Contract Design Under Uncertainty

Tal Alon | PhD Candidate @ Technion, Israel Institue of Technology

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

Computation and Economics Seminar

The Hebrew University of Jerusalem

Feb 25, 2024

THE HEBREW UNIVERSITY OF JERUSALEM

Overview

- Algorithms in society
- Incentivizing **truthfulness** versus incentivizing **effort**
- The **principal-agent** model [GH83]
- Contract design with **types**

Modern Algorithms in Society

- Interactions with **self-interested** individuals
- In algorithmic game theory, we take **incentives** into account

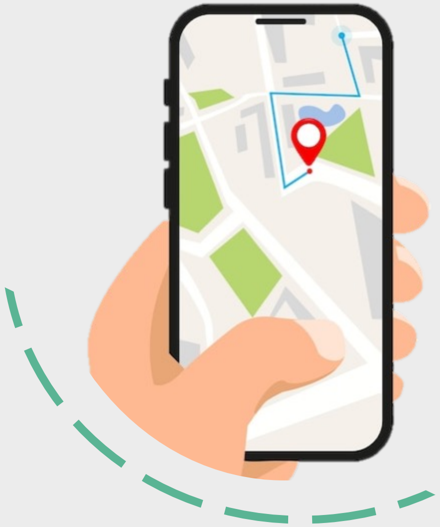

Social Media Marketing

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

Paid Ads

- Algorithmic **auction** - advertisers bid → allocation and payment
- Hidden type or adverse selection

Input ⟶ **Algorithm** ⟶ Output

- ✓ **Incentivize truthfulness**
- ✓ Computational complexity
- ✓ Achieving optimality

Influencer Marketing

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

Input → **Algorithm** → Output

- ✓ **Incentivize effort**
- ✓ Computational complexity
- ✓ Achieving optimality

The Principal-Agent Model [GH83]

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

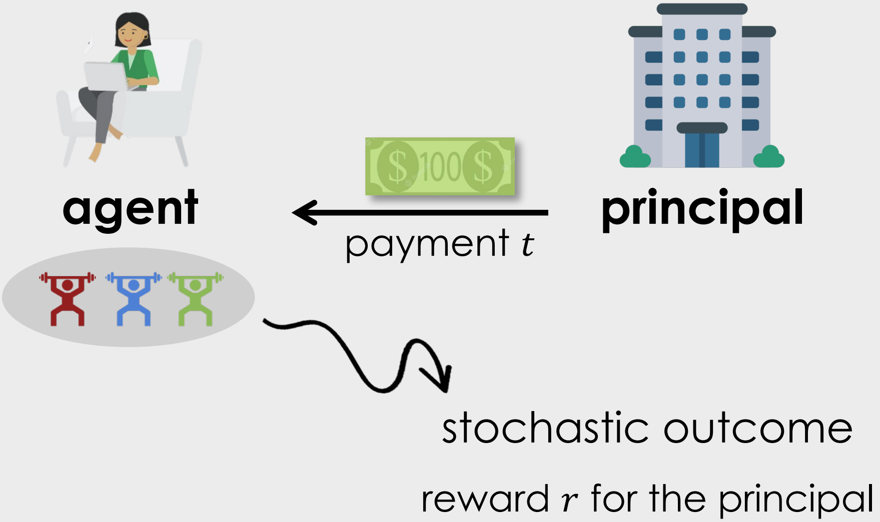

The Principal-Agent Model [GH83]

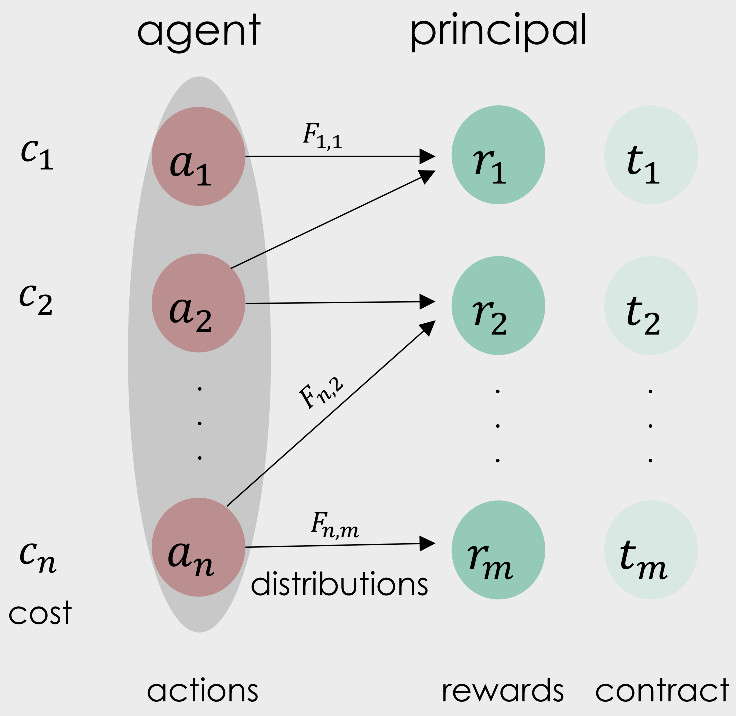

- Agent's **action i^*** maximizes $\sum_{j\in[m]} F_{i,j} t_j - c_i$
- Principal's **revenue** $\sum_{j\in[m]} F_{i^*,j} r_j - \sum_{j\in[m]} F_{i^*,j} t_j$

The Principal-Agent Model [GH83]

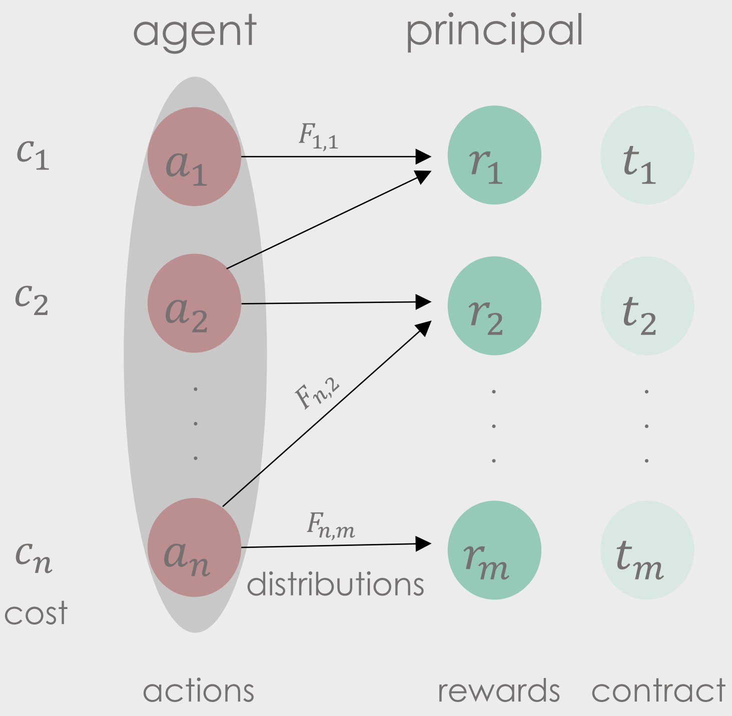

- Agent's **action i^*** maximizes $\sum_{j\in[m]} F_{i,j} t_j - c_i$
- Principal's **revenue** $\sum_{j\in[m]} F_{i^*,j} r_j - \sum_{j\in[m]} F_{i^*,j} t_j$

Min-payment LP for incentivizing action i:

$$\min_{t_1,\ldots,t_m} \quad \sum_{j\in[m]} F_{i,j} t_j$$

$$s.t. \quad \sum_{j\in[m]} F_{i,j} t_j - c_i \geq \sum_{j\in[m]} F_{i',j} t_j - c_{i'} \quad \forall i' \in [n] \quad (IC)$$

$$t_j \geq 0 \quad \forall j \in [m] \quad (LL)$$

Simple Example

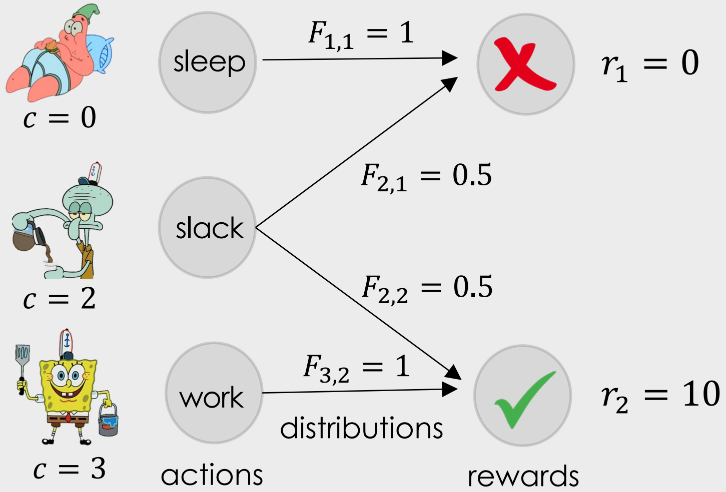

Min-payment LP for incentivizing working:

$$
\begin{aligned}
\min_{t_1,t_2} \quad & t_2 \\
s.t. \quad & t_2 - 3 \geq t_1 - 0 && (IC - sleep) \\
& t_2 - 3 \geq 0.5 \times t_1 + 0.5 \times t_2 - 2 && (IC - slack) \\
& t_1, t_2 \geq 0
\end{aligned}
$$

The optimal contract: $(t_1, t_2) = (0,3)$

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

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

Agenda

- Combine **moral hazard** and **adverse selection**

Input → **Algorithm** → Output

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

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)

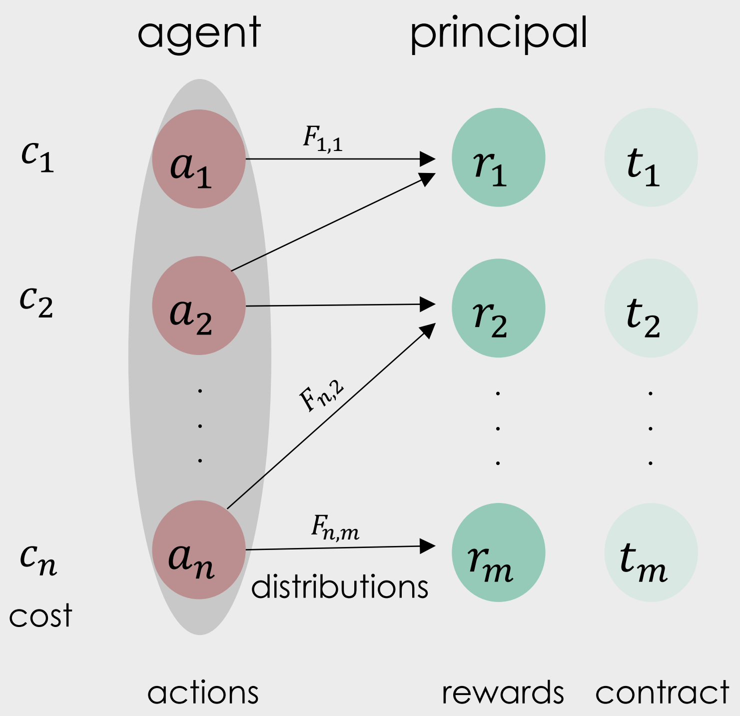

The Model

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

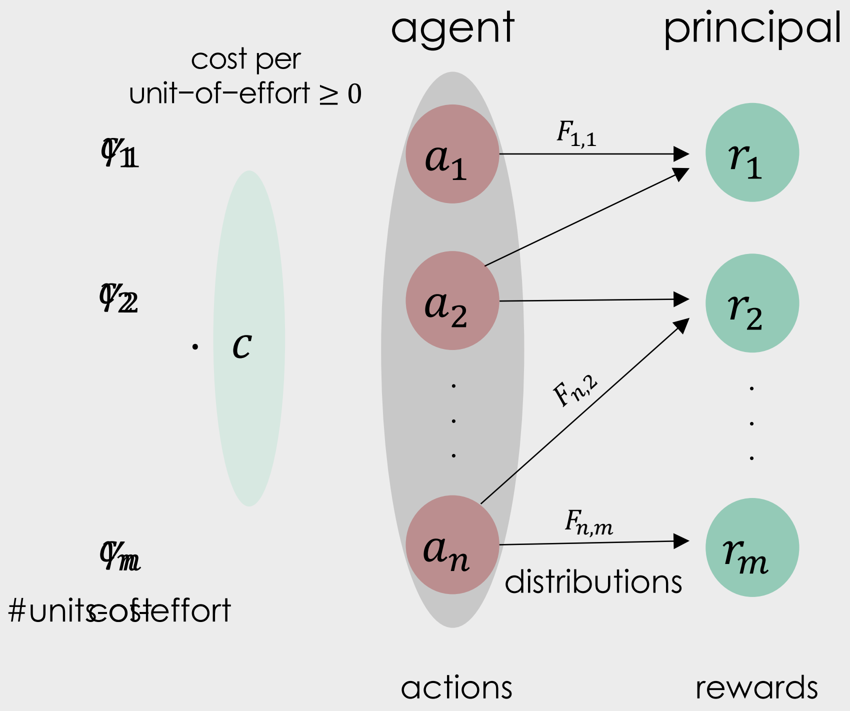

The Model

- A **contract** (x, t) consists of two mappings:
  - **Payment rule** $t: C \rightarrow \mathbb{R}^m$ from types to a **paymet** scheme
  - **Allocation rule** $x: C \rightarrow [n]$ 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 **revenue** $\mathbb{E}_{c \sim G}[R_{x(c)} - T_{x(c)}^c]$ 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. How 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. How do implementable allocation rules look like?

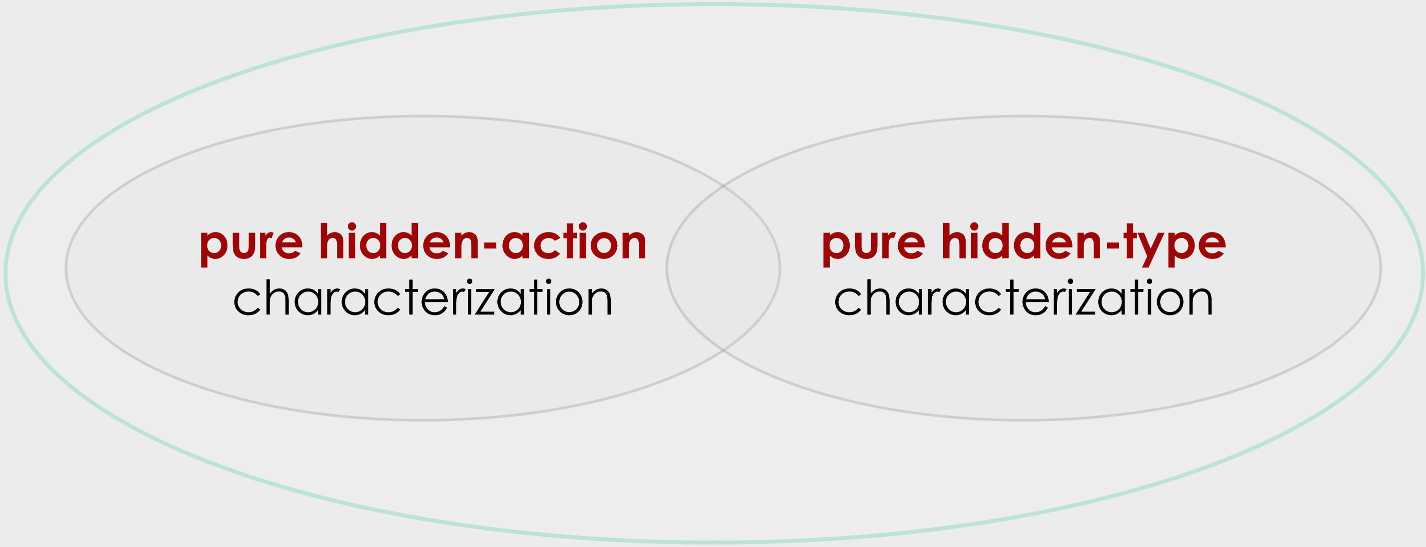

Design Space Characterization

Pure hidden type

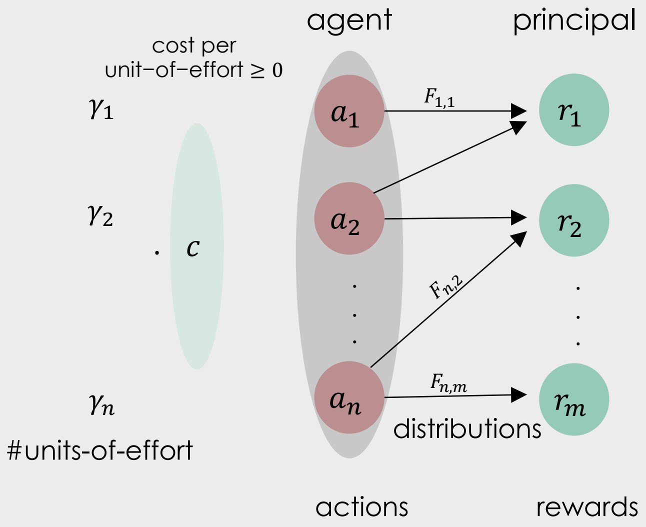

Design Space Characterization

Pure hidden type

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

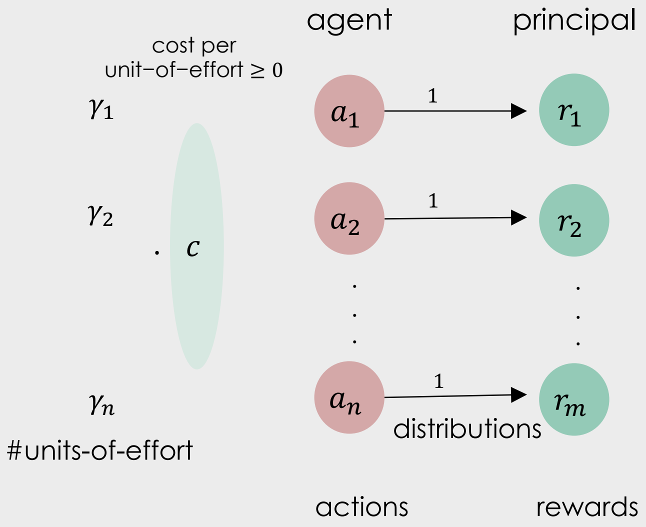

Design Space Characterization

Pure hidden type

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

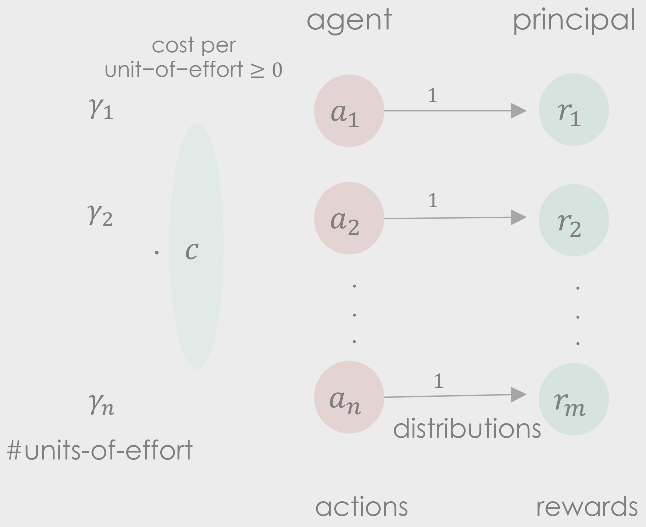

Monotone allocation rule $x: C \rightarrow [n]$

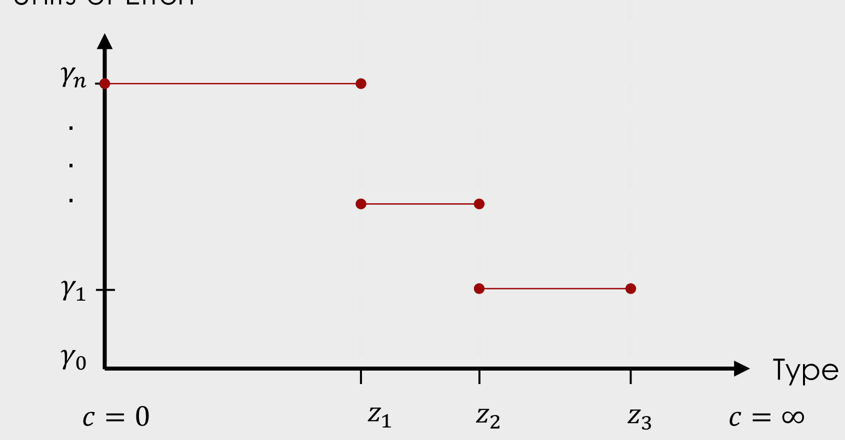

Design Space Characterization

Pure hidden action

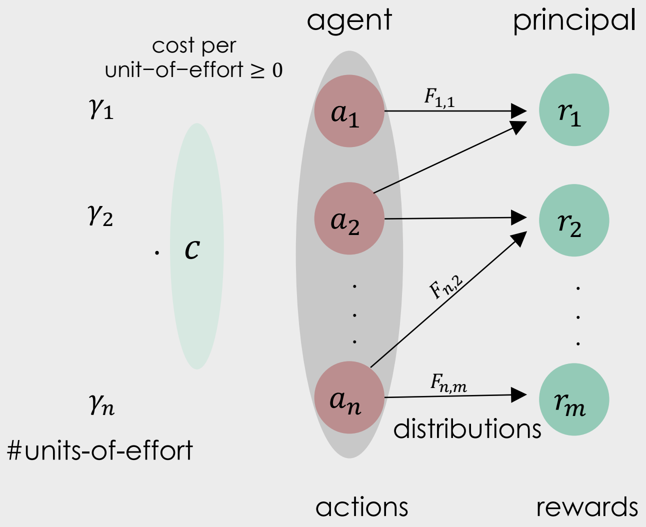

Design Space Characterization

Pure hidden action

- x is **implementable** $\Leftrightarrow \nexists$ **deviation scheme** $\lambda = (\lambda_1, \dots, \lambda_n)$ s.t. $\sum_{k \in [n]} \lambda_k = 1$ and:

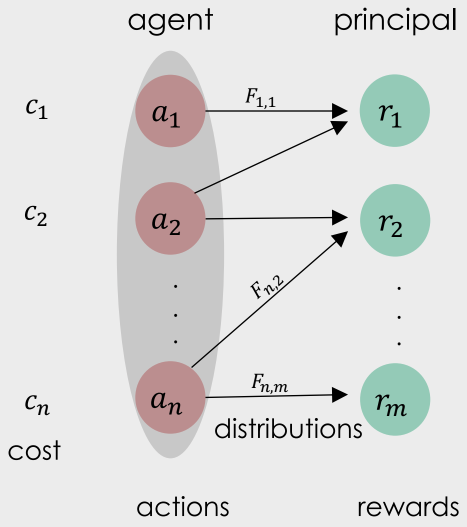

Design Space Characterization

Pure hidden action

- x is **implementable** $\Leftrightarrow$ $\nexists$ **deviation scheme** $\lambda = (\lambda_1, \ldots, \lambda_n)$ s.t. $\sum_{k \in [n]} \lambda_k = 1$ and:

  (1) $\sum_k \lambda_k F_k \geq F_x$ (dominant distribution)

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

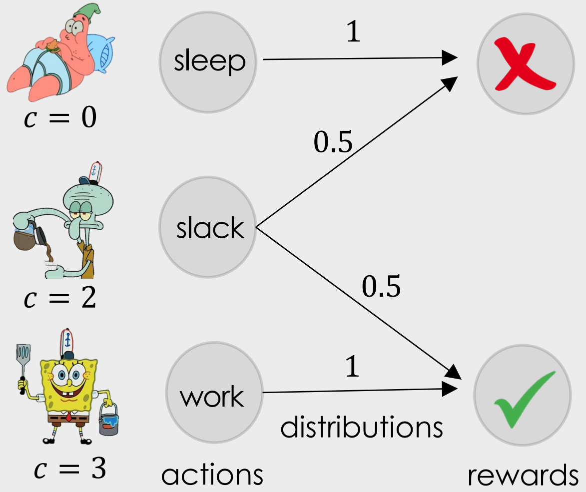

Design Space Characterization

Pure hidden action

- x is **implementable** $\Leftrightarrow \nexists$ **deviation scheme** $\lambda = (\lambda_1, \dots, \lambda_n)$ s.t. $\sum_{k \in [n]} \lambda_k = 1$ and:

  (1) $\sum_k \lambda_k F_k \geq F_x$ (dominant distribution)

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

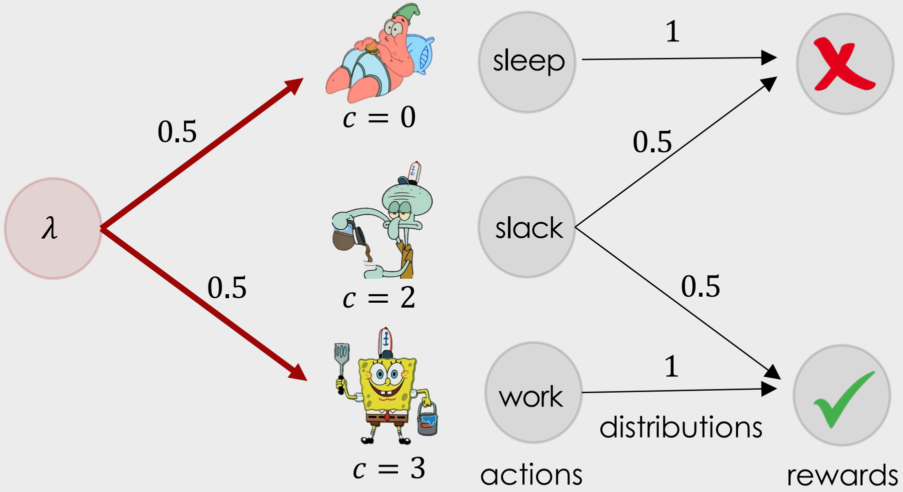

is not implementable

Design Space Characterization

Pure hidden action

- x is **implementable** $\Leftrightarrow \nexists$ **deviation scheme** $\lambda = (\lambda_1, \ldots, \lambda_n)$ s.t. $\sum_{k\in[n]} \lambda_k = 1$ and:

  (1) $\sum_k \lambda_k F_k \geq F_x$ (dominant distribution)

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

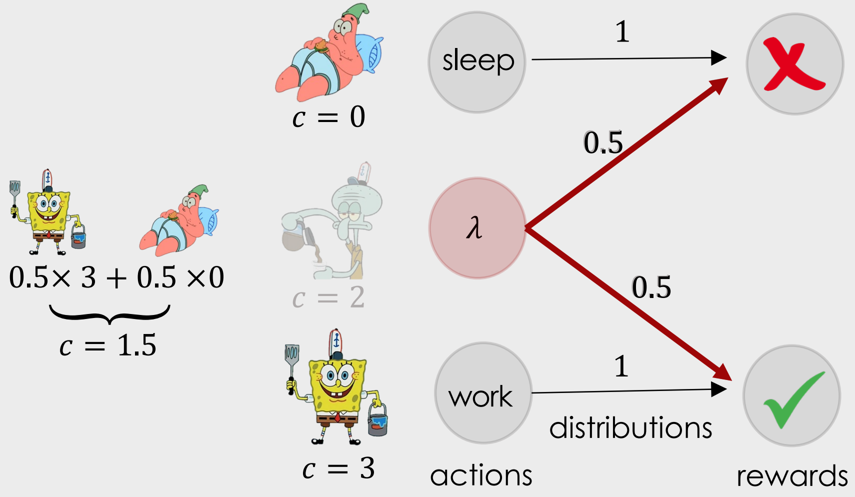

 is not implementable

Design Space Characterization

Hidden type and hidden action

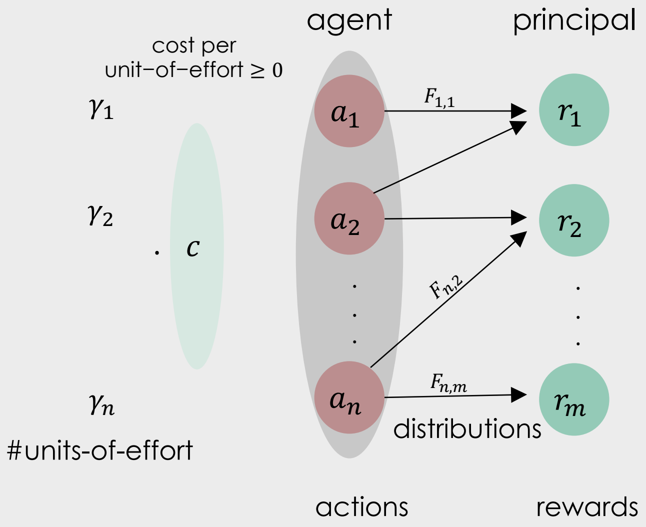

Design Space Characterization

Hidden type and hidden action

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

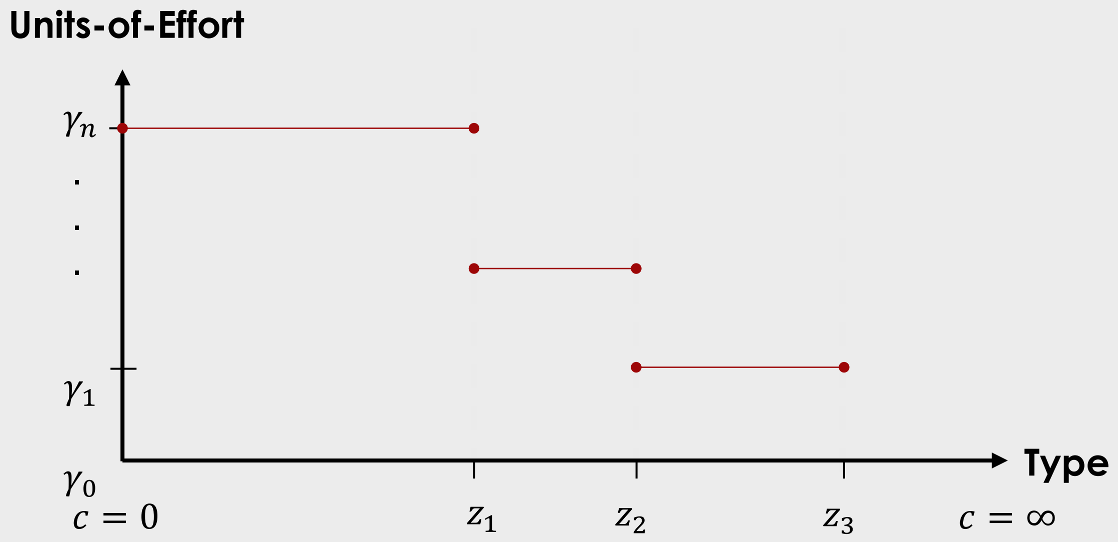

Monotone piecewise-constant allocation rule $x: C \to [n]$

Design Space Characterization

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

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]

Back to Our Example

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

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

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

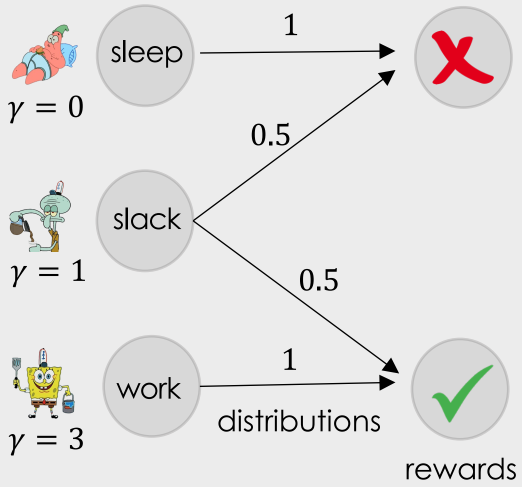

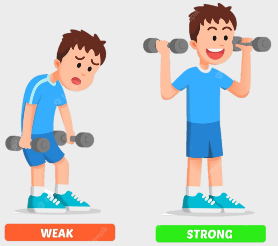

$c \to \infty$ $c = 0$

Back to Our Example

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

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

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

allocation rule: 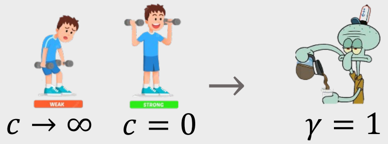

(1) distributions sum (1,1)

(2) joint cost ∞

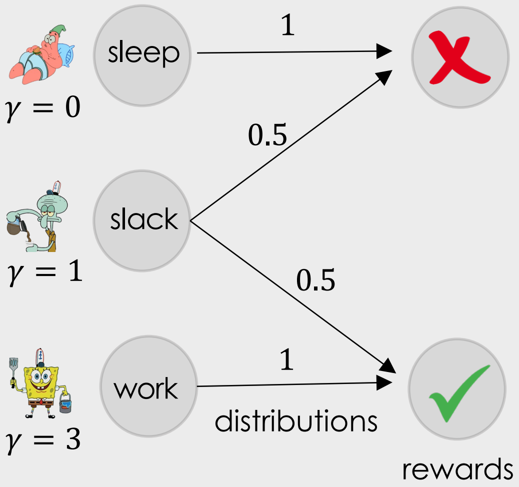

Back to Our Example

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

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

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

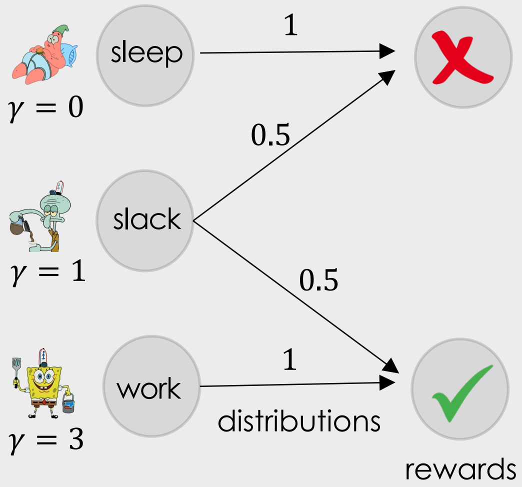

allocation rule:

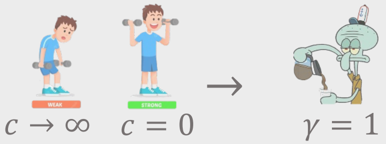

(1) distributions sum (1,1)

(2) joint cost ∞

deviation:

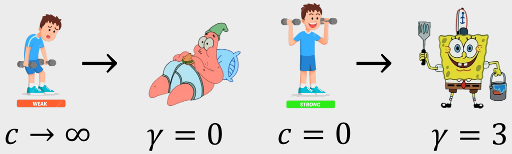

(1) distributions sum (1,1)

(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

"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

- **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 $\text{Wel}_{[\kappa,\bar{c}]} \geq \eta \text{Wel}_{[\underline{c},\bar{c}]}$

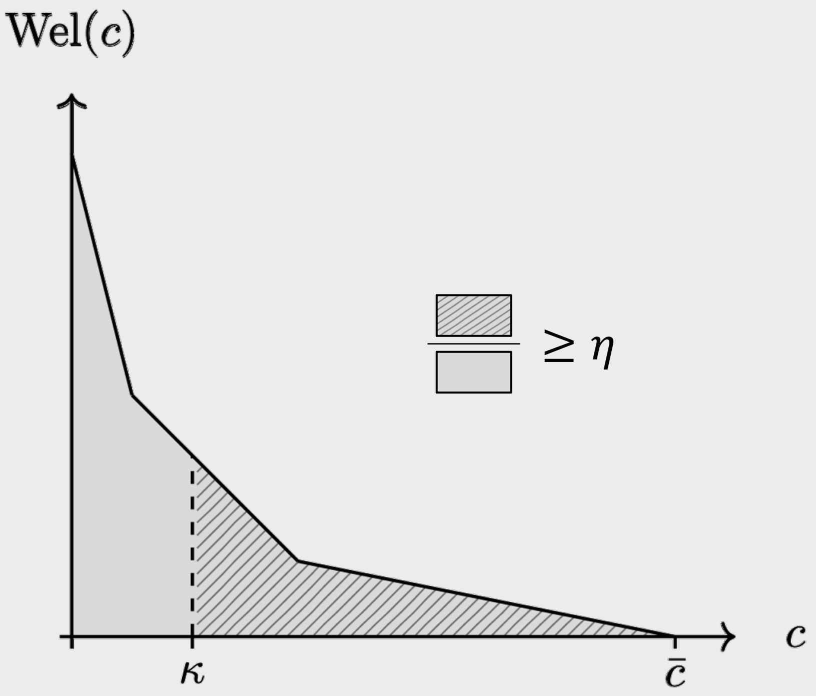

The Small-Tail Assumption

Depends on the entire principal-agent setting

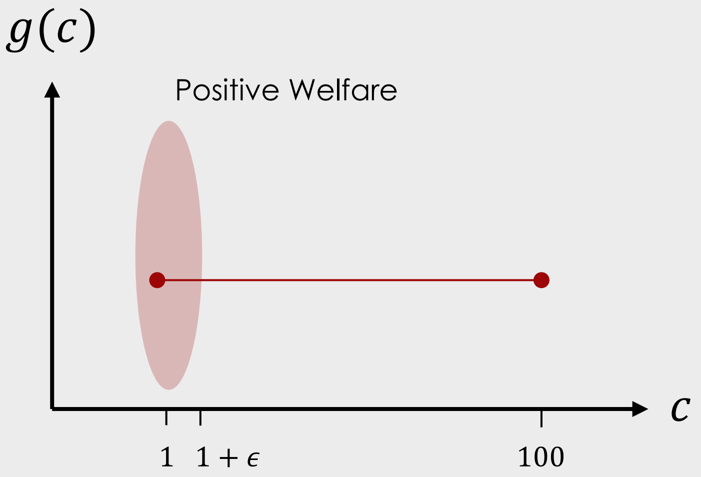

High Uncertainty

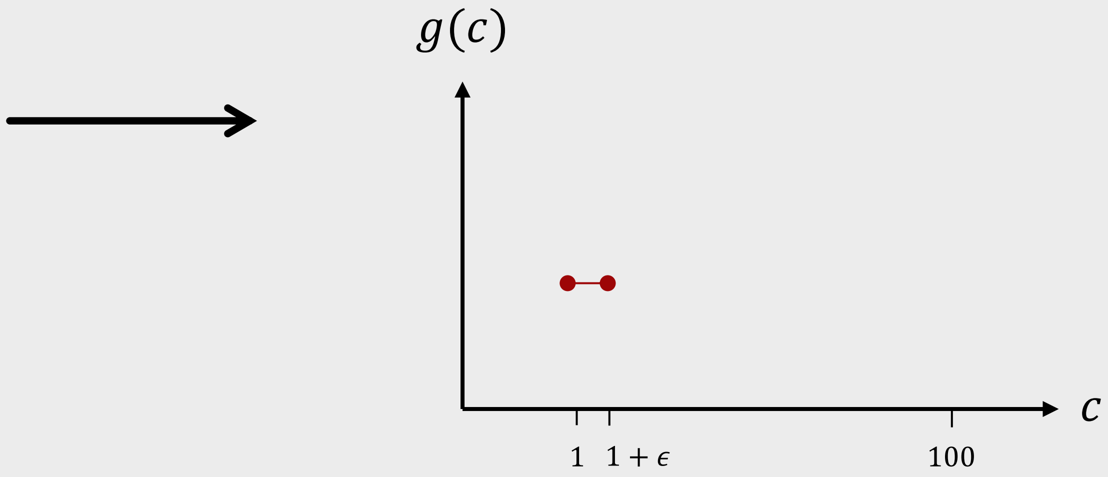

Low Uncertainty

Dütting et al. (2019)

Universal Approximation Guarantee

Theorem [ADLT EC'23]. Let $q \in (0,1)$ and $G\left(c_q\right) = q$. If for $\alpha, \eta \in (0,1)$ the settings is $(\frac{c_q}{\alpha}, \eta)$-small-tail then linear contract α is at least $(\mathbf{1} - \boldsymbol{\alpha})\boldsymbol{\eta}\boldsymbol{q}$ 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 $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\kappa})$-**slowly-increasing** if $G(c) \leq \frac{1}{\beta} G(\alpha c)\ \forall \kappa \leq c$

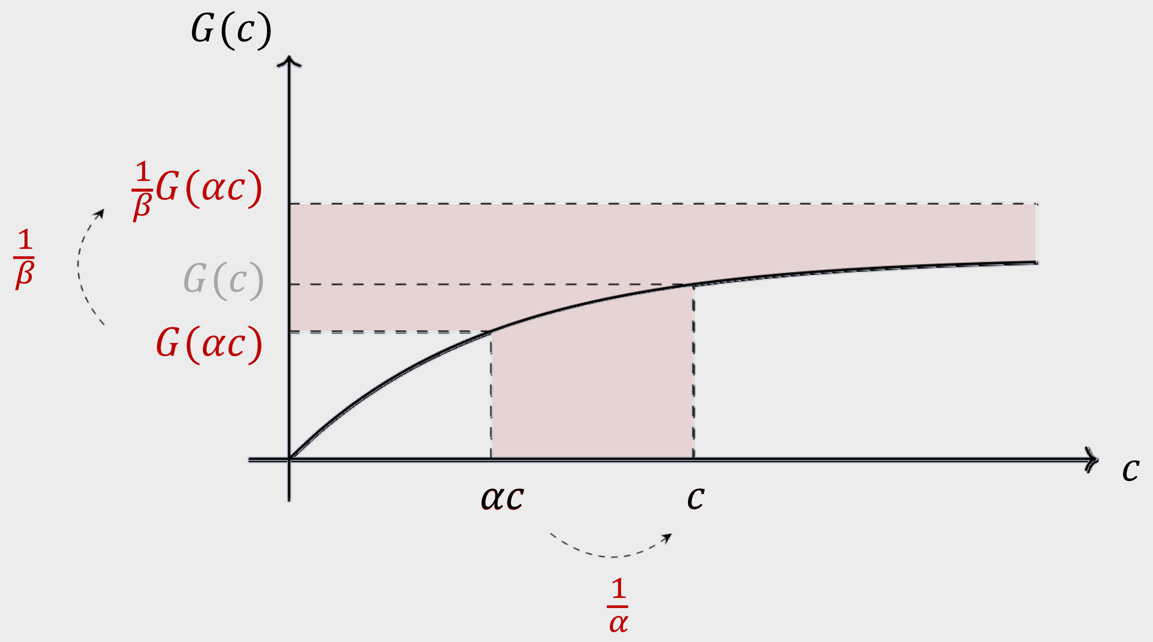

Approximation for Slowly Increasing

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

Under $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\kappa})$**-slowly-increasing** and $(\boldsymbol{\kappa}, \boldsymbol{\eta})$**-small-tail**

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

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

Proof Idea for Slowly Increasing

linear contract revenue $\geq (1-\alpha) \times$ linear contract welfare $\geq (1-\alpha) \times \beta \times$ optimal welfare

Proof Idea for Slowly Increasing

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

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

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

Proof Idea for Slowly Increasing

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

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 $R_i - \gamma_i \frac{c}{\alpha}$

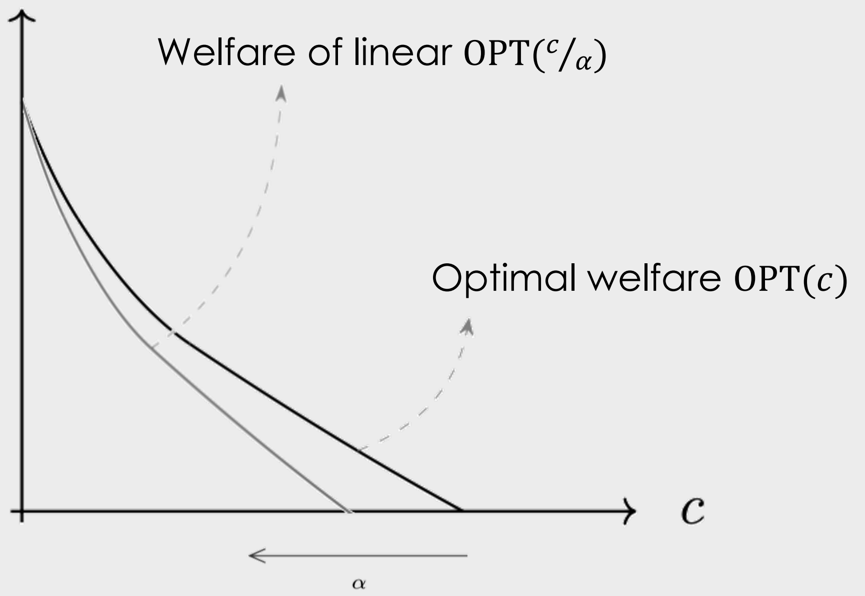

Approximation for Slowly Increasing

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

Under $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\kappa})$**-slowly-increasing** and $(\boldsymbol{\kappa}, \boldsymbol{\eta})$**-small-tail**

linear contract α is $(\mathbf{1} - \boldsymbol{\alpha})\boldsymbol{\beta}\boldsymbol{\eta}$ of the **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**
- **Applications** of this theory

Thank You!

alontal@campus.technion.ac.il

www.talalon.org