Equal-Pay Contracts

Tomek Ponitka
Tel Aviv University

Joint work with



Michal Feldman
Tel Aviv University
Microsoft ILDC

Yoav Gal-Tzur
Tel Aviv University

Maya Schlesinger
Tel Aviv University

Contract Design

manager at paper company

Principal

owns a project and receives its reward.

revenue from selling paper

team of salespeople

Agents

affect the expected reward through their costly actions.

thorough customer research

per-sale commission

Contract

incentivizes agents to take actions by offering them shares of reward.

Model

Known: n agents. Each agent i has set of individual actions A_i.
$f : 2^A \to \mathbb{R}_{\geq 0}$ maps taken actions S to expected reward $f(S)$.

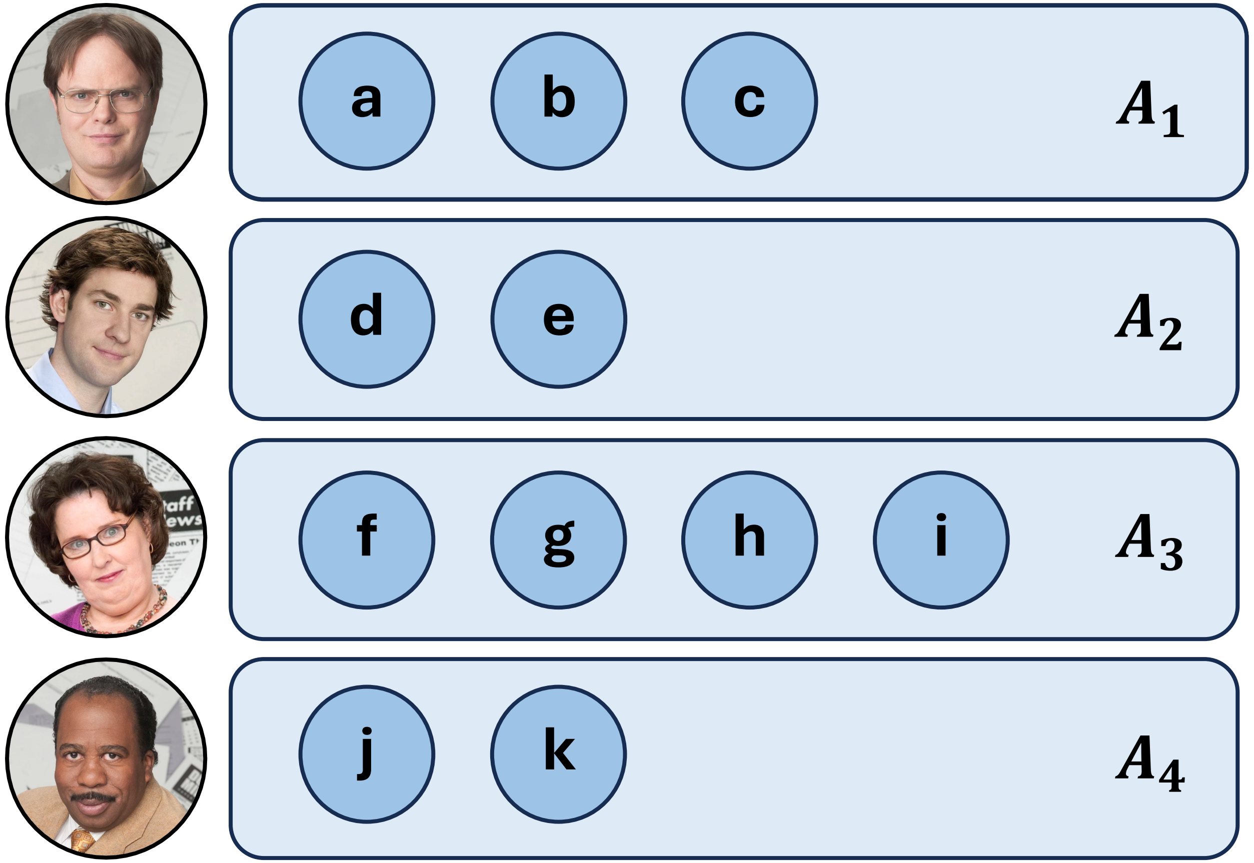

a b c A_1

d e A_2

f g h i A_3

j k A_4

Model

Known: n agents. Each agent i has set of individual actions A_i.
$f : 2^A \to \mathbb{R}_{\geq 0}$ maps taken actions S to expected reward $f(S)$.

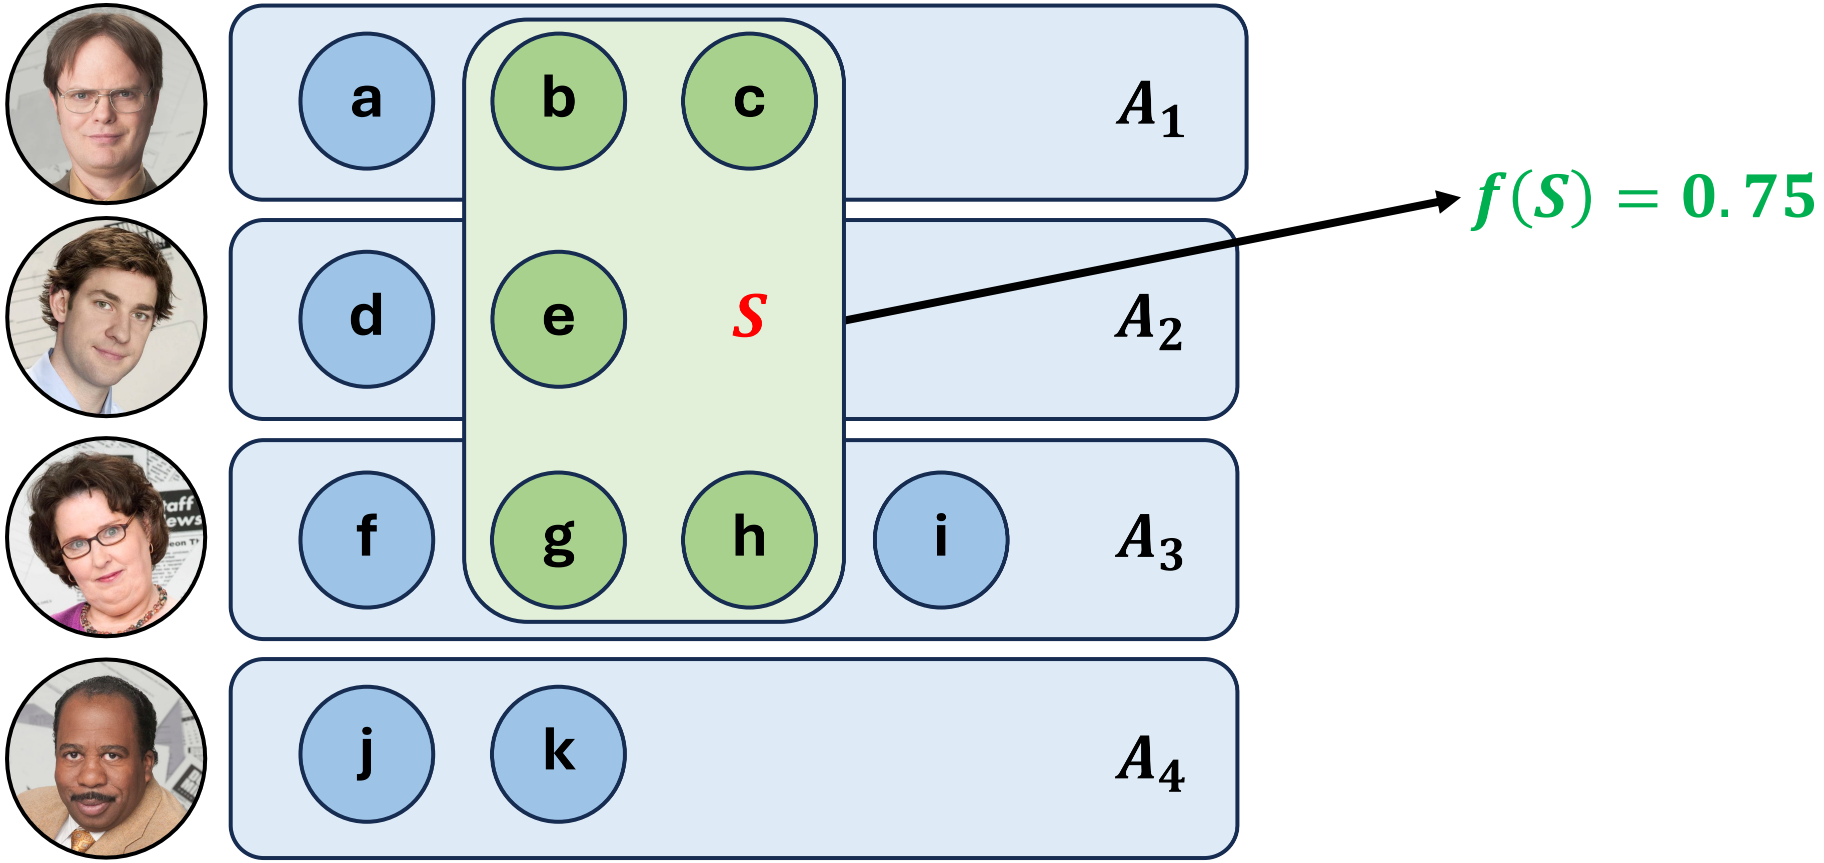

Model

Known: n agents. Each agent i has set of individual actions A_i.
$f : 2^A \to \mathbb{R}_{\geq 0}$ maps taken actions S to expected reward $f(S)$.

Step 1. Principal commits to a contract $\vec{\alpha} = (\alpha_1, \ldots, \alpha_n) \in [0, 1]^n$.
offering agent i an α_i-fraction of the reward.

$\alpha_1 = 0.1$

$\alpha_2 = 0.2$

$\alpha_3 = 0.3$

$\alpha_4 = 0$

Model

Known: n agents. Each agent i has set of individual actions A_i.
$f: 2^A \to \mathbb{R}_{\geq 0}$ maps taken actions S to expected reward $f(S)$.

Step 1. Principal commits to a contract $\vec{\alpha} = (\alpha_1, \dots, \alpha_n) \in [0, 1]^n$.
offering agent i an α_i-fraction of the reward.

Step 2. Each agent i strategically chooses a subset of actions to take.
Agents form a pure Nash equilibrium $S \subseteq A$.

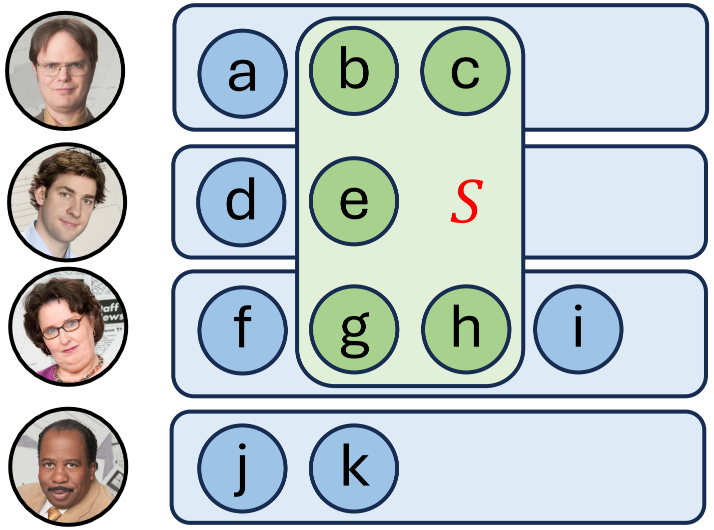

Model

Known: n agents. Each agent i has set of individual actions A_i.
$f : 2^A \to \mathbb{R}_{\geq 0}$ maps taken actions S to expected reward $f(S)$.

Step 1. Principal commits to a contract $\vec{\alpha} = (\alpha_1, \dots, \alpha_n) \in [0, 1]^n$.
offering agent i an α_i-fraction of the reward.

Step 2. Each agent i strategically chooses a subset of actions to take.
Agents form a pure Nash equilibrium $S \subseteq A$.

Agent i's utility:

$$\alpha_i \cdot f(S) - \sum_{a \in S_i} c_a$$

Principal's utility:

$$\left(1 - \sum_{i=1}^{n} \alpha_i\right) \cdot f(S)$$

All About f

The function $f : 2^A \to \mathbb{R}_{\geq 0}$ has exponential representation size.
We access f via either:

Value Queries
Input: $S \subseteq A$
Output: $f(S)$

Demand Queries
Input: $p \in \mathbb{R}^A$
Output: $S \in \operatorname*{argmax}_{S \subseteq A} \left(f(S) - \sum_{a \in S} p_a\right)$

This hierarchy of set functions captures the combinatorial structure of $f : 2^A \to \mathbb{R}_{\geq 0}$.

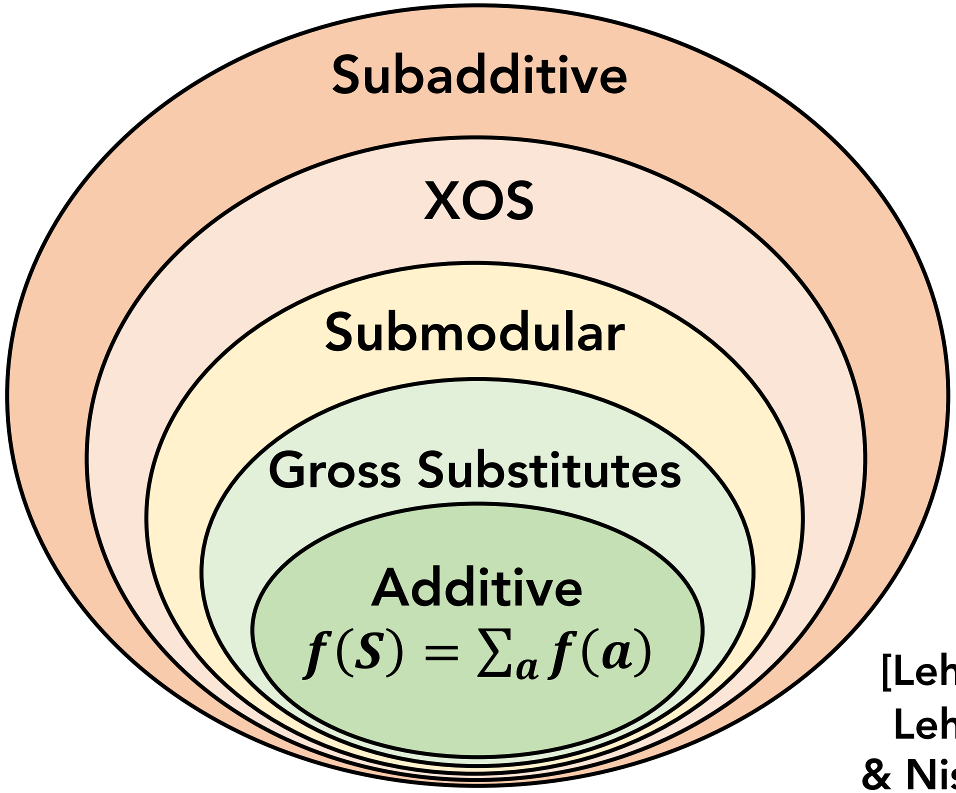

Subadditive

XOS

Submodular

Gross Substitutes

Additive
$f(S) = \sum_a f(a)$

[Lehmann, Lehmann & Nisan '01]

Optimization Problem

maximize **principal's utility** $\left(1 - \sum_{i=1}^{n} \alpha_i\right) \cdot f(S)$ **subject to Nash equilibrium constraints**

Substantial interest in efficient algorithms for finding near-optimal contracts under different structural assumptions on f.

[Babaioff, Feldman, Nisan, EC 2006 & WINE 2006 & SAGT 2009]
[Dütting, Ezra, Feldman, Kesselheim, STOC 2023 & SODA 2025]
[Deo-Campo Vuong, Dughmi, Patel, Prasad, SODA 2024]
[Ezra, Feldman, Schlesinger, ITCS 2024]
[Alon, Castiglioni, Chen, Ezra, Li, Talgam-Cohen, EC 2025]
[Aharoni, Hoefer, Talgam-Cohen, EC 2025]
[Castiglioni, Chen, Li, 2025 & 2025]
[Feldman, Gal-Tzur, Ponitka, Schlesinger, EC 2025 & ITCS 2026]
[Doron-Arad, Shachnai, Shmerler, Talgam-Cohen, EC 2026]
[Dütting, Ezra, Feldman, Kesselheim, EC 2026]
[Ding, Li, Sun, 2026]
[Lavi, Shachnai, Talgam-Cohen, 2026]

Pay Gap

maximize **principal's utility** $(1 - \sum_{i=1}^{n} \alpha_i) \cdot f(S)$ subject to Nash equilibrium constraints

Problem: Optimal contracts might induce a large pay gap.

Discriminatory Payments

$\alpha_1 = 0.5$

$\alpha_2 = 0.05$

$\alpha_3 = 0.005$

$\alpha_4 = 0$

Conflicts with fairness desiderata, institutional rules, legal constraints.

Equal Pay Act of 1963

Knesset News

March 17, 2016

Finance Committee approves bill limiting the pay of executives in financial corporations

In an unprecedented measure, the Knesset Finance Committee unanimously approved Wednesday evening a bill that would lower the maximum salary ceiling for executives in financial corporations from NIS 3.5 million to NIS 2.5 million a year, or 35 times the salary of the lowest-paid worker, whichever is lower.

Equal-Pay Contracts

A contract $\alpha = (\alpha_1, \ldots, \alpha_n)$ is equal-pay
if $\alpha_i = \alpha_j$ for agents i, j with $\alpha_i, \alpha_j \neq 0$.

Equal
Payments

$$\begin{cases} \alpha_1 = 0.2 \\ \alpha_2 = 0.2 \\ \alpha_3 = 0.2 \\ \alpha_4 = 0 \end{cases}$$

All hired agents
receive the same payment.

In the paper: A contract $\alpha = (\alpha_1, \ldots, \alpha_n)$ is γ-equal-pay
if $\alpha_i \geq \gamma \cdot \alpha_j$ for agents i, j with $\alpha_i, \alpha_j \neq 0$.

Research Direction

Question 1 (Computational Aspects):
Can we efficiently compute
nearly-optimal equal-pay contracts?

Question 2 (Price of Equality):
How much utility is lost by restricting contracts to be
equal-pay relative to unconstrained contracts?

Our Contributions

1. We completely characterize the computational complexity of equal-pay contracts.

2. We obtain tight bounds for the price of equality of $\Theta(\log n / \log \log n)$.

3. We resolve two open problems for computational complexity of unconstrained contracts.

Hardness Results

Equal-pay contracts provide a useful lens for studying computational complexity of unconstrained contracts.

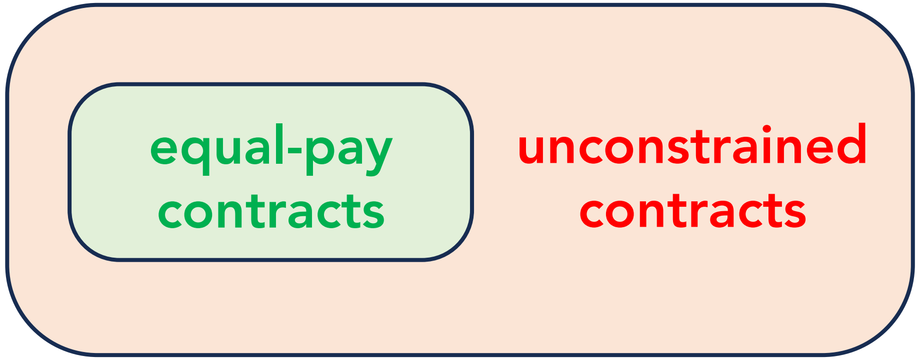

In many hardness results for unconstrained contracts, the optimal contract is equal-pay.

[DEFK ‘23] [EFS ‘24] [DEFK ‘25]

Our analysis of equal-pay contracts provides two new hardness results for unconstrained contracts.

The Computational Landscape*

	Additive	Gross substitutes	Submodular	XOS	Subadditive
Unconstrained	**FPTAS** [FGPS ’26] **No OPT** [DEFK ’23]		**$O(1)$-apx** [DEFK ’25] **No PTAS** [DEFK ’25]		**No $O(\sqrt{n})$** [DEFK ’23]
Equal-Pay					

f

*with demand queries

The Computational Landscape*

f

	Additive	**Gross substitutes**	**Submodular**	**XOS**	**Subadditive**
Unconstrained	**FPTAS** [FGPS '26] **No OPT** [DEFK '23]		$O(1)$**-apx** [DEFK '25] **No PTAS** [DEFK '25]		**No** $O(\sqrt{n})$ [DEFK '23]
Equal-Pay	**OPT** ***NEW***		$O(1)$**-apx** ***NEW*** **No PTAS** [DEFK '25]		**No** $O(\sqrt{n})$ [DEFK '23]

***with demand queries**

The Computational Landscape*

f

	Additive	Gross substitutes	Submodular	XOS	Subadditive
Unconstrained	**FPTAS** [FGPS ’26] **No OPT** [DEFK ’23]	**$O(1)$-apx** **No PTAS** ***NEW***	**$O(1)$-apx** [DEFK ’25] **No PTAS** [DEFK ’25]	**No $O(1)$** ***NEW***	**No $O(\sqrt{n})$** [DEFK ’23]
Equal-Pay	**OPT** ***NEW***	**$O(1)$-apx** **No PTAS**	**$O(1)$-apx** ***NEW*** **No PTAS** [DEFK ’25]	**No $O(1)$**	**No $O(\sqrt{n})$** [DEFK ’23]

***with demand queries**

Our Techniques

Positive Result for Submodular:
Approximate demand queries with agent cardinality constraints

Negative Result for XOS:
Hide a special set

Negative Result for Gross Substitutes:
Encoding a coverage structure over the agents

Our Techniques

Theorem: For submodular f, an $O(1)$-approximation to the optimal equal-pay contract can be computed in poly time with demand queries.

Demand Queries
Input: $p \in \mathbb{R}^A$
Output: $S \in \underset{S \subseteq A}{\operatorname{argmax}} \left(f(S) - \sum_{a \in S} p_a \right)$

Demand queries are agent-agnostic: They ignore ownership of actions.

Payments are generally lower when fewer agents control the actions.

Lemma: Demand queries with agent cardinality constraints can be approximately simulated by polynomially many value queries.

We run the distorted greedy algorithm [HFWK '19] recursively on two levels: at the agent level, and, for each agent, at the action level.

Our Techniques

Information-Theoretic Hardness: For XOS f, no algorithm making polynomially many value or demand queries achieves an O(1)-approx.

(a) The hidden set is required to find a good contract.
(b) Since demand queries are agent-agnostic, identifying the hidden set requires exponentially many such queries.

hidden set of agents T

Our Techniques

Computational Hardness: For gross substitutes f,
there is no PTAS unless P = NP.

Reduction from a promise variant of maximum coverage [EFS '24]

Matroid structure over actions
Coverage structure over agents

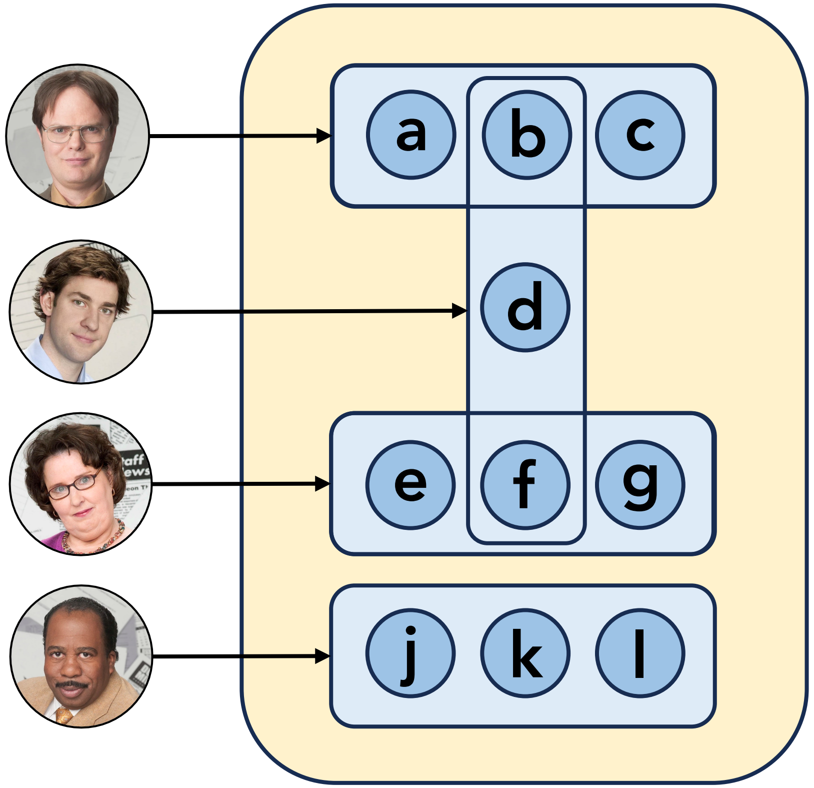

Research Direction

Question 1 (Computational Aspects):
Can we efficiently compute
nearly-optimal equal-pay contracts?

Question 2 (Price of Equality):
How much utility is lost by restricting contracts to be
equal-pay relative to unconstrained contracts?

Price of Equality

Theorem: For XOS f, the optimal equal-pay contract gives an $\Theta(\log n \,/\, \log \log n)$-approx. to the optimal unconstrained contract.

We sort the agents by decreasing $\alpha_i^\star$ and apply a bucketing argument.

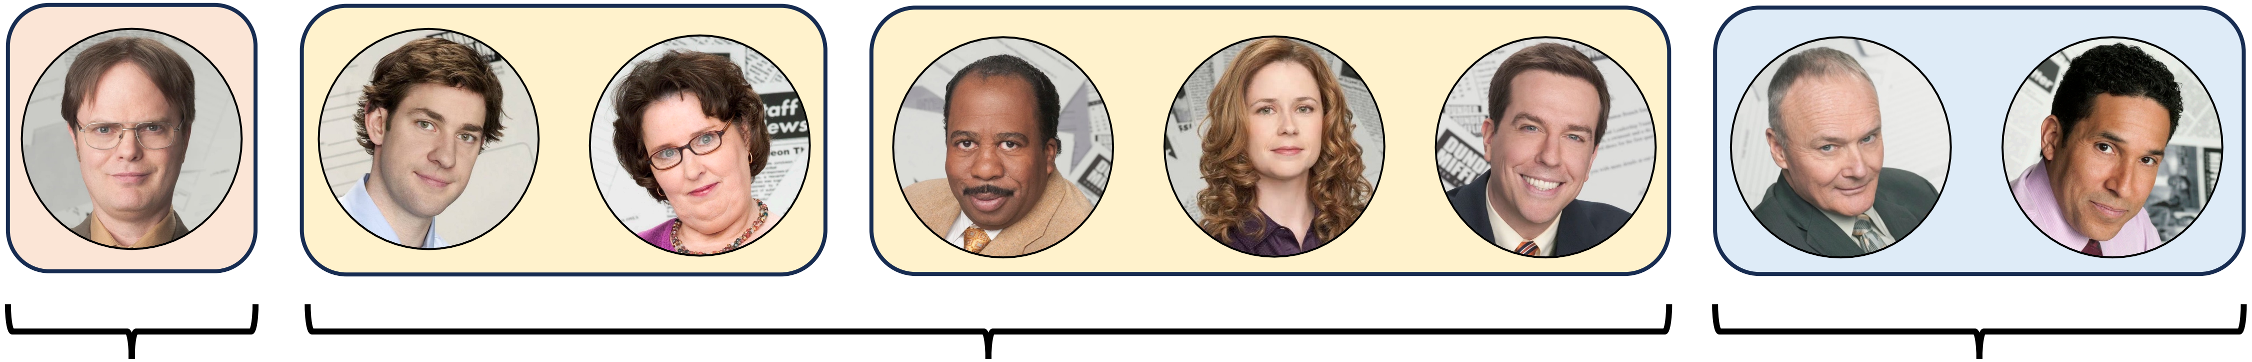

Large Agent

$\alpha_i^\star > 1/2$

1 bucket

Middle Agents

$\Theta(\log n \,/\, \log \log n)$ buckets

Small Agents

$\alpha_i^\star < 1/(2n)$

1 bucket

Price of Equality

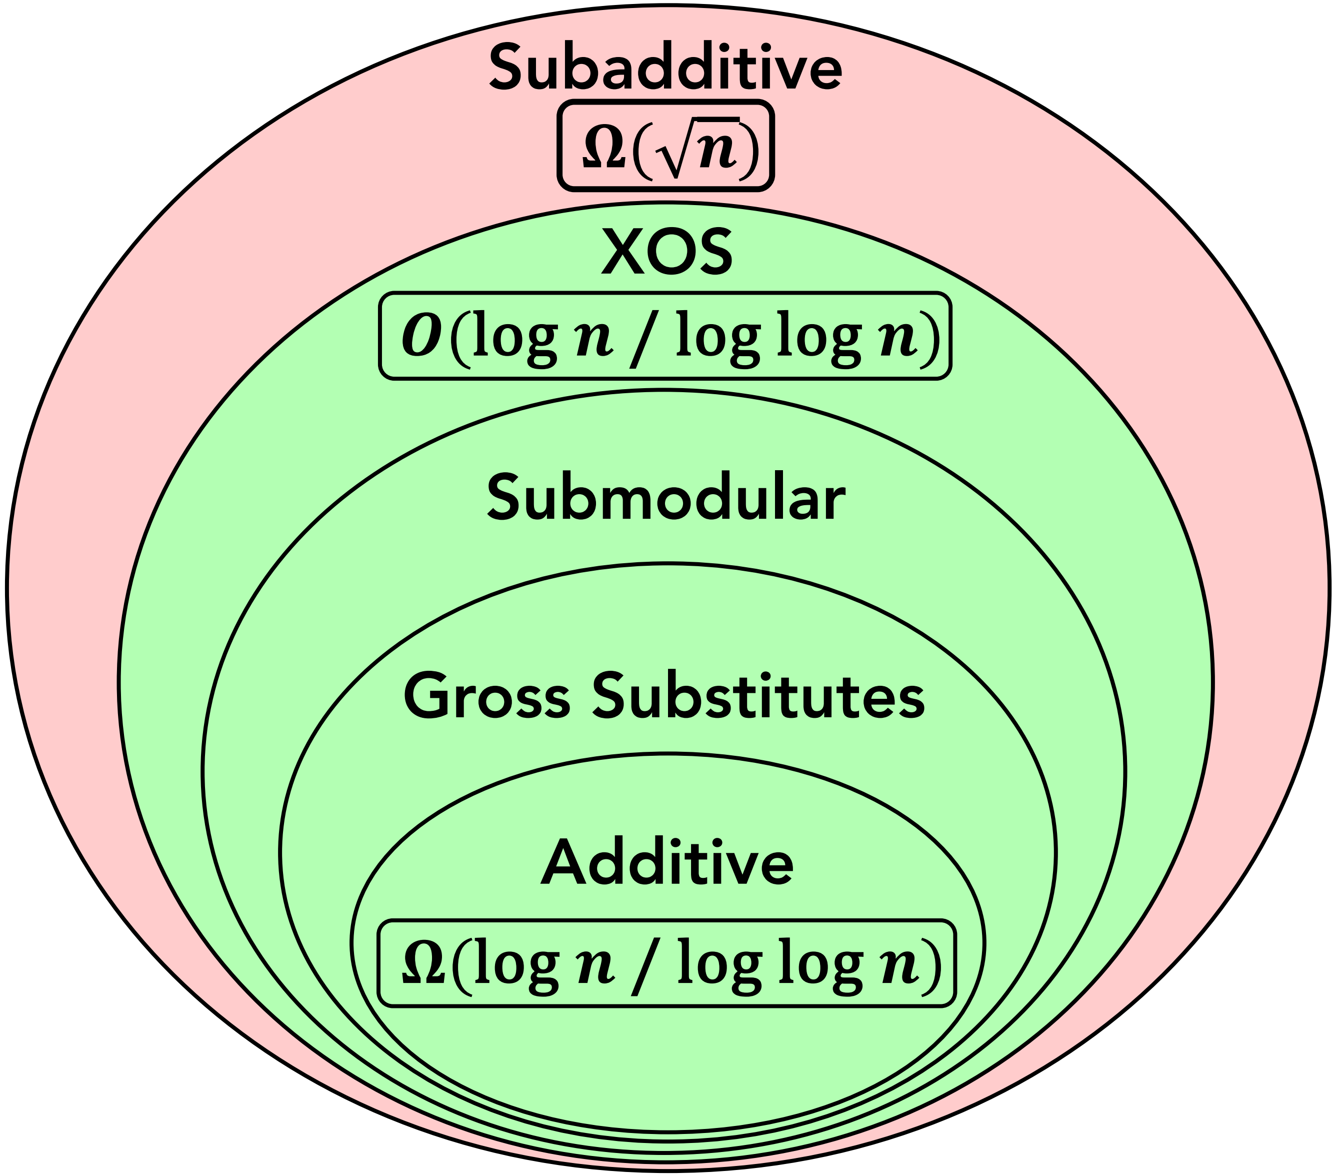

Subadditive

$\boxed{\Omega(\sqrt{n})}$

XOS

$\boxed{O(\log n \,/\, \log\log n)}$

Submodular

Gross Substitutes

Additive

$\boxed{\Omega(\log n \,/\, \log\log n)}$

Other Works on Fair Contracts

Fairness is an **emerging frontier** in algorithmic contract design.

Envy-Free Contracts
[Castiglioni, Chen, Li, 2025]

Nearly-Equal-Pay Contracts
[Ding, Li, Sun, 2026]

Maximizing Social Welfare
[Aharoni, Hoefer, Talgam-Cohen, 2025]
[Feldman, Gal-Tzur, Ponitka, Schlesinger, 2025]

Price of Non-Discrimination in Public Combinatorial Contracts [Feng, Ma, Xiao, 2024]
Fair Contracts [Castiglioni, Chen, Li, 2025]
Envy-Free Contracts with Subsidies [Castiglioni, Chen, Li, 2026]
Anonymous Contracts [Brustle, Dütting, Leonardi, Ponitka, Russo, 2026]

Envy-Free Contracts

[Castiglioni, Chen, Li, 2025] define envy-free contracts.

Equal-Pay $\xrightarrow[\boldsymbol{\Theta(1)}\text{-approximation}]{\subseteq}$ **Envy-Free**

under binary actions and submodular f [CCL '25]

Open Problem:

Bounding the gap for combinatorial actions and/or XOS f.

Nearly-Equal-Pay Contracts

A contract $\alpha = (\alpha_1, \dots, \alpha_n)$ is γ-equal-pay
if $\alpha_i \geq \gamma \cdot \alpha_j$ for agents i, j with $\alpha_i, \alpha_j \neq 0$.

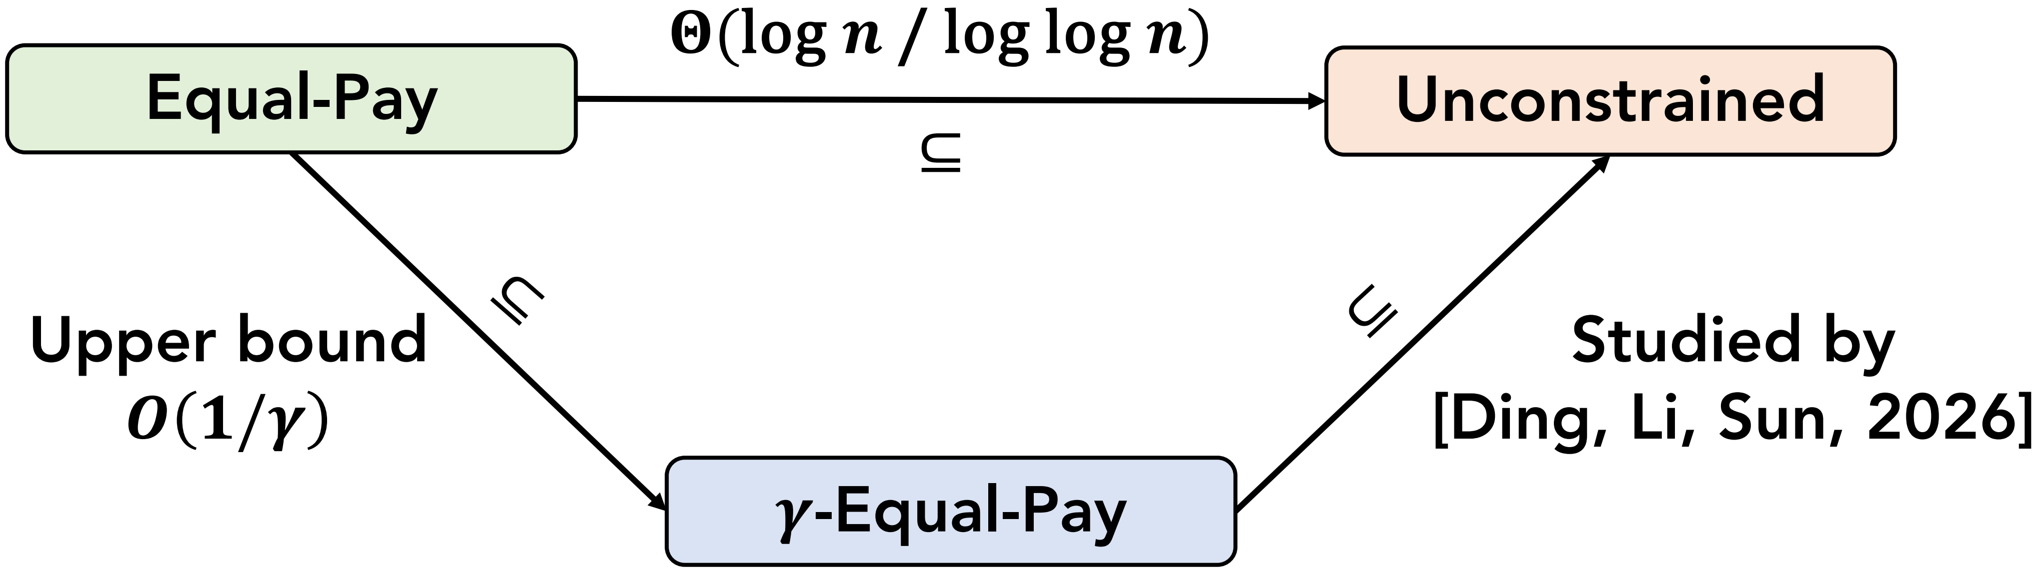

$\Theta(\log n \,/\, \log \log n)$

Equal-Pay

$\subseteq$

Unconstrained

$\subseteq$

Upper bound
$O(1/\gamma)$

$\subseteq$

Studied by
[Ding, Li, Sun, 2026]

γ-Equal-Pay

Open Problem:
Better bounds for the equal-pay vs γ-equal-pay gap.

BEST (BEyond STandard) Objectives

For gross substitutes f, our computational and price-of-equality results for maximizing principal's utility extend to all BEST objectives.

BEST Objectives [FGPS '25]

Principal's Utility (Second Best): $\left(1-\sum_{i=1}^{n}\alpha_i\right)\cdot f(S)$

Expected Reward: $f(S)$

Social Welfare (First Best): $f(S)-c(S)$

Separation: Price of equality for submodular f and maximizing reward is ∞.

Open Problem: Price of equality for submodular f and maximizing welfare.

Takeaways

Equal-Pay Contracts:
(a) motivated by fairness.
(b) useful lens for understanding unconstrained contracts.

The price of equality is $\boldsymbol{\Theta}(\log \boldsymbol{n} \,/\, \log\log \boldsymbol{n})$.

Fairness is an emerging frontier in algorithmic contract design.

Thank you!