CS 405/6001: Game Theory and Algorithmic Mechanism Design **Jul–Nov 2026**

Lecture 21: Affine Maximizers and Single Object Allocation

Lecturer: Swaprava Nath *Scribe(s): Claude Scribe*

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor.*

21.1 Recap and Motivation

Last week we met the VCG mechanism. It is truthful, and it is elegant. But it has three practical limitations.

- It is not budget balanced.
- It treats all agents symmetrically (anonymously). This may not be what the designer wants.
- Its output is always the efficient allocation. The designer may have other goals.

To address these concerns, we generalize VCG. We give agents different **weights**, and we give outcomes a **translation** (a bonus). This lets the designer give different importance to different agents, prefer some allocations over others, and possibly do better on budget balance.

21.2 Affine Maximizers

Definition 21.1 (Affine Maximizer) *An allocation rule f is an **affine maximizer (AM)** if there exist (potentially different) weights $w_i \geqslant 0$, $\forall i \in N$, not all zero, and a function $\kappa : A \to \mathbb{R}$ such that*

$$f^{AM}(\theta) \in \arg\max_{a \in A} \Big(\sum_{i \in N} w_i \theta_i(a) + \kappa(a) \Big).$$

Special cases. The family contains several rules we already know.

1. **Efficient (VCG) allocation:** $\kappa \equiv 0$ and $w_i = 1$ for all $i \in N$. This maximizes social welfare.
2. **Dictatorial allocation:** $\kappa \equiv 0$, $w_d = 1$ for a dictator d, and $w_i = 0$ for all $i \neq d$. The rule picks d's most preferred allocation.
3. **Weighted welfare:** $\kappa \equiv 0$, but the w_i differ. Some agents count more than others. Such a rule is not anonymous.
4. **Allocation-dependent preferences:** κ is non-constant. Different allocations get an intrinsic value, independent of the agents' valuations.

Example 21.2 (The surge policy vote) *A ride-hailing app picks one city-wide pricing policy from $A = \{Surge, Flat, Hybrid\}$. Agent 1 is the drivers' union and agent 2 is the riders' forum. A city regulator likes the Hybrid policy.*

	Surge	*Flat*	*Hybrid*
θ_1 *(drivers)*	*6*	*1*	*3*
θ_2 *(riders)*	*1*	*4*	*2*
κ *(regulator)*	*0*	*0*	*3*

With $w_1 = w_2 = 1$ and $\kappa \equiv 0$, the scores are $7, 5, 5$. Surge wins; this is the efficient choice. With $w_2 = 2$ and $\kappa \equiv 0$, the scores are $8, 9, 7$. Flat wins. Swapping the names of the two agents would now change the outcome, so the rule is not anonymous. With $w_2 = 2$ and the regulator's bonus $\kappa(\text{Hybrid}) = 3$, the scores are $8, 9, 10$. Hybrid wins.

Properties.

- Affine maximizers with different weights are generally not anonymous: permuting the agents can change the outcome.
- They are a strict superclass of VCG (efficient) allocations, so they can satisfy properties that VCG cannot.

Like the Gibbard–Satterthwaite question in voting, we can now ask a characterization question in the quasi-linear setting. In the GS theorem, preferences could be any total ordering. Here, valuations can take any arbitrary real values. Which allocation rules are implementable in such a domain? Before the answer, we need one tie-breaking condition.

Definition 21.3 (Independence of Non-influential Agents) *An AM rule f^{AM} with weights w_i, $i \in N$, and function κ satisfies **independence of non-influential agents (INA)** if for all $i \in N$ with $w_i = 0$,*

$$f^{AM}(\theta_i, \theta_{-i}) = f^{AM}(\theta_i', \theta_{-i}), \quad \forall\, \theta_i, \theta_i', \theta_{-i}.$$

INA is a tie-breaking requirement. If the affine maximization returns a set of allocations, ties must be broken consistently, and never by an agent whose report does not enter the objective. A zero-weight agent does not influence the decision, so it should not break any tie either.

Why INA is needed. Suppose $w_i = 0$, and allocations a and b tie on the affine objective. Agent i prefers a, but the tie is broken in favor of b using i's report. Then i can misreport its type to swing the tie toward a. More formally, if INA is violated, there is a profile (θ_i, θ_{-i}) where several allocations maximize the affine sum, the chosen allocation depends on θ_i although $w_i = 0$, and agent i gains by reporting some θ_i'. This violates dominant strategy incentive compatibility.

Example 21.4 (The silent auditor) *The ride-hailing app adds agent 3, a city auditor bot with $w_3 = 0$. At some profile, Flat and Hybrid tie on the weighted sum. Suppose the app breaks the tie by dropping the alternative that the auditor reports as less preferred. Then the auditor has a lever: by misreporting which of the two it likes less, it pulls the outcome toward the policy it actually prefers. A zero-weight agent should never hold such a lever. INA forbids exactly this.*

Theorem 21.5 *An AM rule satisfying INA is implementable in dominant strategies.*

Proof: We construct a payment p^{AM} that makes (f^{AM}, p^{AM}) DSIC. Let h_i be any function of θ_{-i} only, and set

$$p_i^{AM}(\theta_i, \theta_{-i}) = \begin{cases} \dfrac{1}{w_i}\Big[h_i(\theta_{-i}) - \Big\{\sum_{j \neq i} w_j \theta_j\big(f^{AM}(\theta)\big) + \kappa\big(f^{AM}(\theta)\big)\Big\}\Big], & \forall i : w_i > 0, \\ 0, & \forall i : w_i = 0. \end{cases} \tag{21.1}$$

For agent i with $w_i > 0$, the payoff at the true type θ_i is

$$\begin{aligned}\theta_i(f^{AM}(\theta_i,\theta_{-i})) - p_i^{AM}(\theta_i,\theta_{-i}) &= \frac{1}{w_i}\Big[\Big\{\sum_{j\in N} w_j\theta_j(f^{AM}(\theta_i,\theta_{-i})) + \kappa(f^{AM}(\theta_i,\theta_{-i}))\Big\} - h_i(\theta_{-i})\Big] \\ &\geqslant \frac{1}{w_i}\Big[\Big\{\sum_{j\in N} w_j\theta_j(f^{AM}(\theta_i',\theta_{-i})) + \kappa(f^{AM}(\theta_i',\theta_{-i}))\Big\} - h_i(\theta_{-i})\Big] \\ &= \theta_i(f^{AM}(\theta_i',\theta_{-i})) - p_i^{AM}(\theta_i',\theta_{-i}).\end{aligned}$$

The inequality holds because $f^{AM}(\theta_i,\theta_{-i})$ maximizes the affine objective at the true profile, and $w_i > 0$. So every agent with $w_i > 0$ is truthful.

For an agent with $w_i = 0$, INA gives $f^{AM}(\theta_i,\theta_{-i}) = f^{AM}(\theta_i',\theta_{-i})$ for all $\theta_i,\theta_i',\theta_{-i}$, and the payment is zero. Hence the payoff is identical for all reports, and the agent is weakly truthful. ■

Remark 21.6 *With $w_i = 1$ for all i and $\kappa \equiv 0$, payment (21.1) is exactly the Groves payment. The idea is the same as for Groves: the payment makes agent i's payoff equal to the whole AM objective, scaled by $1/w_i$, minus a term i cannot affect. Truthful reports make the mechanism pick the outcome with the highest objective. A lie can only make it pick a different, weakly lower, one.*

21.3 Roberts' Theorem

What if the valuations are **unrestricted**? An unrestricted valuation is any mapping $\theta_i : A \to \mathbb{R}$, with no further constraint. The answer is the following theorem, which we state without proof.

Theorem 21.7 (Roberts 1979) *Let A be a finite set of alternatives with $|A| \geqslant 3$. Suppose the type space of each agent i contains all valuation functions $\theta_i : A \to \mathbb{R}$. Then every social choice rule $f : \Theta_1 \times \cdots \times \Theta_n \to A$ that is onto and implementable in dominant strategies must be an affine maximizer. That is, there exist weights $w_i \geqslant 0$ and a function $\kappa : A \to \mathbb{R}$ such that*

$$f(\theta) \in \arg\max_{a\in A}\Big\{\sum_{i=1}^{n} w_i\theta_i(a) + \kappa(a)\Big\}, \quad \forall \theta \in \Theta_1 \times \cdots \times \Theta_n.$$

Comparison with the Gibbard–Satterthwaite theorem. Roberts' theorem is the analog of the GS theorem in the quasi-linear domain.

- **GS theorem:** with unrestricted ordinal preferences, only dictatorial social choice functions are strategyproof and onto.
- **Roberts' theorem:** with unrestricted cardinal valuations, only affine maximizers are implementable and onto.

Both theorems say that without domain restrictions, the class of truthful mechanisms is limited. But affine maximizers are richer than dictatorships: they include efficient allocations (VCG), weighted welfare maximization, and allocations with intrinsic preferences over outcomes (through κ).

In short, money buys a richer class. Without money, with unrestricted preferences and $|A| \geqslant 3$, onto and DSIC means dictatorship. With money (quasi-linear utilities) and unrestricted valuations, onto and DSIC

means affine maximizer, a class that contains dictatorships, VCG, weighted welfare, and any bonus κ on outcomes.

We do not want to be restricted to exactly one class of rules; also, unrestricted valuations hardly occur in practice. The way out is again to restrict the domain. In voting, single-peaked preferences gave us the median voter rule. Here, restricting the valuations (for example, to auctions) gives truthful mechanisms beyond VCG and Groves.

Proof reference. The proof of Roberts' theorem is beyond the scope of this course. Interested readers can refer to Lavi, Mu'alem and Nisan [LMN09].

21.4 Single Object Allocation

Let the type set of agent i be $T_i \subseteq \mathbb{R}$, and let $t_i \in T_i$ be the value agent i gets if she wins the object. An allocation a is a vector of probabilities: a_i is the probability that agent i wins, and a_0 is the probability that the object is not sold. The set of allocations is

$$\Delta A = \Big\{ a \in [0,1]^{n+1} : \sum_{i=0}^{n} a_i = 1 \Big\}.$$

An allocation rule is a function $f : T_1 \times T_2 \times \cdots \times T_n \to \Delta A$. Here $f_i(t_i, t_{-i})$ is the probability that i wins the object at type profile (t_i, t_{-i}), and $f_0(t)$ is the probability of not selling. Given an allocation a, the (expected) valuation of agent i has the **product form**

$$v_i(a, t_i) = a_i \cdot t_i.$$

Example 21.8 (The GPU-hour auction) *A cloud provider sells one GPU hour. Each lab has a single number: what the hour is worth to it. The provider may even randomize, for example by giving the hour to Lab A with probability* 0.7*. Then Lab A's expected value is* $0.7\, t_A$*.*

The setting extends to cases where a_i is the *amount* of the object given to i. For example, in sponsored search auctions, a_i is replaced by CTR_i. In general, the analysis carries over whenever every agent's valuation is a product of one factor coming from the allocation and a scalar type of the agent. This domain is called the **single parameter domain**.

21.4.1 The Vickrey (second price) auction

Define $t_{-i}^{(2)} = \max_{j \neq i} t_j$, the highest bid among the others. Agent i wins if $t_i > t_{-i}^{(2)}$, loses if $t_i < t_{-i}^{(2)}$, and a tie-breaking rule decides at equality. The winner pays $t_{-i}^{(2)}$, so the utility in a tie is zero:

$$u_i(t_i, t_{-i}) = \begin{cases} 0 & \text{if } t_i \leqslant t_{-i}^{(2)}, \\ t_i - t_{-i}^{(2)} & \text{if } t_i > t_{-i}^{(2)}. \end{cases}$$

Figure 21.1 shows two observations. First, the utility is a **convex** function of t_i. Its derivative is 0 for $t_i < t_{-i}^{(2)}$ and 1 for $t_i > t_{-i}^{(2)}$, and it is not differentiable at $t_i = t_{-i}^{(2)}$. Second, wherever the derivative exists, it coincides with the probability of winning. Slope of utility equals win probability: this is the key idea behind Myerson's lemma in the next lecture. To make it precise, we need some convex analysis.

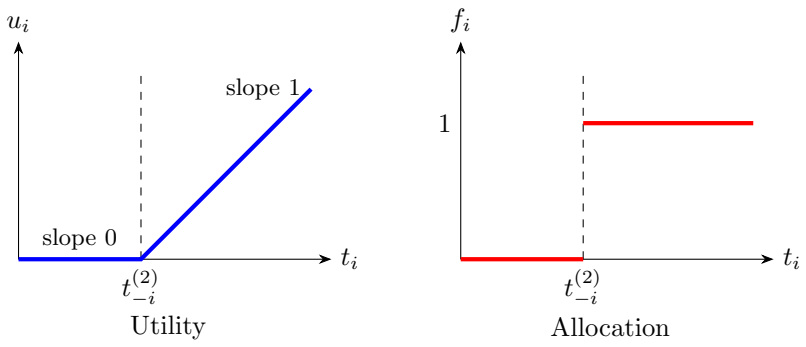


Figure 21.1: Utility and allocation of agent i in the second price auction, with the other bids fixed.

21.5 Some Results from Convex Analysis

We study functions $g : I \to \mathbb{R}$, where $I \subseteq \mathbb{R}$ is an interval.

Definition 21.9 *A function $g : I \to \mathbb{R}$ is **convex** if for every $x, y \in I$ and $\lambda \in [0, 1]$,*

$$\lambda g(x) + (1 - \lambda)g(y) \geqslant g(\lambda x + (1 - \lambda)y).$$

In words, the chord between any two points of the graph lies above the graph. We use two known facts (see, e.g., Rockafellar [Roc70]).

Fact 21.10 *Convex functions are continuous in the interior of their domain. Jumps can occur only at the boundaries.*

Fact 21.11 *Convex functions are differentiable almost everywhere in I. Formally, there is a set $I' \subseteq I$ such that $I \setminus I'$ is countable (so it has measure zero), and g is differentiable at every point of I'.*

If g is differentiable at x, we write $g'(x)$ for the derivative. The next definition extends the idea of a gradient to points where g has a kink.

Definition 21.12 (Subgradient) *For any $x \in I$, s^* is a **subgradient** of g at x if*

$$g(z) \geqslant g(x) + s^*(z - x), \quad \forall z \in I.$$

A subgradient is the slope of a line through $(x, g(x))$ that stays below g everywhere. Two natural questions: does it always exist, and is it unique?

Lemma 21.13 *Let $g : I \to \mathbb{R}$ be convex. Suppose x is in the interior of I and g is differentiable at x. Then $g'(x)$ is the unique subgradient of g at x.*

Proof: Take $z \in I$ with $z > x$ (the case $z < x$ is similar), and h with $z - x > h > 0$. Then $x + h = \frac{h}{z-x} z + \left(1 - \frac{h}{z-x}\right)x$. By convexity,

$$\frac{h}{z - x} g(z) + \left(1 - \frac{h}{z - x}\right) g(x) \geqslant g(x + h) \quad \Rightarrow \quad \frac{g(z) - g(x)}{z - x} \geqslant \frac{g(x + h) - g(x)}{h}.$$

This holds for every such $h > 0$. Letting $h \to 0$ gives $g(z) - g(x) \geqslant g'(x)(z - x)$. So $g'(x)$ is a subgradient at x.

For uniqueness, suppose for contradiction that $s^* \neq g'(x)$ is another subgradient at x. Case 1: $s^* > g'(x)$. By definition, for $h > 0$, $g(x + h) - g(x) \geqslant s^* h$, so $\frac{g(x+h)-g(x)}{h} \geqslant s^* > g'(x)$. Letting $h \to 0$ gives $g'(x) \geqslant s^* > g'(x)$, a contradiction. Case 2: $s^* < g'(x)$. Take $h < 0$; the same argument with the inequality reversed gives a contradiction. ■

Lemma 21.14 *Let $g : I \to \mathbb{R}$ be convex. Then for every $x \in I$, a subgradient of g at x exists.*

Fact 21.15 *At points of $I \setminus I'$, the set of subgradients is a convex set. Define*

$$g'_+(x) = \lim_{z \to x,\, z > x} g'(z), \qquad g'_-(x) = \lim_{z \to x,\, z < x} g'(z).$$

The set of subgradients at $x \in I \setminus I'$ is $[g'_-(x), g'_+(x)]$.

We denote the set of subgradients of g at x by $\partial g(x)$. Lemma 21.13 says $\partial g(x) = \{g'(x)\}$ for $x \in I'$, and Lemma 21.14 says $\partial g(x) \neq \emptyset$ for all $x \in I$.

Example 21.16 (Hinge loss) *The hinge loss of a support vector machine is $g(x) = \max\{0, 1 - x\}$. It is convex, with slope -1 to the left of 1 and slope 0 to the right. So $g'_-(1) = -1$ and $g'_+(1) = 0$, and $\partial g(1) = [-1, 0]$: every line through $(1, 0)$ with slope in $[-1, 0]$ stays below g. At every other point the subgradient is unique. This is exactly what subgradient descent uses to train with non-smooth losses.*

Example 21.17 (Back to the GPU-hour auction) *The Vickrey utility $u_i(t_i) = \max\{0, t_i - t_{-i}^{(2)}\}$ has the same kind of kink. So $\partial u_i(t_{-i}^{(2)}) = [0, 1]$, and $\partial u_i(t_i) = \{f_i(t_i, t_{-i})\}$ at every other point.*

Lemma 21.18 *Let $g : I \to \mathbb{R}$ be convex, and let $\phi : I \to \mathbb{R}$ satisfy $\phi(z) \in \partial g(z)$ for all $z \in I$. Then for all $x, y \in I$ with $x > y$, $\phi(x) \geqslant \phi(y)$.*

Here ϕ picks one value at every point, even where many subgradients exist. The lemma says that any such selection of subgradients is monotone.

Proof: By definition of a subgradient,

$$g(x) \geqslant g(y) + \phi(y)(x - y), \qquad g(y) \geqslant g(x) + \phi(x)(y - x).$$

Adding the two inequalities gives $(\phi(x) - \phi(y))(x - y) \geqslant 0$. Since $x > y$, $\phi(x) \geqslant \phi(y)$. ■

If g were differentiable everywhere, we would have $g(x) = g(y) + \int_y^x g'(z)\, dz$. The same holds for convex functions with any selection of subgradients.

Lemma 21.19 *Let $g : I \to \mathbb{R}$ be convex. Then for any $x, y \in I$,*

$$g(x) = g(y) + \int_y^x \phi(z)\, dz,$$

where $\phi : I \to \mathbb{R}$ satisfies $\phi(z) \in \partial g(z)$ for all $z \in I$.

Even with kinks, g is recovered by integrating any selection of its subgradients. In the next lecture, g will be an agent's utility and ϕ will be her probability of winning.

21.6 Summary

- Affine maximizers generalize VCG with agent weights w_i and an outcome bonus κ. Efficient and dictatorial rules are special cases.
- An affine maximizer satisfying INA is DSIC, with the payment (21.1).
- Roberts' theorem: with unrestricted valuations, $|A| \geqslant 3$, every onto DSIC rule is an affine maximizer. It is the quasi-linear analog of the GS theorem.
- In the single object, single parameter domain, the valuation is $a_i \cdot t_i$. The Vickrey utility is convex, and its slope is the win probability.
- Convex functions have subgradients everywhere. Any selection of subgradients is monotone, and integrates back to the function.

References

[LMN09] R. Lavi, A. Mu'alem and N. Nisan, "Two simplified proofs for Roberts' theorem," *Social Choice and Welfare*, 32, 2009, pp. 407–423.

[Rob79] K. Roberts, "The characterization of implementable choice rules," in *Aggregation and Revelation of Preferences*, J.-J. Laffont (ed.), North-Holland, 1979, pp. 321–349.

[Roc70] R. T. Rockafellar, *Convex Analysis*, Princeton University Press, 1970.