Economic vs. computational efficiency in markets

Volodymyr Kuleshov, supervised by Prof. Adrian Vetta. School of Computer Science, McGill University.

How can we share a scarce good in a fair way?

Philosophers have thought about this question for centuries. In recent years, it has been studied from a new *computational* perspective.

In this project, we want to formally quantify **economic efficiency** in markets that perform **minimal computation**.

Given impractically large (*exponential*) amounts of computation, optimal social welfare is achievable. Instead, we want to quantify economic efficiency in practice, when computation is minimal. Thus we look at the **simplest** possible markets.

Such markets are important in practice for sharing bandwidth between users on a large network...

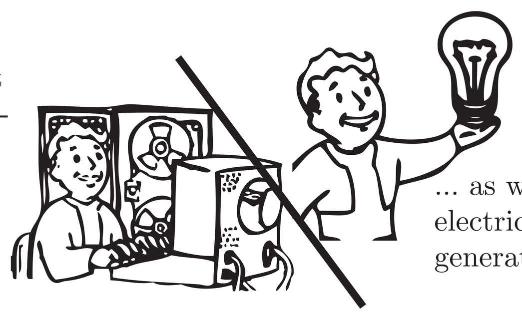

... as well as for organizing electricity production between generators and bulk buyers.

Our understanding is limited to *one-sided* markets

The most studied one-sided market is a mechanism by Kelly [2] for sharing bandwidth over a **single** network link.

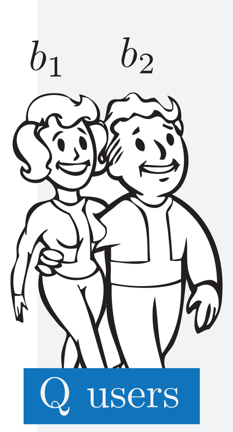

1. Each user submits a bid b_q and pays b_q\$.

2. User q receives a fraction f_q of total capacity

3. The happiness of user q is measured using a *utility function* $U_q(b_q)$.

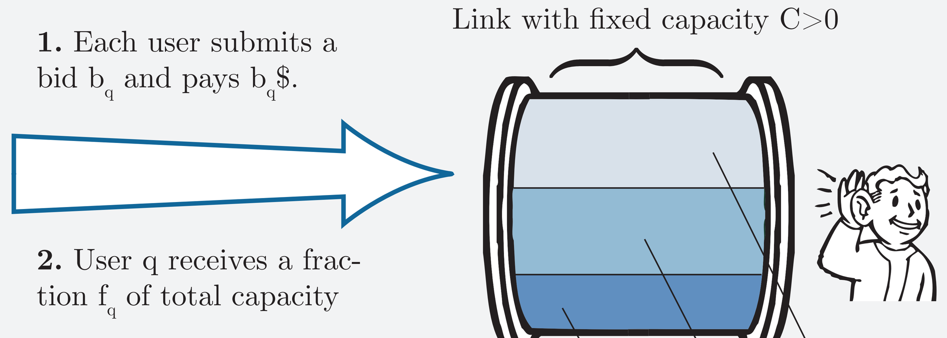

$$f_1 = \frac{b_1}{\sum_{q\in Q} b_q} C \qquad f_2 = \frac{b_2}{\sum_{q\in Q} b_q} C$$

How will users behave? For that, we need additional concepts:

Definition. Bids $b = (b_1, ..., b_Q)$ form a *Nash equilibrium* if $\forall q \in Q$ we have $b_q = \arg\max_{\bar{b}} U_q(\bar{b}, b_{-q})$.

We'll assume that once all the bids are in, we have a Nash equilibrium: each user submits his **best bid**, given what the others have done.

Observe that everyone does what's best for him, and not for society. This can reduce social welfare, defined as the **sum of all the utilities**. We measure this loss using the **price of anarchy**.

Definition. Let b^{NE} be the Nash equilibrium that achieves the lowest social welfare $\sum_{q\in Q} U_q(b^{NE})$. The *price of anarchy [3]* is the ratio

$$\rho = \frac{\sum_{q\in Q} U_q(b^{NE})}{\sup_{\bar{b}} \sum_{q\in Q} U_q(\bar{b})}$$

That's worst Nash equilibrium welfare over best possible welfare.

Theorem. *The price of anarchy in Kelly's mechanism for users is* **3/4** *[1].*

Theorem. *When providers compete for a fixed demand, the price of anarchy is* **1/2** *[5].*

WHERE OUR RESEARCH COMES IN

We understand user and provider competition, but **not both** at once. Establishing formal efficiency guarantees for *two-sided* markets is an important open problem.

We propose the first efficient **two-sided** market:

A **user's** utility function:

$$U_q(b_q, \underbrace{b_{-q}}_{\text{others' bids}}) = \underbrace{V_q(f_q)}_{\text{value}} - \underbrace{pf_q}_{\text{expenses}}$$

Assumption. *Valuation functions $V_q(f_q) : \mathbb{R}^+ \to \mathbb{R}^+$ are strictly increasing, concave, and differentiable.*

A **provider's** utility function:

$$U_r(b_r, b_{-r}) = \underbrace{ps_r}_{\text{revenue}} \underbrace{- C_r(s_r)}_{\text{costs}}$$

Assumption. *The cost functions $C_r(s_r) : \mathbb{R}^+ \to \mathbb{R}^+$ are differentiable and the derivative is continuous, convex, strictly increasing, and $C'(0) = 0$.*

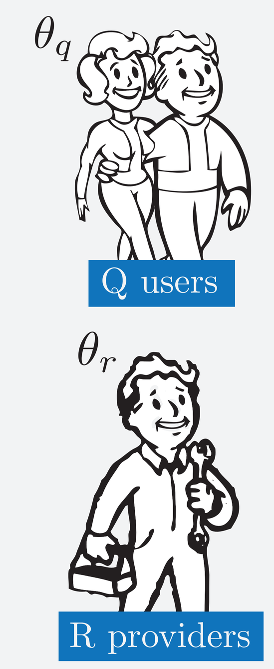

1. Each participant submits a bid b.

2. The market computes a single market-clearing price p.

3. Provider r is told to offer s_r units of bandwidth. User q gets a fraction f_q of total production.

4. User q pays pf_q and provider r is paid ps_r.

$$p = \frac{\sum_{q\in Q} b_q + \sum_{r\in R} b_r}{R} \qquad f_q = \frac{b_q}{p} \qquad s_r = 1 - \frac{b_r}{p}$$

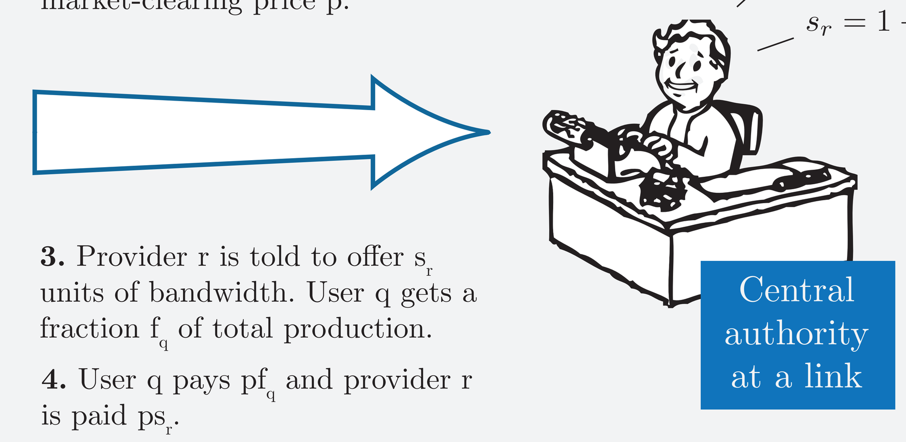

At equilibrium, total supply must **equal** total demand:

$$\sum_{q\in Q} \frac{b_q}{p} = \sum_{r\in R} \left(1 - \frac{b_r}{p}\right)$$

$$\Downarrow$$

$$p = \frac{\sum_{q\in Q} b_q + \sum_{r\in R} b_r}{R}$$

How efficient is this new market?

When market participants have **no market power**, our market achieves optimal social welfare.

Theorem 1. *The two-sided market always has a competitive equilibrium. The welfare achieved at a competitive equilibrium is optimal.* □

When participants **have market power**, they use it in their own interests, and as a whole, society is worse off. In our market, the loss of welfare (the *price of anarchy*) is about 40%.

Theorem 2. *There exists a unique Nash equilibrium in the two-sided market with $R > 1$ providers.* □

Theorem 3. *The price of anarchy of the two-sided market when there are $R > 1$ providers is*

$$\frac{s^2(S^2 + 4Ss + 2s^2)}{S(S + 2s)}$$

where $S = R - 1$, and s is the unique positive root of the polynomial

$$\gamma(s) = 16s^4 + 10S^2 s(3s - 2) + S^3(5s - 4) + Ss^2(49s - 24)$$

Furthermore, this bound is tight.

Corollary 4. *The greatest loss of efficiency occurs when $R = 2$. The price of anarchy then equals approximately* **0.588727**. □

We have defined our market for a single network link, but our efficiency guarantees hold for **general networks**! In that case, the mechanism is slightly different:

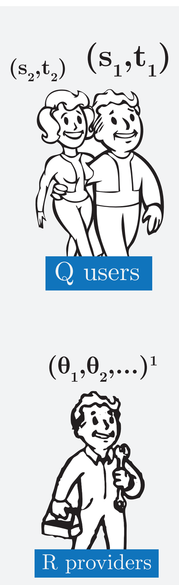

1. Each user wants to send traffic between an (s,t) pair of nodes.

2. A market instance at each link takes a bid from everyone. It independently handles money, allocations.

3. Users send max (s,t)-flows given bought capacities. Utilities are functions of flows, not capacities.

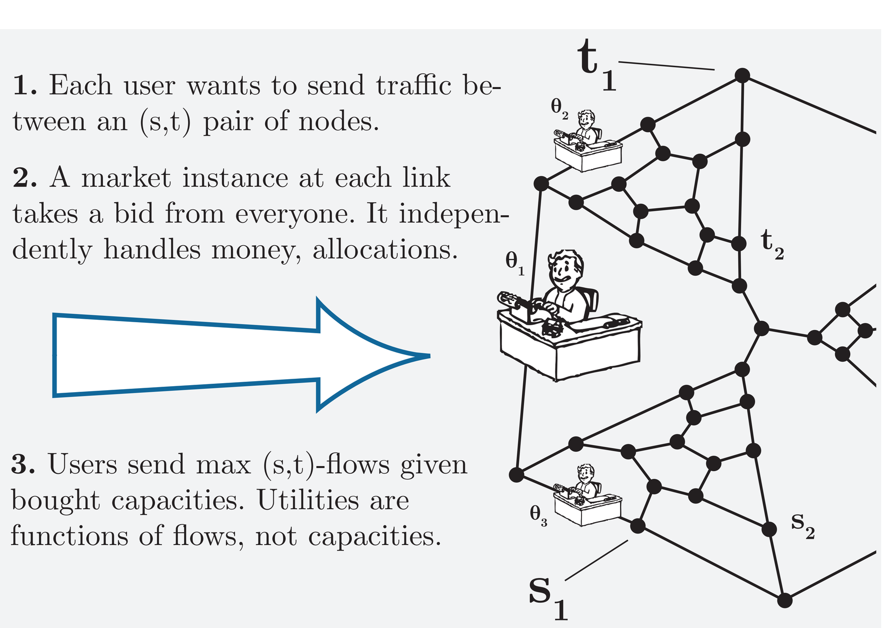

Theorem 5. *When extended to networks, the mechanism retains the same price of anarchy.*

Corollary 6. *When extended to a general market of N goods, the mechanism has the same price of anarchy of* **0.588727**. □

So in fact, this market can be used to share **any type of resource**!

Outline of main proofs

Proof of theorem 3. First, from our differentiability and convexity assumptions, we obtain necessary and sufficient first order Nash equilibrium conditions.

We then show that the worst efficiency occurs when valuations are linear. We can also show that marginal costs are linear in the worst case (harder).

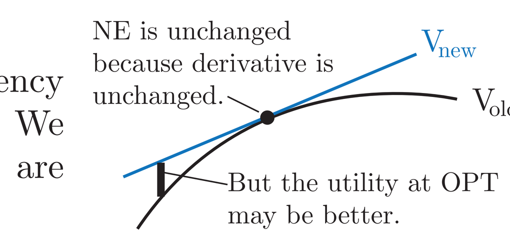

Next, we express the price of anarchy as the solution of an optimization problem. We minimize over all possible allocations and utilities under the constraint that they form a Nash equilibrium:

minimize	$\dfrac{d_1^{NE} + \sum_{i=2}^{Q} \alpha_i d_i^{NE} - \frac{1}{2}\sum_{j=1}^{R} \beta_j (s_j^{NE})^2}{\sum_{j=1}^{R} s_j^{OPT} - \frac{1}{2}\sum_{j=1}^{R} \beta_j (s_j^{OPT})^2}$	(1)	Price of anarchy
such that	$\alpha_i \left(1 - \frac{d_i^{NE}}{R}\right) \geq \mu \ \forall i \text{ s.t. } d_i^{NE} > 0$	(2)	Nash eq. conditions
	$\alpha_i \left(1 - \frac{d_i^{NE}}{R}\right) \leq \mu \ \forall i$	(3)	
	$\beta_j s_j^{NE} \left(1 + \frac{s_j^{NE}}{R-1}\right) \leq \mu \ \forall j \text{ s.t. } 0 < s_j^{NE} \leq 1$	(4)	
	$\beta_j s_j^{NE} \left(1 + \frac{s_j^{NE}}{R-1}\right) \geq \mu \ \forall j \text{ s.t. } 0 \leq s_j^{NE} < 1$	(5)	
	$\sum_{i=1}^{Q} d_i^{NE} = \sum_{j=1}^{R} s_j^{NE}$	(6)	Supply = demand
	$\beta_j s_j^{OPT} \leq 1 \ \forall j \text{ s.t. } 0 < s_j^{OPT} \leq 1$	(7)	Optimality conditions
	$\beta_j s_j^{OPT} \geq 1 \ \forall j \text{ s.t. } 0 \leq s_j^{OPT} < 1$	(8)	
	$d_i^{NE} \geq 0 \ \forall i$	(9)	Non-negativity
	$0 \leq s_j^{NE}, s_j^{OPT}, \alpha_i \leq 1 \ \forall i, j$	(10)	
	$0 \leq \mu$	(11)	

Constraints (2,3) and (4,5) hold with equality at OPT. Incorporating equality constraints into the objective function eliminates part of the variables. The objective function becomes concave and symmetric in the d-variables. At OPT they must be all equal, and can be replaced by only one unknown. Then, applying a technical lemma and lots of algebra, we eventually obtain the two-variable problem below. We proceed to show that its solution yields the desired bound. □

minimize	$\dfrac{(1-\mu)^2 + \mu s - {}^{\mu}/_{2} \frac{s}{1+s/(R-1)}}{\min(\frac{s(1+s/(R-1))}{\mu}, 1) - \frac{\mu}{2s(1+s/(R-1))} \min(\frac{s(1+s/(R-1))}{\mu}, 1)^2}$	(12)
such that	$0 < s \leq 1$	(13)
	$0 \leq \mu < 1$	(14)

Proof of theorem 5. Providers' costs are separable in the links, but not users' valuations. By convexity, U_q has a supporting hyperplane, the subgradient. This hyperplane is linearly separable in the links and bounds the original price of anarchy. The case of linearly separable utilities can now be bounded by the price of anarchy of a single link. □

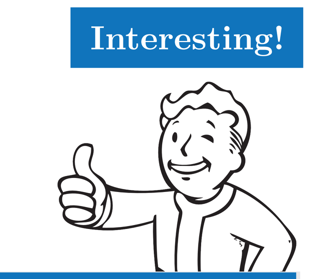

This market is optimal within its class

Kelly's mechanism is optimal within a large class of mechanisms. These mechanisms have a natural extension to two-sided markets.

Definition. A *smooth two-sided market-clearing mechanism* is a tuple of functions (D, S), $D : (0, \infty) \times [0, \infty) \to \mathbb{R}$, $S : (0, \infty) \times [0, \infty) \to \mathbb{R}$ such that for all Q, R, and for all $b \in \mathbb{R}^{Q+R}, b \neq 0, b \geq 0$, there exists a unique $p > 0$ that satisfies the following equation

$$\sum_{i=1}^{Q} D_i(p, b_i) = \sum_{j=1}^{R} S_j(p, b_j)$$

Definition. The class $\mathcal{M}$ contains all smooth two-sided market-clearing mechanisms (D, S) such that:

1. The utility function of every market participant is concave in its bid when they have no market power.
2. The utility function of every market participant in concave in its bid when there is market power.
3. The demand function D is non-negative. The supply function S is bounded above.
4. The competitive equilibrium allocations are optimal.
5. Users and providers have the same expressive power: demand and the supply functions are symmetric in b in the sense that for all $p > 0$, $b \geq 0$,
$$\frac{\partial D}{\partial b}(p, b) = -\frac{\partial S}{\partial b}(p, b)$$

Theorem 7. *Out of all smooth two-sided market-clearing mechanisms $(D, S) \in \mathcal{M}$, our mechanism uniquely achieves the best possible price of anarchy.*

Proof outline. We can deduce from (1-5) that all mechanisms in this class must have the form:

$$\begin{aligned} D(b, p) &= f(p)b \\ S(b, p) &= 1 - f(p)b \end{aligned}$$

where f(p) is a differentiable and invertible function of the price p. We can now derive new necessary and sufficient Nash equilibrium conditions and repeat the argument of the main proof. This time we optimize over all possible functions f(p).

We find that the best price of anarchy occurs when the elasticity of f(p) equals one. This elasticity is uniquely achieved by our mechanism. □

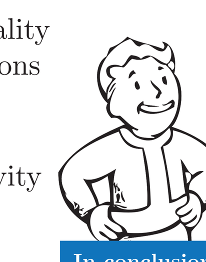

Within our computationally efficienient market system, there is a 40% loss in economic efficiency. Designing this market was non-trivial, and most similar mechanisms are inefficient. This raises questions about the economic efficiency of real markets.

It is not clear how prices are formed in real markets, but a version of our mechanism has been studied in economics as a **model of price formation** [4]. Thus our results are also relevant within an existing body of economic literature on prices.

References

[1] *Ramesh Johari and John N. Tsitsiklis. Efficiency loss in a network resource allocation game.* Math. Oper. Res., *29(3):407–435, 2004.*

[2] *Frank Kelly. Charging and rate control for elastic traffic.* European Transactions on Telecommunications, *1997.*

[3] *Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In* in Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, *pages 404–413, 1999.*

[4] *Lloyd S Shapley and Martin Shubik. Trade using one commodity as a means of payment.* Journal of Political Economy, *85(5):937–68, October 1977.*

[5] *John N. Tsitsiklis, Arthur C. Smith, Ramesh Johari, and Ramesh Johari. Efficiency loss in market mechanisms for resource allocation. Technical report, 2004.*