Welfare Theorems for Democratic Politics

Jon X. Eguia (MSU) Glen Weyl (Microsoft) Dimitrios Xefteris (UCY)

CRETE 2026
Ios | 7-11 July 2026

Big picture

Economy

- 1st Welfare Theorem: Competitive markets yield Pareto efficiency.
- 2nd Welfare Theorem: By reallocating endowments, any Pareto-efficient allocation can be decentralized.
- Decentralized allocation of private goods can be efficient.

Politics

- Voting does *not* generally aggregate intensities.
- Issue-by-issue one-person-one-vote can select outcomes that are Pareto dominated.
- Is decentralized collective choice inherently incompatible with efficiency?

Majority voting across issues can be inefficient

	Issue 1 (A_1 vs. B_1)	Issue 2 (A_2 vs. B_2)	Issue 3 (A_3 vs. B_3)
Voter 1	-10	1	1
Voter 2	1	-10	1
Voter 3	1	1	-10

Under issue-by-issue majority rule

A_k wins in every issue.

Each voter loses the one issue she cares most about and gets utility -8.

Pareto improvement

If B_k wins in every issue, each voter gets 0.

So the majority outcome is **strictly Pareto dominated**.

Existing approaches

Mixing the economic and political sphere

- This is the logic behind VCG-type mechanisms and AGV-type mechanisms.
- More recent proposals such as quadratic voting (LW16) also follow this broader mechanism-design route.
- But this approach relies on the absence of wealth effects (EX21).

Staying within the political sphere

- The main benchmark here is *storable votes* (C05)–voters shift a purely political budget across issues–and *logrolling* (CP19)–voters exchange votes.
- These mechanisms preserve the separation of spheres, but existing approaches do not provide a welfare-theorem foundation.

References: Vickrey (1961); Clarke (1971); Groves (1973); d'Aspremont and Gérard-Varet (1979); Casella (2005); Posner and Weyl (2015); Lalley and Weyl (2016); Goeree and Zhang (2017); Casella and Palfrey (2019); Eguia and Xefteris (2021); Casella and Sanchez (2022); Pivato and Ghosh (2025).

Research question and answer

Question	Answer
Can we recover welfare-theorem logic for collective choice *without* allowing private wealth to purchase political power?	**Yes**, if influence can be transferred across issues with a **quadratic** conversion rate between voice credits and votes.

Voice credits B_i

Issue allocation q_{ik}

Aggregate score s_k

- $\sum_k \alpha |q_{ik}|^r \leq B_i$
- $p_k = G'(s_k)$
- $r \to \infty$: nearly rigid one-vote-per-issue politics.
- $r \to 1$: nearly frictionless transfer of influence.
- $r = 2$: the unique curvature that delivers the welfare theorems.

What the paper contributes

1. **A general mechanism class:** agents allocate a political influence budget across binary issues, subject to an r-power concentration cost.
2. **Equilibrium existence:** a competitive political equilibrium exists for any fixed finite economy under weak regularity conditions.
3. **First Welfare Theorem for Politics:** in large replicated societies, **quadratic costs** imply asymptotic Pareto efficiency.
4. **Uniqueness:** for every $r \neq 2$, there are profiles for which equilibrium is asymptotically Pareto inefficient.
5. **Second Welfare Theorem for Politics:** under the right budget assignment and $r = 2$, any target Pareto-efficient outcome can be achieved in equilibrium.

Mechanism: transferable voice credits with concentration costs

Environment

- n agents, K binary issues.
- Aggregate score in issue k:

$$s_k(q) = \sum_{i=1}^{n} q_{ik}.$$

- Alternative A_k is implemented with prob.

$$y_k = G\big(s_k(q)\big), \qquad G'(x) \equiv g(x) \in \mathbb{R}_{++}.$$

- $G(x) = 1 - G(-x) \to 1$ as $x \to +\infty$.
- $U_i(q) = \sum_{k=1}^{K} v_{ik}\, G\big(s_k(q)\big)$, with $v_{ik} \neq 0$.

Feasibility constraint

$$Q_i^{(r)}(B_i) = \left\{ q_i \in \mathbb{R}^K : \ \alpha \sum_{k=1}^{K} |q_{ik}|^r \leq B_i \right\}.$$

- $r > 1$.
- B_i: voice credits only.
- Not exchangeable for private consumption.
- The design preserves a **separation of spheres**.

Why the exponent r matters

Interpretation

- Higher r: influence is expensive to concentrate.
- Lower r: influence is easier to move across issues.
- The mechanism nests standard benchmarks.

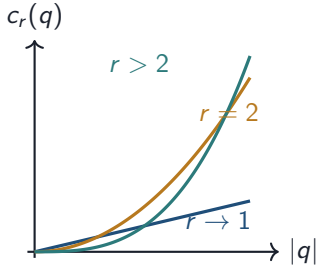

$r \to \infty$	Almost rigid "one-vote-per-issue" allocation
$r = 2$	Quadratic voting benchmark; marginal influence scales linearly in vote purchases
$r \to 1$	Nearly linear transferability, close to storable-votes logic

Competitive political equilibrium (CPE)

Price-taking problem

Given perceived pivotalities $p = (p_1, \ldots, p_K) \in (0, \infty)^K$, agent i solves

$$\max_{q_i \in Q_i^{(r)}(B_i)} \sum_{k=1}^{K} v_{ik} p_k q_{ik}.$$

- p_k is the perceived marginal effect of influence on the probability of A_k.
- Agents take p_k as parametric, exactly like price-taking consumers do in markets.

Equilibrium condition

A pair (p^*, q^*) is a CPE if:

1. each q_i^* solves the price-taking problem at p^*,
2. perceived pivotalities are correct:

$$p_k^* = g\big(s_k(q^*)\big) \quad \forall k.$$

Conceptually: prices of influence clear the political "market".

Closed-form best response

Lemma 1

For $\rho = \frac{r}{r-1}$, the unique best response is

$$q_{ik}(p, B_i) = \left(\frac{B_i}{\alpha}\right)^{1/r} \frac{\operatorname{sgn}(v_{ik}p_k)\,|v_{ik}p_k|^{1/(r-1)}}{\left(\sum_{\ell=1}^{K} |v_{i\ell}p_\ell|^{\rho}\right)^{1/r}}, \qquad \mu_i(p, B_i) = \left(\frac{1}{rB_i}\sum_{\ell=1}^{K} |v_{i\ell}p_\ell|^{\rho}\right)^{(r-1)/r}.$$

Interpretation

Scale is governed by B_i; direction is governed by the relative magnitudes of $v_{ik}p_k$ across issues. The denominator is the common shadow-price normalization.

Quadratic benchmark

At $r = 2$, there exists a scalar $\kappa_i > 0$ such that

$$q_{ik} = \kappa_i v_{ik} p_k.$$

This is the linearity that later drives efficiency.

Why quadratic costs are special

Power transformation

Define

$$\psi_r(x) = \operatorname{sgn}(x)|x|^{1/(r-1)}.$$

Then in any equilibrium, the sign of the aggregate score in issue k is driven by a weighted sum of *transformed* valuations:

$$\operatorname{sgn}\big(s_k(q^*)\big) = \operatorname{sgn}\left(\sum_i \omega_i\, \psi_r(v_{ik})\right), \qquad \omega_i > 0.$$

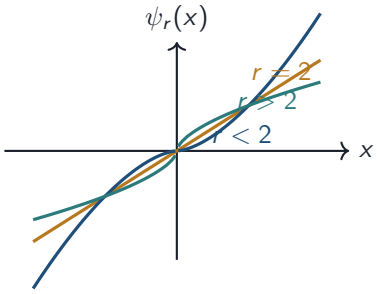

Knife-edge result

ψ_r **is linear if and only if** $r = 2$**.** Only then is the induced decision rule a linear combination of the *original* valuations (v_{ik}).

Existence of a CPE for any finite society

- Budgets imply a uniform bound on feasible scores:

$$|s_k| \le \sum_i \left(\frac{B_i}{\alpha}\right)^{1/r} =: S(B).$$

- Therefore feasible pivotalities live in a compact box

$$P(B) = \left[\underline{p}(B), \overline{p}(B)\right]^K.$$

- Best responses vary continuously in p.

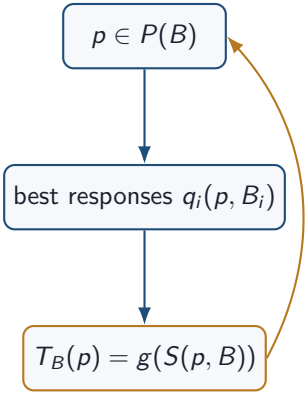

This mirrors general equilibrium logic: define a compact price set, show continuity, and close the system with a fixed point.

Asymptotic approach

Construction

Start from a finite set of agent types $t = 1, \dots, n_0$ with valuations v^t and budgets B_t. Replicate each type m times.

$$(t, a), \qquad a = 1, \dots, m.$$

All clones of a given type face the same problem and therefore choose the same vote vector.

Aggregation under replication

For each issue k,

$$s_k\big(q^{(m)}\big) = m\,\bar{S}_k\big(p^{(m)}\big),$$

where $\bar{S}_k(\cdot)$ is the per-replica type aggregate.

- This allows a clean large-society limit.
- The asymptotic question becomes: what kind of social choice rule is induced when $m \to \infty$?

The induced deterministic decision rule

Lemma 2

At any equilibrium,

$$s_k(q^*) = (p_k^*)^{1/(r-1)} \sum_{i=1}^{n} \omega_i \, \psi_r(v_{ik}), \qquad \omega_i > 0,$$

so the sign of the collective decision depends on a weighted sum of transformed valuations.

If $r = 2$

$$\psi_2(v_{ik}) = v_{ik}.$$

Hence the rule is linear in the original intensities.

If $r \neq 2$

The rule is linear only in the *distorted* objects $\psi_r(v_{ik})$, not in the primitive valuations v_{ik}.

Linear aggregation in the limit iff $r = 2$

Proposition 1

In the replication limit, the collective choice rule is representable as a linear rule in the original valuations if and only if $r = 2$.

- For $r = 2$, the equilibrium score sign is linear in v_{ik}.
- The induced rule respects marginal trade-offs in primitive intensities.
- This is exactly the analogue of price-mediated efficiency in markets.

Structural uniqueness

The result is not merely "quadratic works." It is stronger:

quadratic is the unique member of the r-power family that yields the correct linear aggregation for *all* preference profiles.

The right welfare object in large societies

Endogenous welfare weights

For type t in the replicated economy, define

$$\theta_t^{(m)} := \frac{\overline{p}^{(m)}}{\mu_t^{(m)}} > 0, \qquad \overline{p}^{(m)} := \max_k p_k^{(m)}.$$

The corresponding per-replica weighted welfare is

$$W^{(m)}(y) = \sum_{t=1}^{n_0} \theta_t^{(m)} \sum_{k=1}^{K} v_{tk} y_k.$$

Interpretation

$\mu_t^{(m)}$ is the shadow price of influence; $\theta_t^{(m)}$ is its inverse, normalized by maximal pivotality.

Why this matters

Efficiency is evaluated using **endogenous weights** implied by equilibrium.

First Welfare Theorem for Politics

Definition (Asymptotic Pareto Efficiency)

A sequence of Competitive Political Equilibria is *Asymptotically Pareto Efficient* if the probability that the policy vector is Pareto Efficient converges to one.

Theorem 1

Under **quadratic concentration costs** ($r = 2$), every sequence of competitive political equilibria is asymptotically Pareto optimal.

Proof intuition for the First Welfare Theorem

1. Under $r = 2$, the KKT conditions imply a score identity:

$$\bar{S}_k\big(p^{(m)}\big) = \frac{p_k^{(m)}}{2\alpha\,\bar{p}^{(m)}}\, C_k^{(m)}, \qquad C_k^{(m)} := \sum_t \theta_t^{(m)} v_{tk}.$$

2. Hence the sign of the equilibrium score matches the sign of the weighted aggregate valuation $C_k^{(m)}$.
3. If an issue becomes deterministic, equilibrium picks the welfare-maximizing option.
4. If an issue fails to be deterministic, then $C_k^{(m)} \to 0$: the issue becomes asymptotically welfare-irrelevant.

Bottom line

Quadratic pricing makes decentralized political influence match marginal welfare gains across issues.

What fails when $r \neq 2$

Source of inefficiency

When $r \neq 2$, the score sign depends on $\psi_r(v)$ instead of v.

- If $r < 2$, the transformation amplifies large intensities.
- If $r > 2$, it compresses them.
- Either way, the induced social choice rule is not aligned with primitive welfare trade-offs.

Consequence

We prove a **tight uniqueness result**: for every $r \neq 2$, there exists a replicated economy whose equilibrium outcome is asymptotically Pareto dominated.

Counterexample setup for every $r \neq 2$

Five issues, five types, replicated m times

In each issue, three types support A_k and two types support B_k. Intensities are assigned as follows:

- supporters of A_k: $\{v, \varepsilon, \varepsilon\}$
- supporters of B_k: $\{z, z\}$

with

$$2z > v + 2\varepsilon \qquad \text{and} \qquad v^{\beta} + 2\varepsilon^{\beta} > 2z^{\beta}, \quad \beta = \frac{1}{r-1} \neq 1.$$

Issue k	1	2	3	4	5
Supporters of B_k	$\{4,5\}$	$\{1,5\}$	$\{1,2\}$	$\{2,3\}$	$\{3,4\}$

The construction rotates who is intense on the majority side and who is moderately intense on the minority side.

The parametrized utility table

Voter	Issue 1	Issue 2	Issue 3	Issue 4	Issue 5
1	$+v$	$-z$	$-z$	$+\varepsilon$	$+\varepsilon$
2	$+\varepsilon$	$+v$	$-z$	$-z$	$+\varepsilon$
3	$+\varepsilon$	$+\varepsilon$	$+v$	$-z$	$-z$
4	$-z$	$+\varepsilon$	$+\varepsilon$	$+v$	$-z$
5	$-z$	$-z$	$+\varepsilon$	$+\varepsilon$	$+v$

- The sign tells us whether the voter prefers A_k or B_k.
- The magnitude tells us how intensely she cares.

What one-person-one-vote ($r \to \infty$) does in this family

Majority pattern

On every issue there are three supporters of A_k and two supporters of B_k. Therefore plain majority chooses

$$(A_1, A_2, A_3, A_4, A_5).$$

But the welfare comparison goes the other way if

$$2z > v + 2\varepsilon.$$

Then each voter prefers the all-B vector to the all-A vector.

So majority rule can again select a Pareto-dominated outcome.

A concrete variant: $v = 10$, $z = 8$, $\varepsilon = 1$

Voter	1	2	3	4	5
1	$+10$	-8	-8	$+1$	$+1$
2	$+1$	$+10$	-8	-8	$+1$
3	$+1$	$+1$	$+10$	-8	-8
4	-8	$+1$	$+1$	$+10$	-8
5	-8	-8	$+1$	$+1$	$+10$

Pareto ranking

Every voter gets value -4 from the all-A outcome and 0 from the all-B outcome. So the all-B vector is **strictly Pareto superior**.

Why absence of concentration penalties ($r \to 1$) can prove detrimental

Equilibrium behavior

Each voter spends her five votes on the issue she cares about most. Since $v = 10$ is larger than $z = 8$, each voter's favorite issue is the one on which she strongly prefers A.

Voter	Issue 1	Issue 2	Issue 3	Issue 4	Issue 5
1	5 for A	0	0	0	0
2	0	5 for A	0	0	0
3	0	0	5 for A	0	0
4	0	0	0	5 for A	0
5	0	0	0	0	5 for A

Why the counterexample works

Primitive welfare comparison

Under the all-A outcome, each type gets

$$v + 2\varepsilon.$$

Under the all-B outcome, each type gets

$$2z.$$

Since $2z > v + 2\varepsilon$, the all-B outcome **strictly Pareto dominates** the all-A outcome.

But transformed intensities favor A

Because

$$v^{\beta} + 2\varepsilon^{\beta} > 2z^{\beta},$$

there is a symmetric equilibrium in which every issue has a positive score, so the replicated CPE converges to the all-A outcome.

The uniqueness result is genuinely sharp

Proposition 2

For every $r > 1$ with $r \neq 2$, there exists a replicated five-issue economy such that:

1. the equilibrium outcome converges to all-A,
2. the all-B outcome strictly Pareto dominates all-A.

Hence **quadratic costs are the unique efficient configuration** within the entire r-power family.

Why the example breaks at $r = 2$

If $r = 2$, then $\beta = 1$, and the two inequalities become

$$2z > v + 2\varepsilon \qquad \text{and} \qquad v + 2\varepsilon > 2z,$$

which is impossible. The counterexample collapses exactly at the quadratic benchmark.

Second Welfare Theorem for Politics

Question

Suppose we want to decentralize a given Pareto-efficient collective allocation. Can we do so *without* importing private wealth into politics?

Answer: yes, via voice-credit redistribution

The political analogue of lump-sum transfers is a redistribution of **influence budgets** B_i.

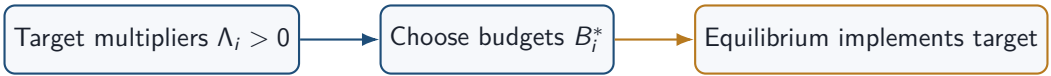

Second Welfare Theorem for Politics

Theorem 2

For any Pareto efficient vector of policies, there exist budgets $B_i^* > 0$ and a corresponding sequence of equilibria under quadratic costs, such that the probability that this policy vector is implemented converges to one.

Construction

For any provisional pivotality vector p and $r = 1$, define

$$q_{ik}(p) = \operatorname{sgn}(v_{ik}p_k)\left(\frac{|v_{ik}p_k|}{\alpha r \Lambda_i}\right)^{1/(r-1)}$$

and choose the budget that makes this vector just feasible:

$$B_i(p) = \alpha \sum_k |q_{ik}(p)|^r.$$

Closing the system

Aggregate these induced votes into scores $S_k(p)$ and solve the fixed point

$$p_k = g(S_k(p)).$$

At the fixed point, the constructed budgets decentralize the target multipliers.

Interpretation: a political welfare-theorem pair

First Welfare Theorem

With quadratic conversion, equilibrium aggregates political intensities efficiently in large societies.

Second Welfare Theorem

With the right distribution of voice credits, any efficient outcome can be supported in equilibrium.

Central conceptual point

The paper recovers welfare-theorem logic *without monetizing politics*. Private wealth never buys influence; only a separate endowment of voice credits does.

- This preserves a normatively attractive **separation of spheres**.
- It also clarifies the exact mechanism-design feature that is doing the work: **quadratic pricing of influence concentration**.
- It does not require assumptions regarding interpersonal comparability of utility.

What this means for political economy

1. **Democratic inefficiency is not inevitable.** The inefficiency of standard voting comes from rigid equal weighting issue by issue, not from the impossibility of decentralized aggregation.
2. **Intensity can be aggregated without money.** The mechanism uses a distinct political budget, so the result is not an argument for vote buying.
3. **Quadratic pricing is not a convenience assumption.** It is a uniquely correct curvature result, not merely a tractable benchmark.
4. **The equilibrium analogy.** Pivotalities play the role of prices; voice credits play the role of endowments; redistribution of voice credits plays the role of lump-sum transfers.

Takeaways

Main message

There is a full welfare-theorem logic for democratic politics—but only when political influence is allocated through **quadratically priced voice credits** rather than rigid one-person-one-vote across issues or frictionless transfer of votes.

- Existence for finite societies.
- Efficiency in replicated economies.
- Tight uniqueness of the quadratic benchmark.
- Implementation of any target multipliers via budgets.
- Strong separation between economic wealth and political power.
- Clear bridge between social choice and general equilibrium.

Thank you

Appendix: best-response derivation in one line

Lagrangian

$$\mathcal{L}(q_i, \mu) = \sum_{k=1}^{K} v_{ik} p_k q_{ik} + \mu \left(B_i - \alpha \sum_{k=1}^{K} |q_{ik}|^r \right).$$

KKT stationarity gives

$$v_{ik} p_k = \mu \alpha r \, \operatorname{sgn}(q_{ik}) |q_{ik}|^{r-1} \quad \Longrightarrow \quad q_{ik} = \operatorname{sgn}(v_{ik} p_k) \left(\frac{|v_{ik} p_k|}{\mu \alpha r} \right)^{1/(r-1)}.$$

Imposing the binding budget constraint pins down the multiplier and yields the closed form used in the text.

Appendix: score sign representation

General equilibrium sign formula

With

$$\omega_i := \left(\frac{B_i}{\alpha}\right)^{1/r} \left(\sum_{\ell=1}^{K} |v_{i\ell} p_\ell|^{\rho}\right)^{-1/r} > 0,$$

we can write

$$q_{ik} = (p_k)^{1/(r-1)} \omega_i \, \psi_r(v_{ik}).$$

Summing over agents yields

$$s_k = (p_k)^{1/(r-1)} \sum_i \omega_i \psi_r(v_{ik}).$$

Why this matters

Everything hinges on whether ψ_r is linear. That is exactly the condition $r = 2$.

Appendix: the counterexample inequalities

Need both

$$2z > v + 2\varepsilon \qquad \text{and} \qquad v^{\beta} + 2\varepsilon^{\beta} > 2z^{\beta}, \qquad \beta \neq 1.$$

- The first inequality says *primitive welfare* prefers the all-B allocation.
- The second says the *transformed* score rule induced by $r \neq 2$ prefers the all-A allocation.
- The paper shows such triples (v, z, ε) exist for every $\beta \neq 1$.
- At $\beta = 1$ (equivalently $r = 2$), the two inequalities are contradictory.