of linear contract
$\mathbb{E}_{c\sim G}[(1-\alpha)R_{x(c)}]$	$\geq (\mathbf{1}-\boldsymbol{\alpha})\times$	$\mathbb{E}_{c\sim G}[R_{x(c)}-\gamma_{x(c)}c]$

linear contract revenue $\geq (\mathbf{1}-\boldsymbol{\alpha})\times$ **linear contract welfare** $\geq (1-\alpha)\times\beta\times$ optimal welfare

Proof Idea for Slowly Increasing

Step 2. welfare of linear contract α is at least β of the optimal welfare

- When maximizing welfare, the agent maximizes $R_i - \gamma_i c$
- For linear contract α:

  The agent maximizes $\alpha R_i - \gamma_i c$ which is $R_i - \gamma_i \frac{c}{\alpha}$

linear contract revenue $\geq (1-\alpha) \times$ **linear contract welfare** $\geq (1-\alpha) \times \boldsymbol{\beta} \times$ **optimal welfare**

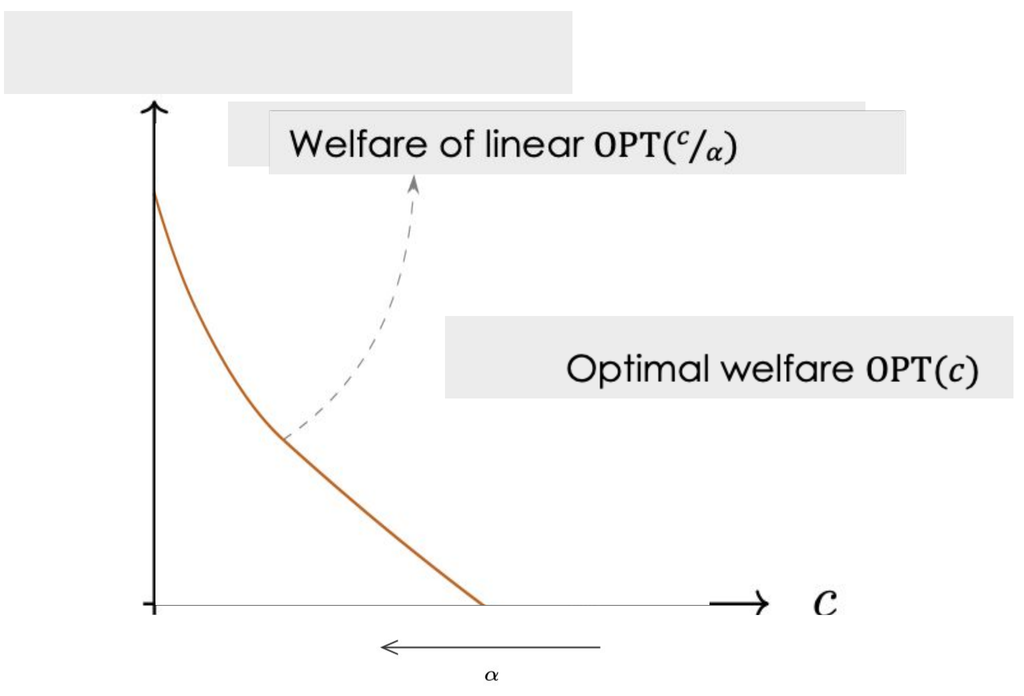

Summary and Future Directions

- **Single-parameter** model of types
- **Characterization** of the design space
- Counter-intuitive and **undesirable** properties of **optimal contracts**
- **Linear contracts** are near-optimal

Future directions:

- Other **forms of simple contracts** that are near-optimal
- Contracts that involve **multiple agents**
- Other models for incentivizing effort

!Thank You

alontal@campus.technion.ac.il

www.talalon.org

Research Agenda

Generalizations of the classic model:

- **Personalization** for participants from diverse population (with **types**)
- **Multilateral** contracts involving multiple principals or agents
  - Bad **welfare performance** for independent contracts.

    Identify unique welfare-maximizing for interdependent contracts [ALST EC'21,OR'23]
- The need for **simple** contracts

Research Agenda

Generalizations of the classic model:

- **Personalization** for participants from diverse population (with **types**)
- **Multilateral** contracts involving multiple principals or agents
  - Bad **welfare performance** for independent contracts.
    Identify unique welfare-maximizing for interdependent contracts [ALST EC'21,OR'23]
- The need for **simple** contracts

Design Space Characterization

Pure hidden type

- x is **implementable** $\Leftrightarrow$ x is **monotone** [Myerson 1981]

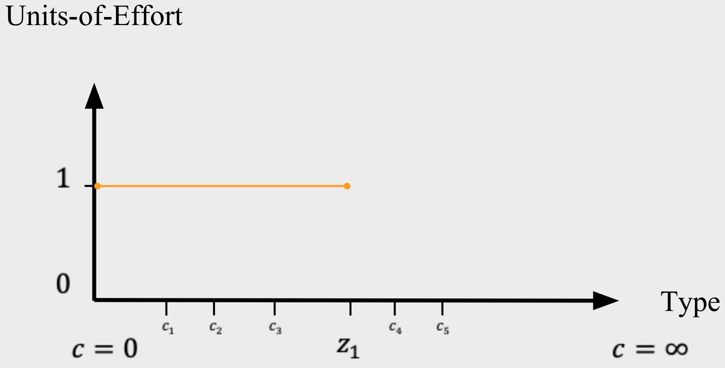

Notation. $S = \{c \mid x(c) = 1\}$

aggregated service $\sum_{c \in S} 1$

joint cost $\sum_{c \in S} c$

Exists no deviation of types $\lambda_{c,\gamma}$:

(1) same aggregated service $\sum_c \lambda_{c,1} = \sum_{c \in S} 1$

(2) lower joint cost $\sum_c \lambda_{c,1} c < \sum_{c \in S} c$

TODO!!!

Fix citations

Design Space Characterization

Absent hidden action

- x is **implementable** $\Leftrightarrow$ x is **monotone** [Myerson 1981]

Monotone allocation rule $x\colon C \to [n]$

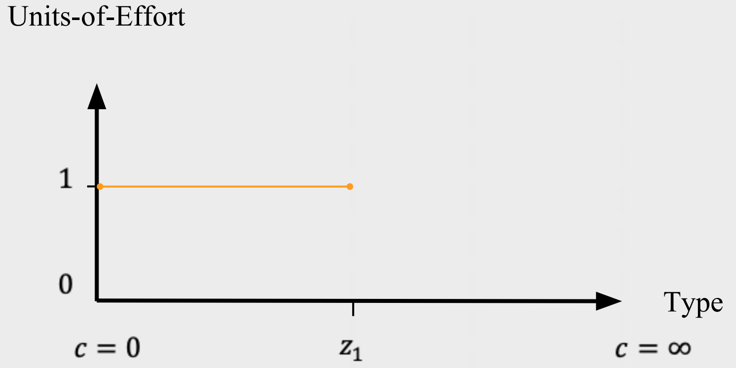

- **allocation rule:** $x(c) = \begin{cases} 1, & c \in [0, z_1), \\ 0, & c \in [z_1, \infty) \end{cases}$
- **payment rule:** $p(c) = \begin{cases} z_1, & c \in [0, z_1), \\ 0, & c \in [z_1, \infty) \end{cases}$
- **utility** of serving: $z_1 - c$, not serving: 0

Principals with Types

Common Agency

- **Multiple principals** hire a **single agent**
- What if the **principals** have different interests?

  Interests form principals' privately-known **types**

TODO

- We study welfare-maximizing contracts with hidden principal types
  - Such contracts incentivize the most rewarding action in aggregate

Common Agency

Theorem [Alon Lavi Shamash & T.-C. EC'21].

<u>The bad:</u>

- In some settings welfare maximization is **impossible**

<u>The good:</u>

- The class of settings with possibility is **large and identifiable**
- For this class, welfare maximization is **tractable**
- The contracts are **unique** and make economic sense

Summary

Some more work

Design Space Characterization

Absent hidden type

- x is **implementable** $\Leftrightarrow$ [GH 1983]

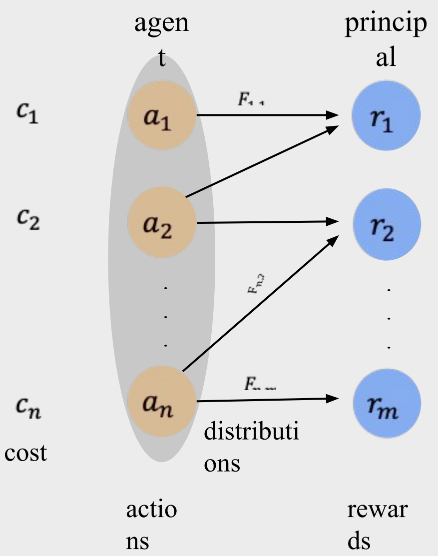

Cost

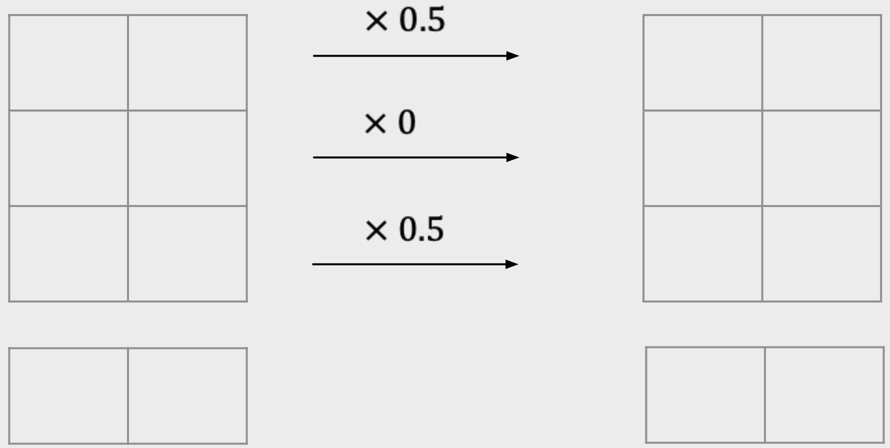

Intuition: Monotonicity Isn't Sufficient

Because

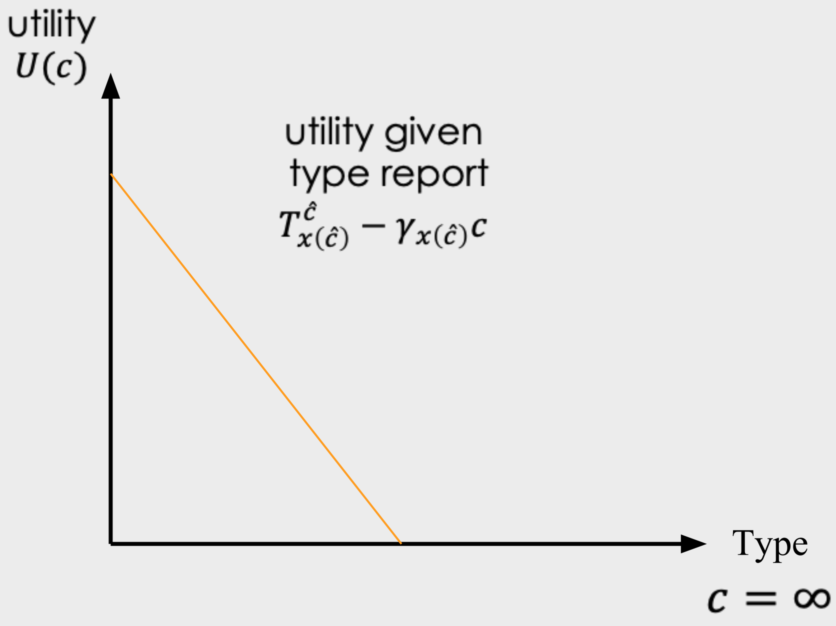

max over
linear
functions

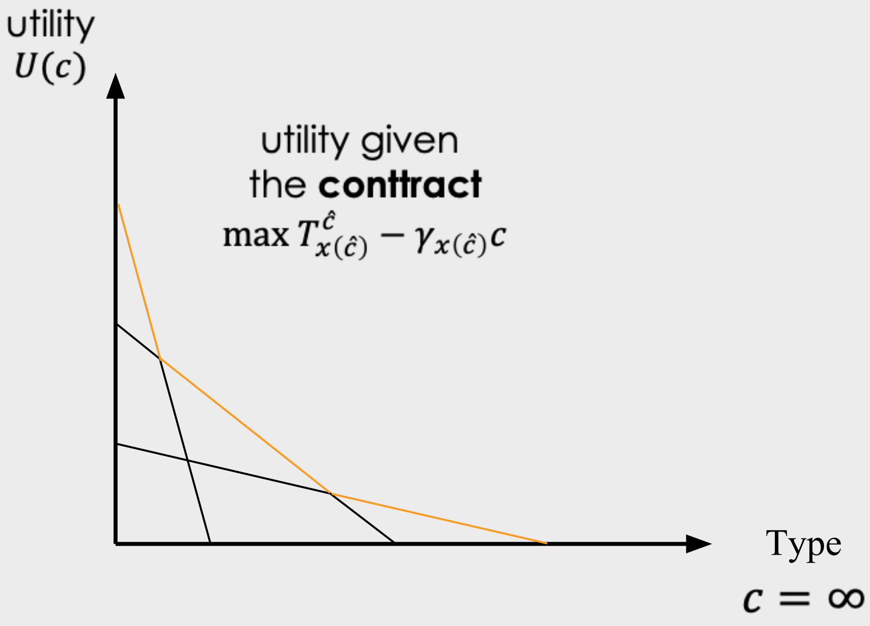

Design Space Characterization

Absent hidden type

- x is **implementable** $\Leftrightarrow$ [GH 1983]

Cost	✗	✓

Design Space Characterization

Absent hidden type

- x is **implementable** ⟺ [GH 1983]

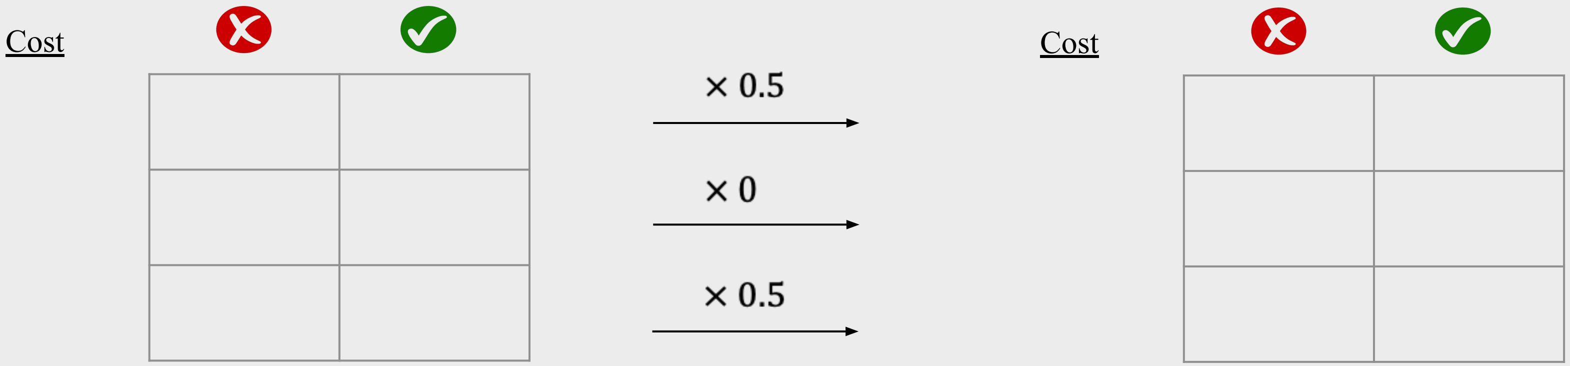

Design Space Characterization

Absent hidden type

- x is **implementable** $\Leftrightarrow$ [GH 1983]

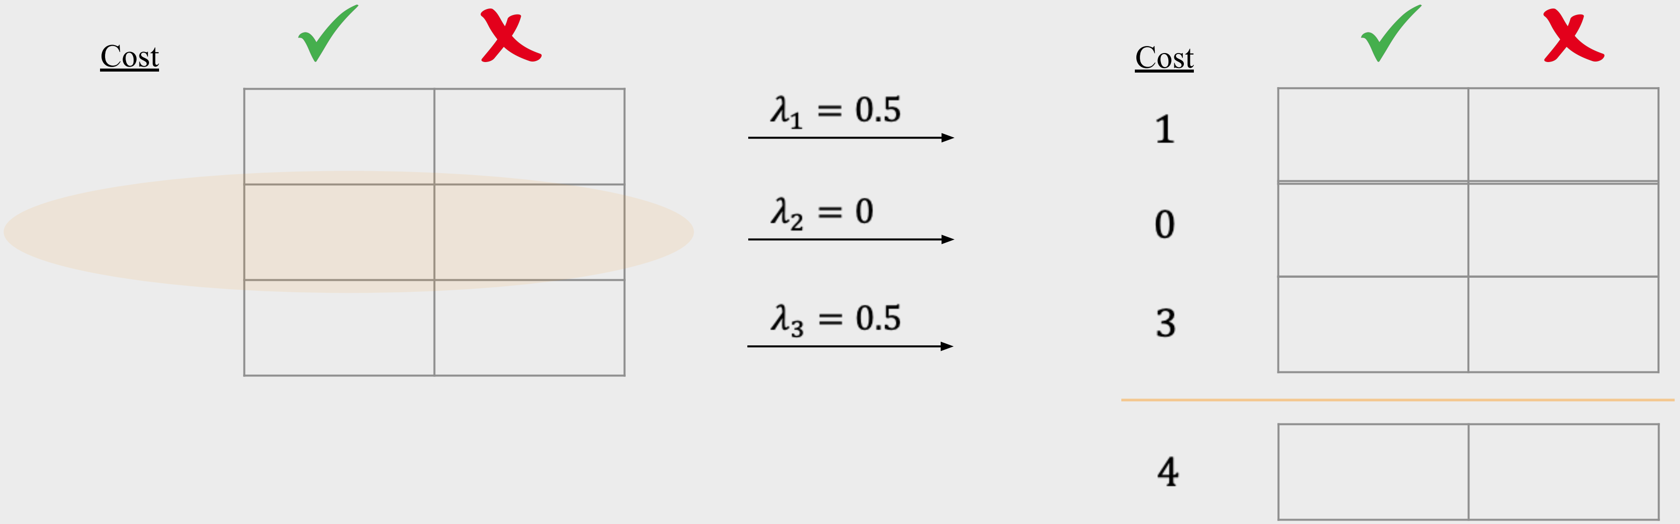

Design Space Characterization

Absent hidden type

- x is **implementable** $\Leftrightarrow$ [GH 1983]

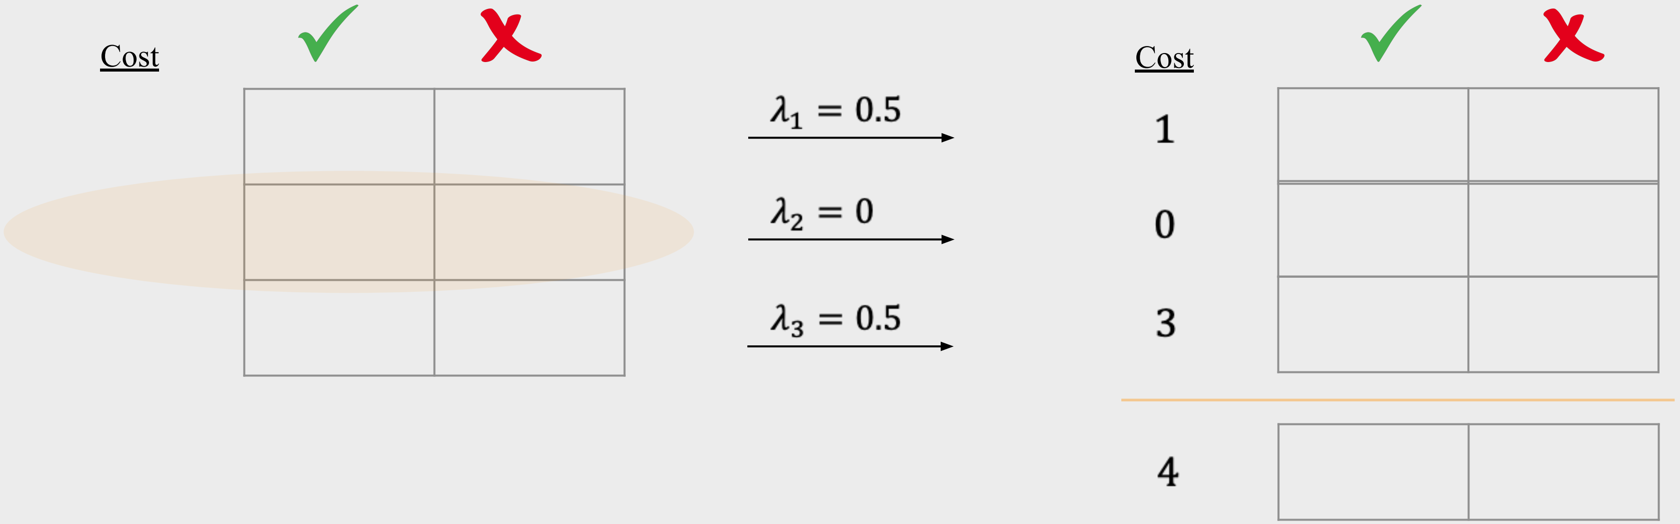

Intuition: Monotonicity Isn't Sufficient

No hidden-action

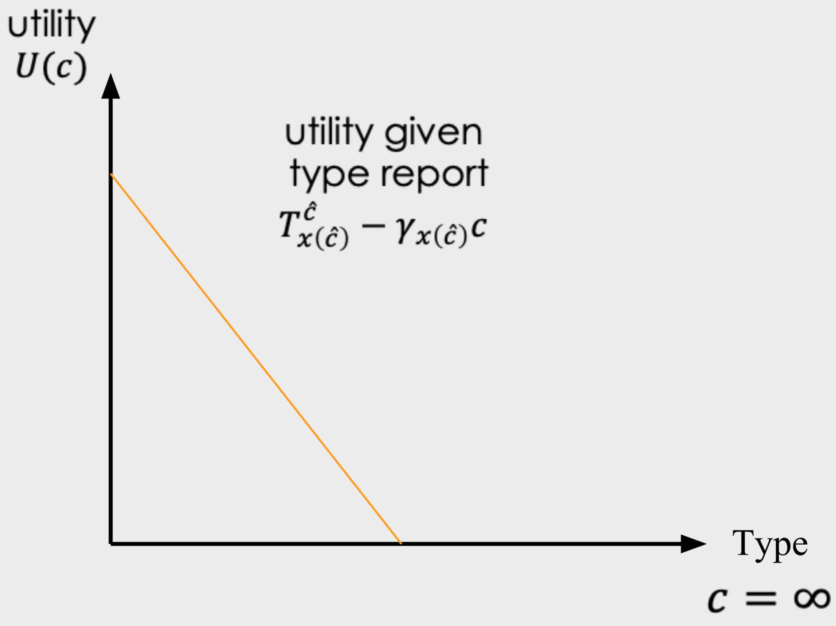

max over
linear
functions

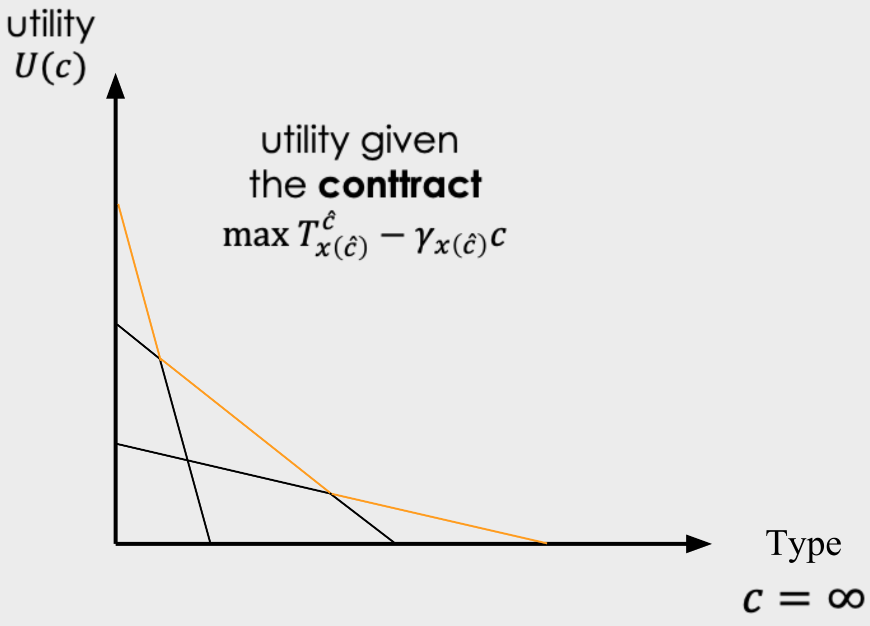

Intuition: Monotonicity Isn't Sufficient

With hidden-action

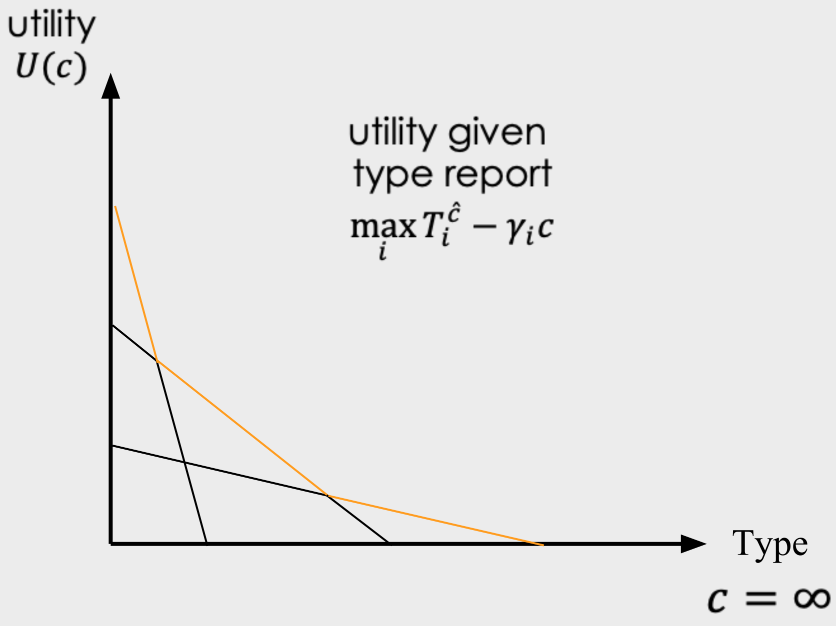

max over
convex
functions

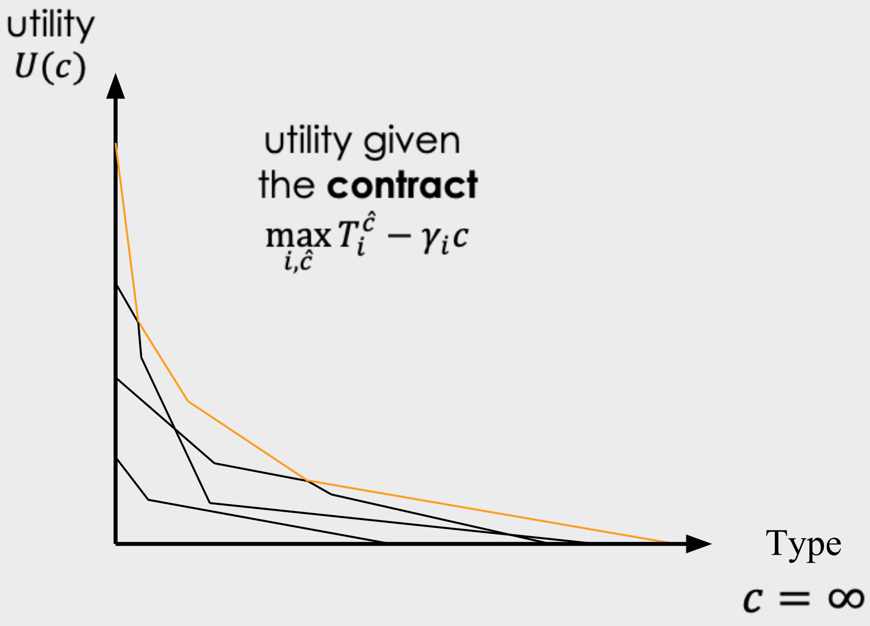

The Principal-Agent Problem [GH83]

- **Moral hazard** - the agent's actions cannot be observed
- **Objective:** a contract maximizing expected **rewards minus payment** $\mathbb{E}[r - t]$

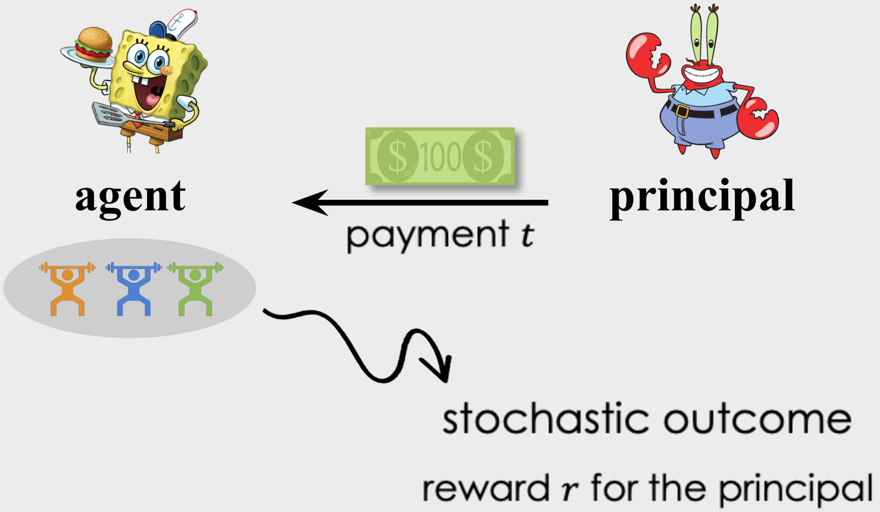

Contract Theory Approach

Proposition [Alon Dütting and T.-C. EC'21]. The following LP computes if x is implementable

$$
\begin{aligned}
\max \quad & 0 \\
\text{s.t.} \quad & F_{x(z)} t^{z} - \gamma_{x(z)} z \geq F_k t^{z'} - \gamma_k z \; \forall z, z' \in C, k \in [n] \\
& t_j^{z} \geq 0
\end{aligned}
$$