Automated Market Making for Peer-to-Peer Energy Sharing

Michele Fabi   Viraj Nadkarni
Leonardo Leone   Matheus X. V. Ferreira

Telecom Paris & CREST | Princeton University | University of Virginia

December 17, 2025

Motivation: Microgrids

- Energy markets are shifting from a **top-down, producer-centric** to a **bottom-up, prosumer-centric** approach

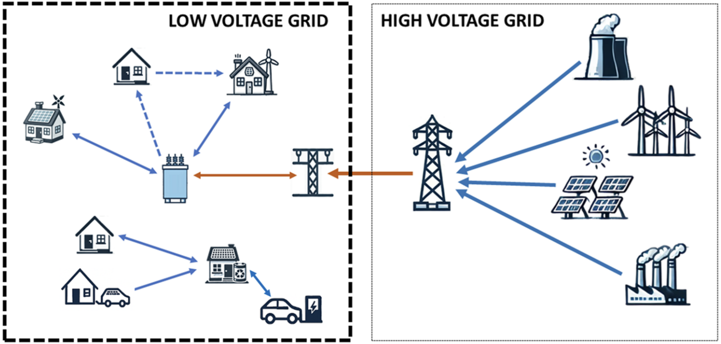

Motivation: Blackouts

Spain-Portugal Blackout

- After some frequency oscillations, the grid records a sudden drop of 1.5GW in power supply from solar
- This reduces frequency from 50Hz to 49Hz causing a security shutdown of the national grid

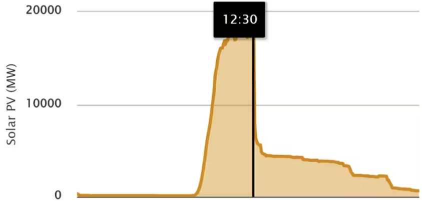

Solar photovoltaic (PV) generation on 28 April 2025. REE

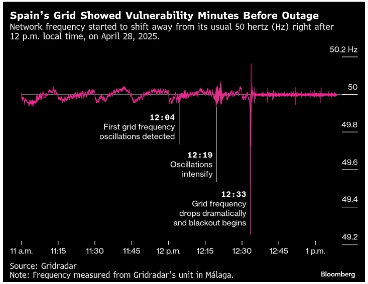

Motivation: Price gaps

Actual prices - EU

Country	Avg. Buying Price (€/kWh)	Highest Selling Price (€/kWh)
Germany	0.42	0.13–0.18 (feed-in tariff)
France	0.25	0.13–0.18 (feed-in tariff)
Italy	0.31	0.12–0.16 (feed-in tariff)
Spain	0.32	0.06–0.10 (feed-in tariff)
Netherlands	0.32	0.08–0.12 (feed-in tariff)
Denmark	0.40	0.10–0.13 (feed-in tariff)
Sweden	0.24	0.07–0.10 (feed-in tariff)
Poland	0.18	0.07–0.10 (feed-in tariff)

Big Picture

Goal: design a simple, decentralized pricing mechanism for peer-to-peer (P2P) energy sharing that:

- respects grid prices (no-arbitrage / autonomy),
- is fair and anonymous (order-agnostic, coalition-proof),
- induces peak shaving and remains budget-balanced.

Approach: adapt AMM logic (à la Uniswap) to energy trading with *concentrated liquidity* and *batch clearing*.

Result: an AMM that quotes a local price around the *arithmetic or geometric mean* of grid bid/ask, adjusted for imbalances.

Contributions

1. Axiomatic theory for local energy pricing (anonymity, coalition-proofness, budget-balance, etc.).
2. **AMM construction** satisfying the axioms, using batch clearing, concentrated liquidity, and re-anchored bonding curves.
3. **Characterization** and **Computation** of a Markov Perfect Equilibrium (MPE) for prosumer community using a **Mean-Field Game (MFG)** framework.
4. **Numerical experiments** using data from the Paris metropolitan area (IDF).

Roadmap

Optimization and Equilibrium of Power Systems

Unlike standard ADMM, axioms impose rationality even **away from convergence** Boyd et al. (2011); Erseghe (2014); F. Moret and Pinson (2024).

Battery aggregation and optimization Prat et al. (2024); Berger and Kassoul (2025).

From bilevel/MPEC models to **routing/congestion games** (Decentralized coordination) Rosenthal (1973); Monderer and Shapley (1996).

Blockchain & DeFi

Principles of **CFMMs** (Angeris et al., 2020, 2023; Schlegel et al., 2023; Fabi and Prat, 2025) applied to **DePIN** (Decentralized Physical Infrastructure Networks) Milionis et al. (2025).

P2P Prosumer Communities

Advances on P2P energy sharing Sousa et al. (2019); Crowley et al. (2025); Pinson et al. (2020) sharing by full **axiomatic formulation**.

Mean Field Games (MFG)

Equilibrium via **Discrete MFGs** Doncel et al. (2019); Lasry and Lions (2007); Guéant et al. (2011).

Roadmap

Microgrid and Power Flows

Network Nodes:

N prosumers;

aggregator (node A);

main grid (node 0).

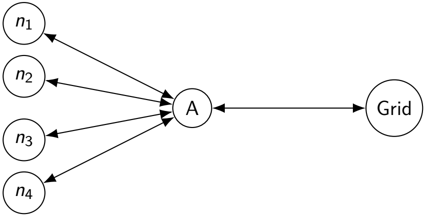

Network structure: prosumers $\leftrightarrow$ aggregator $\leftrightarrow$ grid.

Prosumer Power Flows at time t:

Each prosumer posts a *net* flow $x_{nt} \in \mathbb{R}$:

$$x_{nt} = s_{nt} - d_{nt}, \qquad s_{nt} = \max\{x_{nt}, 0\},\ d_{nt} = \max\{-x_{nt}, 0\}.$$

Aggregate supply and demand: $s_t = \sum_n s_{nt},\ d = \sum_n d_{nt}.$

Community surplus or deficit: $s_{At} = \max\{s_t - d_t, 0\},\ d_{At} = \max\{d_t - s_t, 0\}.$

Market & Payments

Grid prices at time t: $\underline{\lambda}_t$ (sell to grid), $\overline{\lambda}_t$ (buy from grid), with $\overline{\lambda}_t > \underline{\lambda}_t$.

We design a market within the prosumer community for a given time-step t (e.g., in a 15-minute time interval).

Market Definition: A market is $(\boldsymbol{x}, \boldsymbol{P}(\boldsymbol{x}))$, where $\boldsymbol{x} = (x_1, \ldots, x_N)$ and $\boldsymbol{P} : \mathbb{R}^N \to \mathbb{R}^N$.

Payment function: $P_n = R_n - C_n$ (revenues for selling, costs for buying).

Marginal prices:

$$r(s, d) \triangleq \partial R_n / \partial s_n, \qquad c(s, d) \triangleq \partial C_n / \partial d_n.$$

Axioms

1. **Anonymity / Fairness:** order-agnostic; identical actions $\Rightarrow$ identical terms.
2. **Coalition-proofness:** no group gains by pooling/splitting; $\Rightarrow$ linear in own s_n, d_n.
3. **Individual-rationality:** internal trades happen *within* $[\underline{\lambda}, \overline{\lambda}]$ $\Rightarrow$ concentrated liquidity.
4. **Budget-balance:** $\sum C \geq \sum R$ (optionally exact).
5. **No-arbitrage:** $r(s, d) \leq c(s, d)$ for all (s, d).
6. **Monotonicity & Responsiveness:** prices move in the "right" directions (peak shaving).
7. **Homogeneity:** price is scale-invariant; depends on s/d ratio.

Anonymity $\Rightarrow$ batch clearing

Definition (Anonymity). For any permutation π of other agents:

$$P_n(x_n, x_{-n}) = P_n(x_n, \pi(x_{-n})).$$

Implications

Order books / bilateral matching violate anonymity.

AMM with **pooled state** and **session-level batching** satisfies it.

Aggregation-based pricing satisfies anonymity: $P_n = \Psi(x_n, s, d)$.

Transfer depends only on (x_n, s, d).

Coalition-proofness ⇔ Linear individual terms

Claim. Coalition-proofness holds *iff*

$$P_n(x) = s_n\, r(s, d) - d_n\, c(s, d),$$

i.e., **linear** in s_n, d_n with common r, c that depend only on aggregates.

Intuition:

Linear sharing rules are neutral with respect to grouping/splitting.

Nonlinear per-agent terms create incentives to merge/split.

Individual Rationality ⇒ Concentrated Liquidity

Within-spread trading

$$\begin{aligned} s < d: &\quad r(s,d) > \underline{\lambda}, \qquad s \geq d: \; r(s,d) \leq \underline{\lambda}, \\ s \leq d: &\quad c(s,d) \geq \overline{\lambda}, \quad s > d: \; c(s,d) < \overline{\lambda}. \end{aligned}$$

Consequence: AMM must only quote prices inside $[\underline{\lambda}, \overline{\lambda}]$ ⇒ *concentrated liquidity* interval.

Budget balance, No-arbitrage, Responsiveness & Homogeneity

Budget balance: $\sum C \geq \sum R$ (fees optional).

No-arbitrage: $r(s, d) \leq c(s, d)$.

Responsiveness (peak-shaving):

$$\frac{\partial c}{\partial s} < 0, \ \frac{\partial c}{\partial d} > 0, \ \frac{\partial r}{\partial d} > 0, \ \frac{\partial r}{\partial s} \leq 0.$$

Homogeneity (scale-invariance): $r(\alpha s, \alpha d) = r(s, d), \ c(\alpha s, \alpha d) = c(s, d)$ for all $\alpha > 0$.

$$\iff r(s/d) \equiv r(s, d), \ c(s/d) \equiv c(s, d)$$

Axiom-Compliant Price Functions: Summary

Let the Supply-to-Demand ratio (SDR) be

$$y \triangleq s/d.$$

These conditions are necessary and sufficient for an axiom-compliant AMM:

1. **Grid price bounds.**
$$\underline{\lambda} \;\le\; r(y) \;\le\; c(y) \;\le\; \overline{\lambda}.$$
2. **Local price bounds.**
$$c(y) \;\ge\; y\, r(y).$$
3. **Decreasing in y.**
$$r'(y) < 0 \quad \text{and} \quad c'(y) < 0.$$

Roadmap

From invariant to price: re-anchored CFMM

Trading function (session t): $\psi_t(E, M) = K_t$, with state $(E, M) = $ (energy, money).

Price of locally-traded power:

$$\rho_t(E, M) \;=\; \left| \frac{\partial_E \psi_t}{\partial_M \psi_t} \right| .$$

Re-anchoring: choose K_t so initial point $(E, M) = (s_t, 0)$ sits on $\psi_t = K_t$, and $\pi \in [\underline{\lambda}_t, \overline{\lambda}_t]$.

Batch session:

1. Pool is "charged" with supply s_t.
2. All internal trades clear on ψ_t (order-agnostic).
3. Residual imbalances $s_t - d_t$ settled with grid at $\underline{\lambda}_t / \overline{\lambda}_t$.

AMM Workflow (session t)

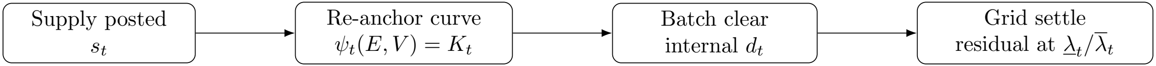

Payments: three imbalance regimes

Let $M_t(s, d)$ be the pool's monetary value after internal trades.

Proportional repartition ensures exact budget-balance and coalition-proofness:

Case I: $d = s$ (balanced)

$$c = r = \frac{M_t(s, s)}{s}.$$

Case II: $d < s$ (excess supply)

$$c = \frac{M_t(s, d)}{d},$$

$$r = \frac{M_t(s, d) + \underline{\lambda}_t(s - d)}{s}.$$

Case III: $d > s$ (excess demand)

$$c = \frac{M_t(s, s) + \overline{\lambda}_t(d - s)}{d},$$

$$r = \frac{M_t(s, s)}{s}.$$

AMM geometry I

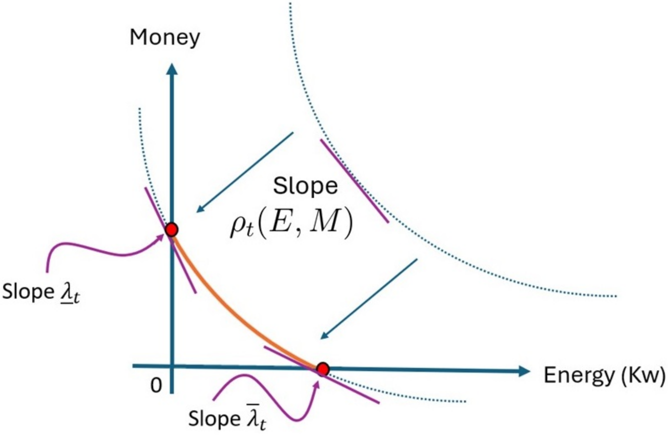

Bonding curve supported only between the lower/upper price bounds λ_t and $\bar{\lambda}_t$.

AMM geometry II

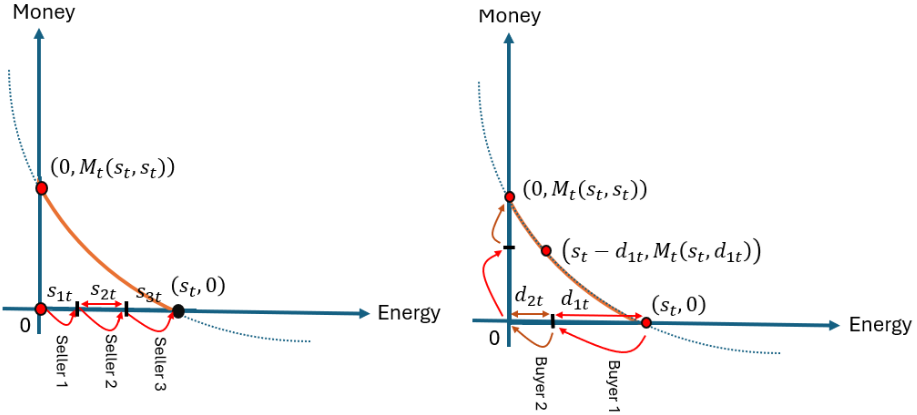

(Left) pool re-anchoring with initial supply; (Right) depletion of capacity in exchange for money

AMM geometry III

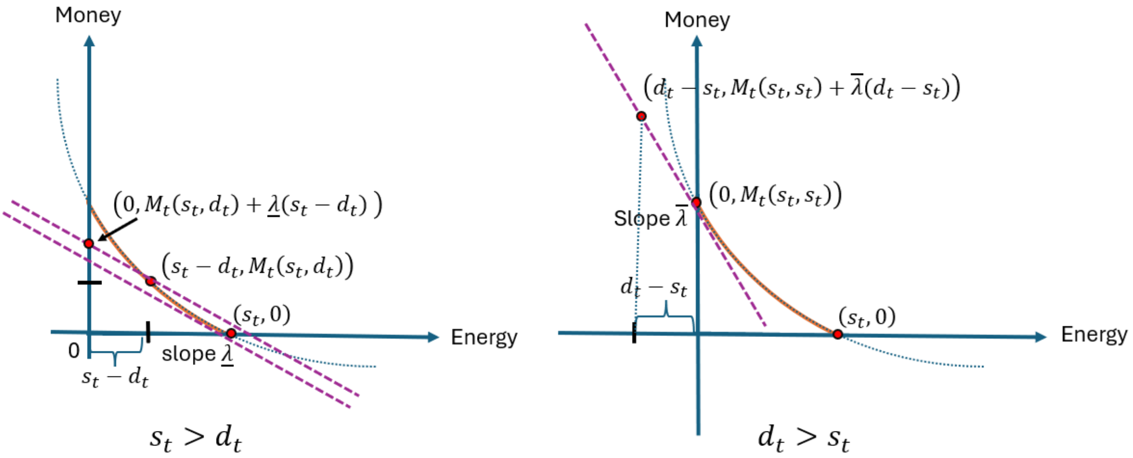

Managing Surpluses (Left) and Deficits (Right) via Tangent-Line Shifts

Bonding curves

Linear $\psi_t(M,E) = M + \lambda_t E$

$$K_t = \lambda_t s_t, \qquad M_t(E, \Delta E) = \lambda_t \, \Delta E.$$

Induced prices (with bounds):

$$c_t(s,d) = \lambda_t + (\overline{\lambda}_t - \lambda_t)\,(1 - s/d)^+, \qquad r_t(s,d) \quad = \lambda_t - (\lambda_t - \underline{\lambda}_t)\,(1 - d/s)^+.$$

Anchors for λ_t:

Mid Market Rate (MMR): $\lambda_t^{\text{arith}} = (\underline{\lambda}_t + \overline{\lambda}_t)/2$, Geometric Market Rate (GMR): $\lambda_t^{\text{geom}} = \sqrt{\underline{\lambda}_t \, \overline{\lambda}_t}$.

Hyperbolic (Uniswap-style, re-anchored)

$$\psi_t(M,E) = \left(M + \kappa_t \sqrt{\underline{\lambda}_t}\right)\left(E + \kappa_t / \sqrt{\overline{\lambda}_t}\right) = \kappa_t^2. \qquad \kappa_t := s_t \, \frac{\sqrt{\overline{\lambda}_t \, \underline{\lambda}_t}}{\sqrt{\overline{\lambda}_t} - \sqrt{\underline{\lambda}_t}},$$

Axiom Compliance via Geometric Construction

Proposition

Let the AMM use batch execution and proportional payments, with liquidity concentrated *strictly* within the grid bounds, $\rho_t(E, M) \in (\underline{\lambda}_t, \overline{\lambda}_t)$. IF the trading function $\psi(E, M)$ is:

(i) strictly increasing in E and M, (ii) homothetic, (iii) quasi-concave,

THEN, the induced market satisfies all the desired axioms.

Intuition (design $\Rightarrow$ axioms).

- **Batch + proportional payments** $\Rightarrow$ Anonymity, Coalition-Proofness.
- **Concentrated liquidity (within bounds) + re-anchoring** $\Rightarrow$ IR, No-Arbitrage, Budget-Balance (internal trades strictly better than grid).
- ψ **quasi-concave** $\Rightarrow$ convex iso-value curves $\Rightarrow$ Responsiveness (peak-shaving).
- ψ **homothetic** $\Rightarrow$ prices depend on reserve ratios $\Rightarrow$ Homogeneity.
- ψ **strictly increasing** $\Rightarrow$ positive marginal values $\Rightarrow$ Monotonicity.

Standard properties for DeFi AMMs (Angeris et al, Milionis et al, Fabi and Prat etc.) gives desired behavior of energy AMM

Roadmap

Prosumer model

How shall a prosumer interact with the AMM?

We consider a single prosumer n with batteries and local generation technologies (solar panels, wind turbines).

Time: Epochs $e = 1, 2, \ldots$; each epoch has T steps of length Δ
($\Delta = 0.25$ if sessions last 15min and prices are quoted in kWh).

Prices (price-taking): $(\bar{r}_t, \bar{c}_t)$ with $\underline{\lambda}_t \leq \bar{r}_t \leq \bar{c}_t \leq \overline{\lambda}_t$.

State at start: $b_{n0}(e) \in [0, B_n]$: battery State of Charge (SoC).

Control variables:
s_{nt}: sell energy,
d_{nt}: buy energy,
k_{nt}: charge(+)/dischage(-) battery,
p_{nt}: consume energy.

Best-Response Linear Program (LP)

The optimization problem of a prosumer can be stated as an (epoch-based) Bellman equation:
Within-epoch LP (Objective)

$$\max_{\mathbf{s},\mathbf{d},\mathbf{k},\mathbf{p},\mathbf{b}} \sum_{t=1}^{T} (\bar{r}_t s_{nt} - \bar{c}_t d_{nt})\Delta + \gamma \bar{\Pi}_n(b_{nT})$$

Subject to the following constraints:

Battery charge constraints

$$b_{nt} = b_{n,t-1}(e) + k_{nt}\Delta \quad \text{(kWh)}$$
$$0 \leq b_{nt} \leq B_n \quad \text{(kWh)}$$
$$-K_n \leq k_{nt} \leq K_n \quad \text{(kW)}$$

Resource constraint (ω: power production)

$$\omega_{nt} = p_{nt} + k_{nt} + s_{nt} + d_{nt} \quad \text{(kW)}$$

Power consumption constraints

$$p_{nt} \geq \alpha_{nt}^{\text{base}} \quad \text{(kW)}$$
$$\textstyle\sum_{t=1}^{T} (p_{nt} - \alpha_{nt}^{\text{base}})\Delta = \alpha_n^{\text{flex}} \quad \text{(kWh)}$$

Demand and supply constraints

$$0 \leq s_{nt} \leq X_n \quad \text{(kW)}$$
$$0 \leq d_{nt} \leq X_n \quad \text{(kW)}$$

Optimal plan & shadow price

Linear program $\Rightarrow$ optimal plan is an extreme-point of the action set.

Let θ_{nt} be the multiplier on power balance:

$$\theta^*_{nt} \;=\; (\mu^*_{nt} - \nu^*_n) \;=\; \gamma(\pi^b_n)^* \;+\; \sum_{\tau=t}^{T}(\iota^{L*}_{n\tau} - \iota^{U*}_{n\tau}).$$

Interpretation: internal marginal value of energy $=$ marginal utility of consumption $=$ discounted continuation value $+$ intra-epoch storage shadow value.

Trading rule (price-taking):

$$x^*_{nt} = \begin{cases} X_n, & \theta^*_{nt} < \bar{r}_t \\ \in (0, X_n), & \theta^*_{nt} = \bar{r}_t \\ 0, & \bar{r}_t < \theta^*_{nt} < \bar{c}_t \\ \in (-X_n, 0), & \theta^*_{nt} = \bar{c}_t \\ -X_n, & \theta^*_{nt} > \bar{c}_t \end{cases} \qquad (x_{nt} = s_{nt} - d_{nt})$$

Transmission-constrained at bounds; internally constrained on equalities.

Rolling-horizon MPC

To approximate the solution to the dynamic programming problem, proceed as follows:

Loop each epoch e:

1. Fix $b_{n0}(e)$ and forecast $(\bar{\mathbf{r}}, \bar{\mathbf{c}})$ over L epochs.

2. Solve L-epoch LP $\Rightarrow$ plan $\mathbf{z}^{*(L)}$; increase L until terminal b stabilizes.

3. Implement only epoch-e actions (s^*, d^*, k^*, p^*); set $b_{n0}(e+1) = b_{nT}(e)$.

Note: small L is myopic; larger L captures storage option value. Converges quickly in practice.

Rolling-horizon MPC converges to the DP solution for sufficiently large L (Prat 2024)

Equilibrium Problem (price-mediated coordination game)

How shall a community of prosumers interact with the AMM?

Players & Information

Prosumers $n = 1, \ldots, N$ (N large).

Public state z_e (weather, forecasts) $\Rightarrow$ correlated types.

Private state: $(\theta_n, b_{n0}(e))$; price-taking.

Timeline per epoch e

1. Observe z_e; form beliefs $F(\boldsymbol{\theta} \mid z_e)$.
2. Choose mixed strategy over *LP extreme points*.
3. AMM sets session prices $(\bar{\boldsymbol{r}}, \bar{\boldsymbol{c}})$ via aggregate net order flow.

Actions and Strategies

Action Set ($\mathcal{A}_n$): The set of all feasible epoch-plans $\mathbf{a}_n$ for prosumer n. This is a convex set defined by the LP constraints.

Because the prosumer's best-response problem is an LP, any optimal *pure* strategy $\mathbf{a}_n^*$ will be an **extreme point**, $\mathbf{a}_n^* \in \text{ext}(\mathcal{A}_n)$.

Mixed Strategy (σ_n): A probability distribution over this support. Formally, it's a map from the prosumer's type θ_n to a distribution over their actions:

$$\sigma_n : \Theta_n \to \Delta(\text{ext}(\mathcal{A}_n))$$

where $\Delta(\cdot)$ denotes the simplex formed over over the set $\text{ext}(\mathcal{A}_n)$.

Mixed Strategy Profile: $\boldsymbol{\sigma} \in \bigtimes_{n=1}^{N} \Delta(\text{ext}(\mathcal{A}_n))$

Game Form

Stage game (epoch) $\Rightarrow$ BNE mixing over extreme points.

Dynamic game $\Rightarrow$ MPE across epochs.

Stage-game BNE (compact definition)

Given z_e, a pair $(\boldsymbol{\sigma}^*, \bar{\boldsymbol{p}}^*)$ with $\bar{\boldsymbol{p}}^* = (\bar{\boldsymbol{r}}^*, \bar{\boldsymbol{c}}^*)$ is a Bayes-Nash Equilibrium (BNE) if:

Individual Optimality: For each n, $\sigma_n^*(\cdot \mid \theta_n)$ must place all probability on the best-performing actions:

$$\sigma_n^*(\cdot \mid \theta_n) \in \arg\max_{\sigma_n} \sum_{\mathbf{a}_n} \sigma_n(\mathbf{a}_n)\, \pi_n(\mathbf{a}_n; \bar{\boldsymbol{p}}^*)$$

where $\mathbf{a}_n$ ranges over the extreme points of the feasible set $\mathcal{A}_n$.

Rational Expectations: The prices $\bar{\boldsymbol{p}}^*$ must be the ones induced (in expectation over $F(\boldsymbol{\theta} \mid z_e)$) by the aggregate behavior under $\boldsymbol{\sigma}^*$.

A mixed strategy is an equilibrium only if if the agent is indifferent between all pure strategies in its support.

Equilibrium characterization

Proposition (ex-ante welfare equivalence)

Under *Budget-Balance*, *Individual Rationality*, and *Responsiveness*, a stage-game BNE maximizes the expected value of grid-trade profits:

$$\max_{\boldsymbol{\sigma}} \; \mathbb{E}\Big[\sum_{t=1}^{T} \big(\underline{\lambda}_t \, s_{At} - \overline{\lambda}_t \, d_{At}\big)\Big]$$

That is, the equilibrium *minimizes* a (possibly weighted) L_1 norm of net trades with the grid.

Intuition

The AMM first matches peers inside the spread.

Only residual net trades clear against the grid at $(\underline{\lambda}_t, \overline{\lambda}_t)$.

The equilibrium problem is therefore: "make the residuals as small and as valuable as possible."

Equilibrium characterization

Proposition (ex-ante welfare equivalence)

Under *Budget-Balance*, *Individual Rationality*, and *Responsiveness*, a stage-game BNE maximizes the expected value of grid-trade profits:

$$\max_{\boldsymbol{\sigma}} \ \mathbb{E}\Big[\sum_{t=1}^{T} \big(\underline{\lambda}_t \, s_{At} - \overline{\lambda}_t \, d_{At}\big)\Big]$$

That is, the equilibrium *minimizes* a (possibly weighted) L_1 norm of net trades with the grid.

Implication

The game can be solved as a Planner Problem (PP) and then decentralized over feasible allocations .

Equilibrium characterization

Proposition (ex-ante welfare equivalence)

Under *Budget-Balance*, *Individual Rationality*, and *Responsiveness*, a stage-game BNE maximizes the expected value of grid-trade profits:

$$\max_{\boldsymbol{\sigma}} \ \mathbb{E}\Big[\sum_{t=1}^{T}\big(\underline{\lambda}_t\, s_{At} - \overline{\lambda}_t\, d_{At}\big)\Big]$$

That is, the equilibrium *minimizes* a (possibly weighted) L_1 norm of net trades with the grid.

Mean-Field View

With many price-taking prosumers and public z_e, actions are conditionally i.i.d.

The welfare of the realized average net flow converges to welfare at the average community net-flow.

The BNE is characterized by the mean-field fixed point.

Mechanism comparison: AMM vs VCG

The **ex-post** optimal welfare (Planner) is an upper bound on the **ex-ante** welfare (AMM):

$$\underbrace{\mathbb{E}_\theta\left[\max_\sigma W(\sigma,\theta)\right]}_{\mathcal{W}_{\text{Planner (VCG)}}} \geq \underbrace{\max_\sigma \mathbb{E}_\theta[W(\sigma,\theta)]}_{\mathcal{W}_{\text{AMM (ex-ante)}}}$$

VCG (first-best, ex-post)

Strategy-proof; uses full types.

Achieves $\mathcal{W}_{\text{Planner}}$.

Not budget-balanced in general (requires subsidies).

Centralized, requires full data revelation *before* trading.

AMM (this paper, ex-ante)

Budget-balanced; types obfuscated.

Implements BNE that maximizes expected welfare under our axioms.

Not strategy-proof (Bayesian Incentive Compatible).

Decentralized, payments determined *after* trading.

Trade-off:

AMM sacrifices ex-post efficiency for implementability (budget balance, lower communication complexity), yet attains the best *ex-ante* welfare possible within our market axioms.

Equilibrium computation

Goal: Approximate BNE/MPE under price-taking and mean-field aggregation.

Computation Pipeline

1. **State-dependent bins:** Create J representative types $(\theta_j, b_j(e))$ by clustering agents (by hardware + current SoC).

2. **Strategy banks $\mathcal{B}_j^L(e)$:** For each type j, solve L-epoch best-response LPs on synthetic prices; store the resulting extreme point plans.

3. **Planner QP (mean-field):** Find the mixed strategies $\{\sigma_j\}$ over the banks that maximize ex-ante welfare (with small Tikhonov regularization).

4. **Decentralize & project:** Agents sample plans from their bank $\mathcal{B}_j^L(e)$ using σ_j, project to their *exact* individual constraints, and implement the first step.

Rolling Horizon

The terminal battery state $b_{nT}(e)$ is carried over to the next epoch $e + 1$.

Planner Problem

Let $\bar{\mathbf{y}}^{(l)}(\boldsymbol{\sigma})$ be the expected net trades for horizon step l.

This is a linear aggregation of the mixed strategies σ_j for all J types, weighted by their population mass w_j:

$$\bar{\mathbf{y}}^{(l)}(\boldsymbol{\sigma}) = \sum_{j=1}^{J} w_j \big(S_j^{(l)} \sigma_j - D_j^{(l)} \sigma_j \big)$$

We define exports $\mathbf{y}^+ \geq 0$ and imports $\mathbf{y}^- \geq 0$ such that $\bar{\mathbf{y}} = \mathbf{y}^- - \mathbf{y}^+$.

Objective Function and Constraints (Regularized Welfare):

$$\max_{\boldsymbol{\sigma}, \mathbf{y}^+, \mathbf{y}^-} \sum_{l=1}^{L} \gamma^{h-1} \big(\underline{\boldsymbol{\lambda}}^{(l)\top} \mathbf{y}^{+(l)} - \overline{\boldsymbol{\lambda}}^{(l)\top} \mathbf{y}^{-(l)} \big) \; - \; \lambda \|\Delta \mathbf{y}(\boldsymbol{\sigma})\|_2^2$$

Net trade balance: $\bar{\mathbf{y}}(\boldsymbol{\sigma}) - \mathbf{y}^+ + \mathbf{y}^- = 0$, **Non-negativity:** $\mathbf{y}^+, \mathbf{y}^- \geq 0$

Probability simplex: $\mathbf{1}^\top \sigma_j = 1, \; \sigma_j \geq 0, \;\; \forall j$

Decentralization & feasibility projection

Implementation Steps

1. Agent n finds its representative bin j.
2. It draws one L-epoch plan $\mathbf{a}_{nk}$ from the bank $\mathcal{B}_j^L(e)$ according to the optimal strategy $\sigma_j^\star(e)$.
3. **Projection:** The agent projects the first-epoch plan $\mathbf{a}_{nk}^{(1)}$ onto their *true* individual constraints:

$$\widetilde{\mathbf{a}}_n^{(1)} = \mathrm{Prj}\left(\mathbf{a}_{nk}^{(1)}; \theta_n, b_{n0}(e)\right).$$

4. Implement $\widetilde{\mathbf{a}}_n^{(1)}$ and update the battery state $b_{n0}(e+1) = b_{nT}(e)$.

Remarks

Projection guarantees feasibility for every agent.

Incentive compatibility is *approximate* (due to binning error).

$L=1$ gives a myopic stage-game; $L>1$ captures the option value of storage across epochs.

Roadmap

Weather data: France

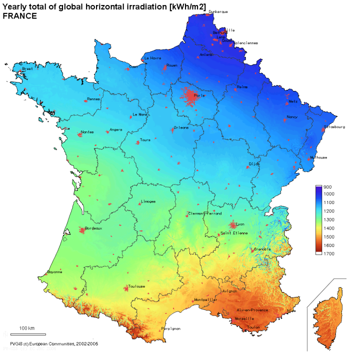

(a)

(b)

Figure: Weather characteristics in metropolitan France. a. Solar irradiation. b. Wind speed at an altitude of 100 meters.

Prosumer Data: Paris and Nice

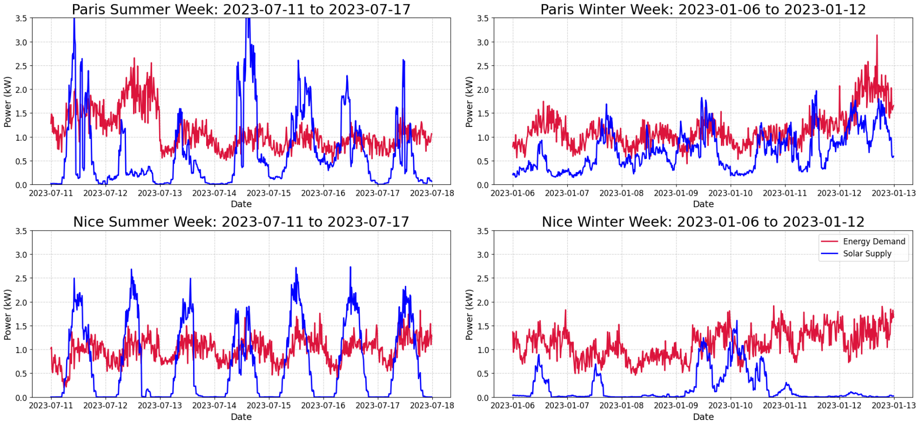

Figure: Community household energy profiles for demand and solar supply for prosumers in Paris and Nice during a summer week (07/10/2023 to 07/16/2023) on the left side, and a winter week (01/06/2023 to 01/12/2023) on the right side. Red lines indicate community demand, while blue and black lines represent solar and Eolic energy supply respectively.

Prosumer Data: Paris and Nice

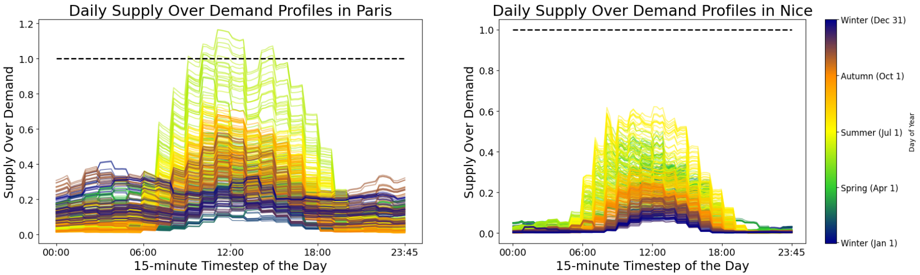

Figure: Yearly supply over demand for Paris and Nice, respectively left and right graphs. Each line represents the average over two weeks and their colors indicate the season. A dashed line represents the equality between supply and demand.

Prosumer Data: Paris and Nice

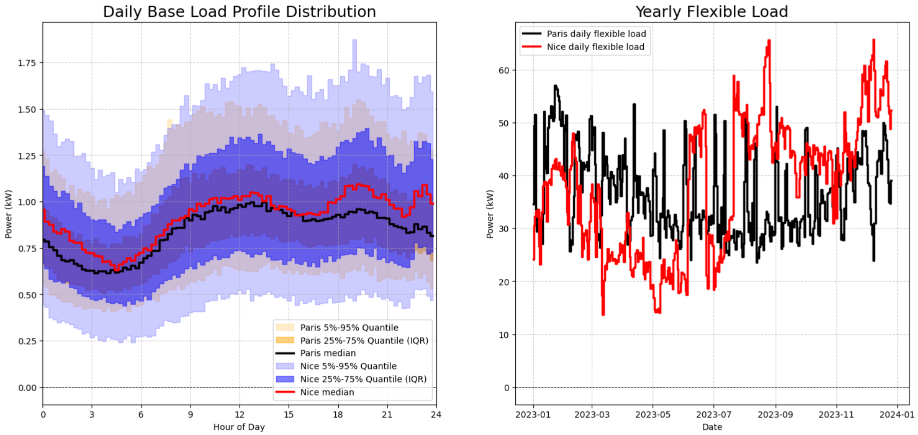

Figure: Consumption profile. On the left side, daily base load profile distribution in Nice, line in red with blue quantiles, and in Paris, line in black with orange quantiles. On the right side, yearly flexible load for Nice in red and Paris in black.

Optimized Prosumer Behavior: Paris and Nice

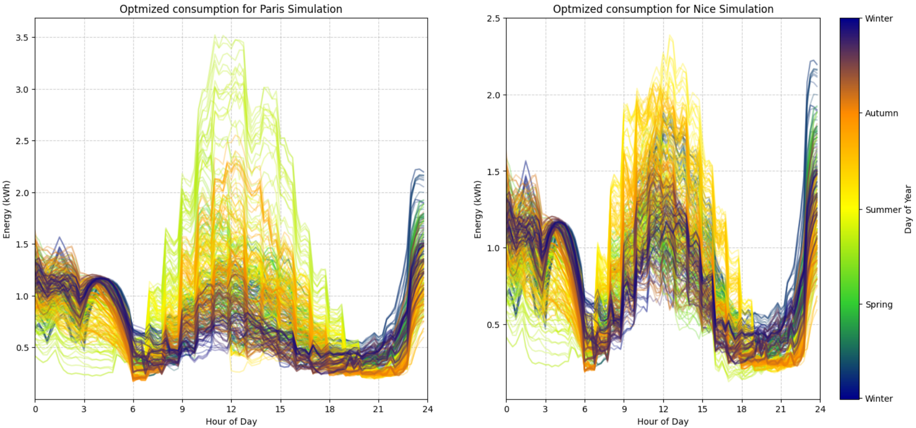

Figure: Optimized consumption for Paris, on the right side, and Nice, on the left side.

Optimized Prosumer Behavior: Paris and Nice

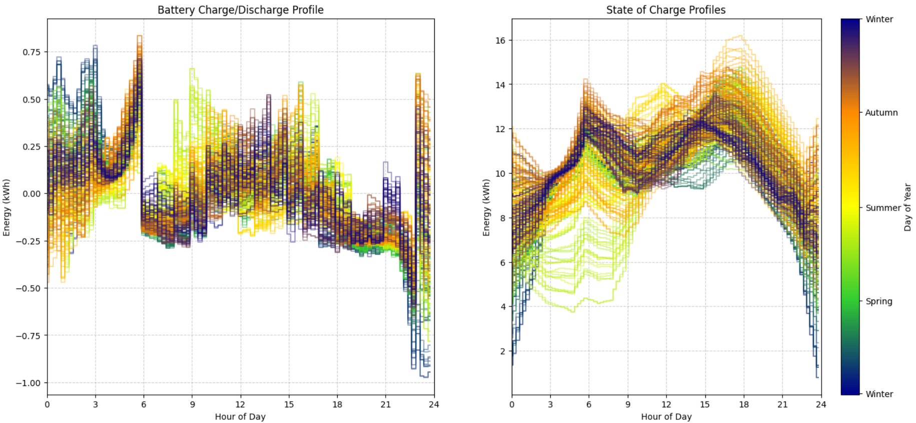

Figure: Two weeks mean battery profiles for Paris (left) and Nice (right). Colors represent annual seasons.

Gains from AMM use: Paris and Nice

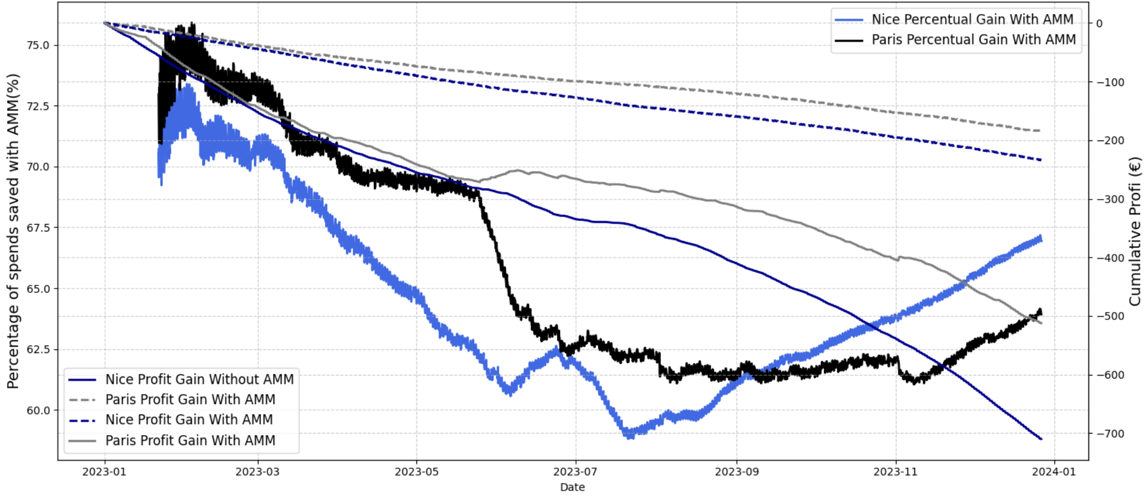

Figure: Cumulative gain using linear pricing function for Paris on the left side and Nice on the right side.

Equilibrium Computation & Synthetic Profiles

Objective: Model grid-level gains from decentralization within a community facilitated by an Automated Market Maker (AMM).

Agent Composition (1,000 Total):

- 30% Solar Prosumers.
- 30% Wind Prosumers.
- 40% Pure Consumers.

Setup: Agents possess batteries, base load, and flexible load.

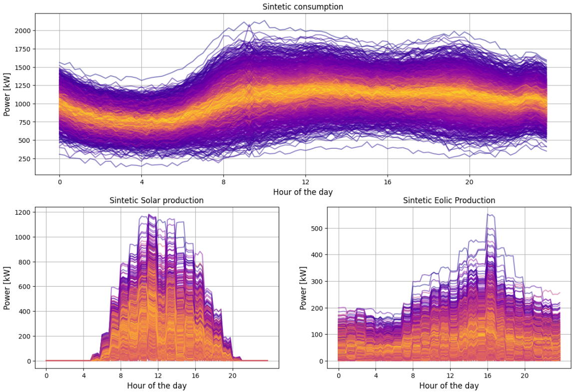

Figure: Synthetic daily curves (Summer) based on Paris data. Colors represent depth (centrality).

Representative Profiles:

The center of each bin is chosen as the profile with depth closest to the center depth.

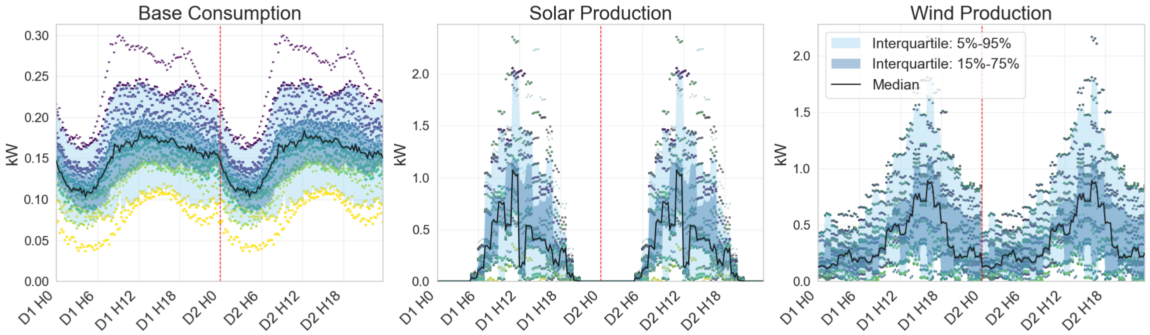

Figure: Profile distribution for base supply and demand.

Mixed Strategy

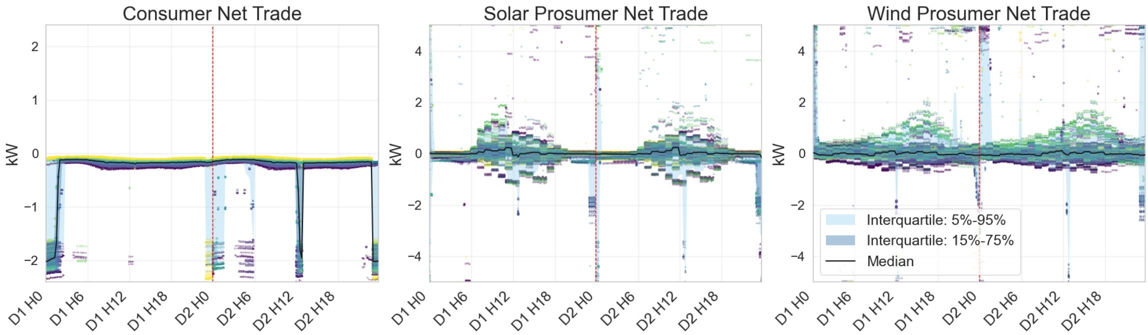

Figure: Net trading profiles: Consumers (left), Solar (middle), Wind (right).

Gains from Trade

Comparison: Green bars: Profits with AMM. Red bars: Benchmark profits (fixed prices).

The dynamic equilibrium approach yields a **42% total gain** for the community compared to the benchmark.

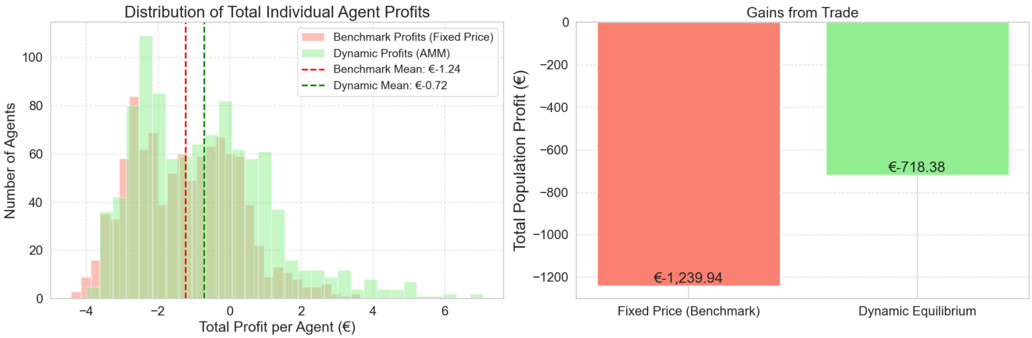

Figure: Distribution of individual agent profits (left) and total gains (right).

AMM Price Dynamics

Dynamic Pricing:

- **Selling Prices (Green):** Generally follow the benchmark trend but fluctuate based on agent effort to match supply/demand.
- **Buying Prices (Red):** Adjusted dynamically via the AMM mechanism.

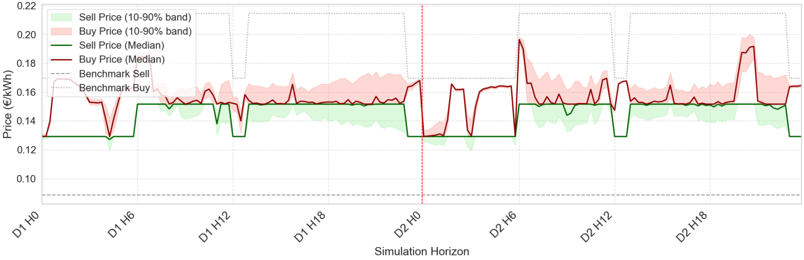

Power Allocation

Flexible Consumption

Dynamic allocation shifts consumption to low-price regions.

Mostly determined by grid prices impacting community buy/sell rates.

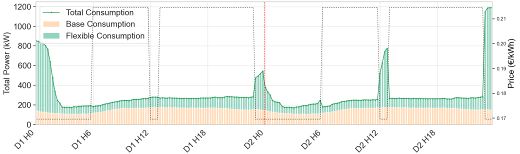

Grid Interaction

Net interaction (blue line) shows the community buying from the grid to match deficits.

Agents prioritize buying from the grid when it is cheaper.

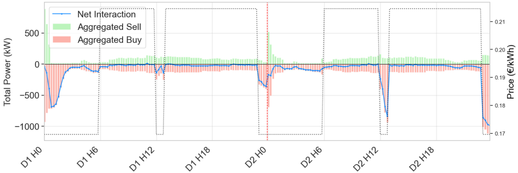

Battery Dynamics

Dispatch Logic:

Battery behavior is highly sensitive to internal community supply and grid prices.

Charging (Blue): Occurs mainly during solar hours or low grid prices.

Discharging (Peach): Occurs when production is low.

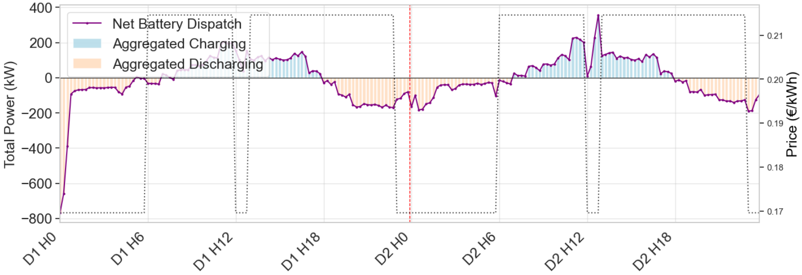

Figure: Aggregate battery dispatch profile.

Conclusion & Future Directions

Summary of Contributions

1. **Axiomatic theory** for P2P energy market.
2. **AMM construction** (batching & concentrated liquidity, standard DeFi properties).
3. Model community as a **Mean-Field Game**

Empirical Validation

Numerical experiments using real prosumer data from the **Paris (IDF) and Nice**.

Future Directions

Forecast market

Interconnected microgrids

continuous-time version of the energy AMM.

Thanks!

References I

Guillermo Angeris, Alex Kao, Ricky Chiang, Tarun Noyes, and Tarun Chitra. Improved price oracles: Constant function market makers. In *Proceedings of the 21st ACM Conference on Economics and Computation*. ACM, 2020.

Guillermo Angeris, David Evans, and Tarun Chitra. The geometry of constant function market makers. *arXiv preprint*, 2023.

Fabien Berger and Myriam Kassoul. Aggregating distributed plug & play battery storage units: effects on local grid dynamics and efficiency. *IET Conference Proceedings*, 2025(14):566–569, 2025. Also presented at CIRED 2025 as Paper 208.

Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

Bennevis Crowley, Jalal Kazempour, and Lesia Mitridati. How can energy communities provide grid services? a dynamic pricing mechanism with budget balance, individual rationality, and fair allocation. *Applied Energy*, 382:125154, 2025. ISSN 0306-2619. doi: https://doi.org/10.1016/j.apenergy.2024.125154. URL https://www.sciencedirect.com/science/article/pii/S0306261924025388.

Josu Doncel, Nicolas Gast, and Bruno Gaujal. Discrete mean field games: Existence of equilibria and convergence. *Journal of Dynamics and Games*, 6(3):221–239, 2019. ISSN 2164-6074. doi: 10.3934/jdg.2019016. URL http://dx.doi.org/10.3934/jdg.2019016.

Tomaso Erseghe. Distributed optimal power flow using admm. *IEEE Transactions on Power Systems*, 29(5):2370–2380, 2014. doi: 10.1109/TPWRS.2014.2306495.

E. Sorin F. Moret, T. Baroche and P. Pinson. Negotiation algorithms for peer-to-peer electricity markets: Computational properties. *IEEE Transactions on Smart Grid*, 15(4):1257–1269, 2024. doi: 10.1109/TSG.2024.3156874.

Michele Fabi and Julien Prat. The economics of constant function market makers. *Journal of Corporate Finance*, 91:102737, 2025. doi: 10.1016/j.jcorpfin.2025.102737.

Olivier Guéant, Jean-Michel Lasry, and Pierre-Louis Lions. Mean field games and applications. In *Paris-Princeton Lectures on Mathematical Finance 2010*, volume 2003 of *Lecture Notes in Mathematics*, pages 205–266. Springer, Berlin, Heidelberg, 2011. ISBN 978-3-642-14660-2. doi: 10.1007/978-3-642-14660-2_3.

Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007. doi: 10.1007/s11537-007-0657-8.

Jason Milionis, Jens Ernstberger, Joseph Bonneau, Scott Duke Kominers, and Tim Roughgarden. Incentive-compatible recovery from manipulated signals, with applications to decentralized physical infrastructure. *arXiv preprint*, 2025.

Dov Monderer and Lloyd Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

References II

Pierre Pinson et al. Peer-to-peer electricity trading: Market design and operational principles. *Energy Systems*, 2020.

Eléa Prat, Richard M. Lusby, Juan Miguel Morales, Salvador Pineda, and Pierre Pinson. How long is long enough? finite-horizon approximation of energy storage scheduling problems. *arXiv preprint*, 2024.

Robert W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973. doi: 10.1007/BF01737559.

Jan Christoph Schlegel, Mateusz Kwaśnicki, and Akaki Mamageishvili. Axioms for constant function market makers. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, pages 1079–1109. ACM, 2023. doi: 10.1145/3580507.3597720.

Tiago Sousa, Tiago Soares, Pierre Pinson, Fabio Moret, Thomas Baroche, and Etienne Sorin. Peer-to-peer and community-based markets: A comprehensive review. *Renewable and Sustainable Energy Reviews*, 104:367–378, 2019. doi: 10.1016/j.rser.2019.01.036.