Simple Mechanisms for Agents with Non-linear Utility

Yiding Feng Jason D. Hartline Yingkai Li

Northwestern University

One item, n agents with independent non-identically distributed private type t. Allocation $x \in \{0, 1\}$, payment $p \in \mathbb{R}_+$. An agent has the utility function $u(t, x, p)$, i.e., a mapping from her private type and the outcome to her von Neumann-Morgenstern utility.

Specific Utility Models

- *linear utility:* private type $t = v$, i.e., the value of the item; utility $u(v, x, p) = vx - p$.
- *public-budget utility:* utility u encodes a public non-identical budget w; private type $t = v$; utility $u_w(v, x, p) = vx - p$ if $p \leq w$ and $-\infty$ otherwise.
- *private-budget utility:* private type $t = (v, w)$; utility $u(v, w, x, p) = vx - p$ if $p \leq w$ and $-\infty$ otherwise.
- *capacitated utility:* private value $t = v$ of the item; non-identical capacity C encoded in utility function; utility $u_C(v, x, p) = \min\{vx - p, C\}$.

Payoff Curves

Payoff Curve: a mapping from a constraint q on the ex ante probability of sale, over randomness in the agent's type and the mechanism, to the payoff (e.g. revenue, welfare) of the optimal mechanism with the ex ante constraint, within a fixed class of mechanisms. Specifically,

Definition. The *price-posting payoff curve* P is the payoff curve generated by the price-posting mechanisms. The *ironed price-posting payoff curve* $\bar{P}$ is the concave hull of P.

Definition. The *optimal payoff curve* R is the payoff curve generated by all possible mechanisms.

Ex Ante Relaxation

- *Optimal Auction:* optimal mechanism that maps types to allocations and payments subject to feasibility constraint; sophisticated; involves competition and discrimination;
- *Ex Ante Relaxation (EAR):* relax feasibility (selling at most one item) to ex ante feasibility (selling at most one item ex ante); upper bounds the payoff of optimal mechanism.

Fact. The optimal payoff curves $\mathbf{R} = \{R_i\}_{i=1}^n$ uniquely determine the payoff of EAR, i.e.,

$$\text{EAR}(\mathbf{R}) = \max_{\mathbf{q}:\|\mathbf{q}\|_1 \leq 1} \sum_{i=1}^{n} R_i(q_i).$$

Reduction Framework for Pricing-based Mechanisms

Myerson (1981) shows that every mechanism for agents with linear utilities is a pricing-based mechanism.

Definition. A mechanism is a *pricing-based mechanism* if its payoff is determined by the price-posting payoff curves $\mathbf{P} = \{P_i\}_{i=1}^n$ of agents.

For non-linear agents, however, mechanisms (e.g., revenue-optimal mechanism) are not uniquely pinned down by the pricing-posting payoff curves in general even for a single-agent setting. The ζ-*closeness* (defined below) of an agent measures how close her ironed price-posting payoff curve is to her optimal payoff curve.

Definition. An agent's ironed price-posting payoff curve $\bar{P}$ is ζ-*close* to her optimal payoff curve R, if for all $q \in [0, 1]$, there exists a quantile $q^\dagger \leq q$ such that $\zeta \cdot \bar{P}(q^\dagger) \geq R(q)$. Such an agent is ζ-*close*.

Based on the definition of ζ-closeness, we present the main result: a reduction framework that converts mechanisms for agents with linear utilities to agents with non-linear utilities, and approximately preserves its payoff approximation guarantee.

Definition. Fix any set $\mathcal{A}$ of (non-linear) agents with price-posting payoff curves $\mathbf{P}$. The *linear agents analogy* $\mathcal{A}_L$ is an set of linear agents whose price-posting payoff curves are also $\mathbf{P}$, and optimal payoff curves are $\bar{\mathbf{P}}$.

Main Theorem. Fix any set $\mathcal{A}$ of (non-linear) agents with price-posting payoff curves $\mathbf{P}$ and optimal payoff curves $\mathbf{R}$. For any DSIC, IIR, deterministic mechanism $\mathcal{M}_L$ for agents with linear utility, there is a pricing-based mechanism $\mathcal{M}$ for non-linear agents $\mathcal{A}$ that is DSIC, IIR, and satisfies

1. Identical payoff: mechanism $\mathcal{M}$ for non-linear agents $\mathcal{A}$ has the same payoff as mechanism $\mathcal{M}_L$ for the linear agents analog $\mathcal{A}_L$. Denote the payoff of mechanism $\mathcal{M}_L$ and $\mathcal{M}$ as $\mathcal{M}(\mathbf{P})$ and$\mathcal{M}_L(\mathbf{P})$.
2. Identical feasibility: mechanism $\mathcal{M}$ for non-linear agents $\mathcal{A}$ has the same distribution over outcomes as mechanism $\mathcal{M}_L$ for the linear agents analog $\mathcal{A}_L$.

Denote by γ the approximation of mechanism $\mathcal{M}_L$ for the linear agents analog $\mathcal{A}_L$ to the ex ante relaxation of $\mathcal{A}_L$, i.e., $\mathcal{M}_L(\mathbf{P}) \geq {}^1\!/\!_\zeta \cdot \text{EAR}(\bar{\mathbf{P}})$. If each non-linear agent in $\mathcal{A}$ is ζ-close, then mechanism $\mathcal{M}$ for non-linear agents $\mathcal{A}$ is $\gamma\,\zeta$-approximation to the ex ante relaxation of $\mathcal{A}$, i.e., $\mathcal{M}(\mathbf{P}) \geq {}^1\!/\!_{\gamma\zeta} \cdot \text{EAR}(\mathbf{R})$.

Simple Mechanisms for Agents with Linear Utility

Here we list some simple mechanism for agents with linear utility and their approximation (w.r.t. ex ante relaxation) guarantee.

- *Sequential Posted Pricing*: agents in sequence (specified by mechanisms) are offered take-it-or-leave-it prices. It is a $e/(e-1)$-approximation to ex ante relaxation for linear agents (Chawla et al., 2010).
- *Oblivious Posted Pricing*: agents in sequence (unknown to mechanisms in advance) are offered take-it-or-leave-it prices. It is a 2-approximation to ex ante relaxation for linear agents (Chawla et al., 2010).
- *Marginal Payoff Maximization (a.k.a., Myersonian Auction)*: Marginal revenue mechanism in Bulow and Roberts (1989) is a $e/(e-1)$-approximation to ex ante relaxation for linear agents (Alaei et al., 2013).

ζ-Closeness for Non-linear Utility

Table: Summary of results for ζ-closeness under various assumptions.

	public budget		independent private budget			capacitated utility	
	revenue		revenue		welfare	revenue	welfare
assumption	regular		regular† value	MHR budget‡		regular, support $[0, \bar{v}]$, capacity at least $\bar{v}/\eta$	
ζ-closeness	1*	2*	3	$1 + 3e - {}^1\!/\!_e$	2	$2 + \ln \eta$	1*

$*$ indicates tight ratio. $\dagger$ distribution F is regular if $v - \frac{1-F(v)}{f(v)}$ is non-decreasing in v. $\ddagger$ distribution F is MHR if the hazard rate $\frac{f(v)}{1-F(v)}$ is non-decreasing in v.

Implementation of Main Theorem

For any deterministic DSIC, IIR mechanism $\mathcal{M}_L$ for linear agents, it can be represented by a mapping from the quantiles of other agents to a threshold quantile for each agent. The agent wins when her quantile is below the threshold and loses when her quantile is above the threshold. Denote the function that maps the profile of other agent quantiles $\{q_j\}_{j\in N\setminus\{i\}}$ to a quantile threshold for agent i by $\hat{q}_i^{\mathcal{M}_L}\left(\{q_j\}_{j\in N\setminus\{i\}}\right)$.

For any agent with non-linear utility, the single-agent pricing problem identifies the per-unit (market clearing) price p^q to offer the agent for any ex ante allocation constraint q. Denote the allocation probability selected by an agent with type t as $x^q(t)$ when offered per-unit price p^q. Under mild assumption, $x^q(t)$ is non-decreasing in quantile q for all type t.

Implementation.

- **Input:** A set $\mathcal{A}$ of (non-linear) agents; and deterministic, DSIC, IIR mechanism $\mathcal{M}_L$ for linear agents.
- For each agent i with private type t_i, map the type to a random quantile q_i according to the distribution H_i with cdf $H_i(q) = x_i^q(t_i)$.
- For each agent i, calculate quantile threshold as $\hat{q}_i = \hat{q}_i^{\mathcal{M}_L}\left(\{q_j\}_{j\in N\setminus\{i\}}\right)$.
- For each agent i, set payment $p_i = p^{\hat{q}_i} x_i^{\hat{q}_i}(t_i)$, and allocation $x_i = 1$ if $q_i < \hat{q}_i$ and $x_i = 0$ otherwise.

Extension

The reduction framework can be generalized from single-item environments to any downward-closed environments, e.g., multi-unit, matroid.

Reference

Alaei, S., Fu, H., Haghpanah, N., and Hartline, J. (2013). The simple economics of approximately optimal auctions. In *Proc. 54th IEEE Symp. on Foundations of Computer Science*.

Bulow, J. and Roberts, J. (1989). The simple economics of optimal auctions. *The Journal of Political Economy*, 97:1060–90.

Chawla, S., Hartline, J., Malec, D., and Sivan, B. (2010). Sequential posted pricing and multi-parameter mechanism design. In *Proc. 41st ACM Symp. on Theory of Computing*.

Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6:58–73.