Machine learning architectures for mean-field games models of price formation

Julián Gutiérrez Pineda

Joint work with Diogo Gomes (Kaust) and Mathieu Laurière (NYU Shanghai)

King Abdullah University of Science and Technology

جامعة الملك عبدالله للعلوم والتقنية

SITE Research Center Seminar

September 2022

- Mean-field games price-formation model
- A posteriori estimates
- Price as Lagrange multiplier
- Neural networks approach
- Numerical results
- Mean-field games price-formation model with common noise

Goal: determine the price ϖ of a commodity with supply Q (known):

$$\varpi = F(Q).$$

Two cases:

- Q is deterministic (standard energy production),

$$dQ(t) = b^s(Q(t), t)dt, \quad t \in [0, T],$$

- Q is stochastic (renewable energy production)

$$dQ(t) = b^s(Q(t), t)dt + \sigma^s(Q(t), t)dW(t), \quad t \in [0, T].$$

$\varpi = F(Q)$, F given:

- Shen and Basar (2011) Stackelberg games.
- Djehiche et al. (2020) Cournot models for stochastic dynamics.
- Alasseur et al. (2020) mean-field equilibria. Alasseur et al. (2021) random jumps.
- Aïd et al. (2021) MFG of optimal stopping.
- MFG for intraday electricity markets Féron et al. (2020), Féron et al. (2021).

Price formation:

- Basar and Srikant (2002) Stackelberg games.
- Market clearing conditions Shrivats et al. (2020).
- Fujii and Takahashi (2020) market clearing conditions. Fujii and Takahashi (2021) major players.
- Aid et al. (2020) Noisy demand forecast.

Gomes and Saúde (2021): Find (u, m, ϖ) solving, on $[0, T] \times \mathbb{R}$,

$$\begin{cases} -u_t + H(x, \varpi(t) + u_x) = 0, \quad u(T, x) = u_T(x) \\ m_t - \left(H_p(x, \varpi(t) + u_x)m(t,x)\right)_x = 0, \quad m(0,x) = m_0(x) \\ \int_{\mathbb{R}} -H_p(x, \varpi(t) + u_x)m(t,x)dx = Q(t). \end{cases}$$

Here,

$u_T,$ the terminal cost,
$m_0 \in \mathcal{P}(\mathbb{R}),$ the initial distribution,
$Q,$ the supply,
$H,$ the Hamiltonian, $H(x,p) = \sup_v \left(-pv - L(x,v)\right),$

are given.

Optimal control derivation: Take ϖ. Control $d\mathbf{x}(t) = v(t)dt$ so as to minimize

$$\int_0^T (L(\mathbf{x}, v) + \varpi v)\, ds + u_T(\mathbf{x}(T)).$$

The value function

$$u(t,x) = \inf_v \int_t^T (L(\mathbf{x}, v) + \varpi v)\, ds + u_T(\mathbf{x}(T))$$

solves $-u_t + H(x, \varpi(t) + u_x) = 0$, and provides the optimal control

$$v^*(t,x) = -H_p(x, \varpi(t) + u_x(t,x)).$$

m_0 evolves under v^* (push-forward) according to m in the transport equation

$$m_t - \left(H_p(x, \varpi(t) + u_x)m(t,x)\right)_x = 0.$$

Check the balance condition

$$\int_{\mathbb{R}} -H_p(x, \varpi(t) + u_x)dm(t, x) = Q(t), \quad t \in [0, T].$$

Existence of a solution (u, m, ϖ) by a fixed point argument: Update

$$\frac{d}{dt}\tilde{\varpi} = f(\varpi, u(\varpi), m(\varpi)).$$

Then,

$$\frac{d}{dt}\varpi = f(\varpi, u, m).$$

The Euler-Lagrange equation of a representative player is

$$\begin{cases} L_x(X(t), v(t)) - \frac{d}{dt}\left(L_v(X(t), v(t)) + \varpi(t)\right) = 0 \quad t \in [0, T], \\ L_v(X(T), v(T)) + \varpi(T) + u_T'\left(X(T)\right) = 0, \end{cases}$$

which characterizes v^* trajectory-wise.

The balance condition characterizes ϖ.

Remark: Equivalently, for $P(t) := -\left(L_v(X(t), v(t)) + \varpi(t)\right)$,

$$\begin{cases} \frac{d}{dt}P(t) = H_x(X(t), P(t) + \varpi(t)) & t \in [0, T], \\ P(T) = u_T'(X(T)), \\ \frac{d}{dt}X(t) = -H_p(X(t), P(t) + \varpi(t)) & t \in [0, T], \\ X_0 = x_0. \end{cases}$$

The Hamiltonian system is an optimal criteria even for stochastic processes.

To solve $\inf_{x\in\mathbb{R}^n} f(x)$ subject to $g(x)=0$, let

$$\mathcal{L}(x,\lambda) = f(x) + \lambda \cdot g(x), \quad x \in \mathbb{R}^n, \lambda \in \mathbb{R}^m.$$

Then, $x^\star$ and $\lambda^\star$ are characterized by the problems

$$\inf_{x\in\mathbb{R}^n} \sup_{\lambda\in\mathbb{R}^m} \mathcal{L}(x,\lambda) \quad \text{and} \quad \sup_{\lambda\in\mathbb{R}^m} \inf_{x\in\mathbb{R}^n} \mathcal{L}(x,\lambda).$$

Dual ascent method:

$$x^{j+1} = \text{argmin}_{x\in\mathbb{R}^n} \mathrm{L}(x,\lambda^j), \quad \lambda^{j+1} = \lambda^j + h_j g(x^{j+1}), \quad h_j > 0.$$

Arrow-Hurwicz-Uzawa method (Ryu and Boyd (2015)):

$$x^{j+1} = x^j - h_j \left(f_x(x^j) + \lambda^j g_x(x^j)\right), \quad \lambda^{j+1} = \lambda^j + h_j g(x^{j+1}).$$

Consider $v = v(t, x)$.
Given $\tilde{\varpi} : [0, T] \to \mathbb{R}$ and $\tilde{v} : [0, T] \times \mathbb{R} \to \mathbb{R}$, let $\tilde{m} : [0, T] \times \mathbb{R} \to \mathbb{R}$ solve

$$\begin{cases} \tilde{m}_t + (\tilde{v}(t, x)\tilde{m})_x = 0 & [0, T] \times \mathbb{R}, \\ \tilde{m}(0, x) = m_0(x) & x \in \mathbb{R}. \end{cases}$$

Let

$$\mathcal{I}(t) = \tilde{\varpi}(t) \left(\int_{\mathbb{R}} \tilde{v}(t, x)\tilde{m}(t, x)\mathrm{d}x - Q(t) \right), \quad t \in [0, T].$$

$\mathcal{I} \equiv 0$ for (u, m, ϖ).

Notice: We encode (u, m, ϖ) using (ϖ, v^*).

For N players,

$$\tilde{m}^N(t,x) = \frac{1}{N}\sum_{i=1}^{N} \delta_{X^i(t)}(x), \quad t \in [0,T],\ x \in \mathbb{R},$$

$$\mathcal{I}^N(t) = \tilde{\varpi}^N(t) \left(\frac{1}{N}\sum_{i=1}^{N} v^i(t, X^i(t)) - Q(t) \right), \quad t \in [0,T].$$

We consider

$$\sup_{\tilde{\varpi}^N} \inf_{\mathbf{v}} \underbrace{\frac{1}{N}\sum_{i=1}^{N} \int_0^T \left(L(X^i(t), v^i(t)) + \tilde{\varpi}^N(t) \left(v^i(t) - Q(t) \right) \right) dt + u_T \left(X^i(T) \right)}_{\mathcal{L}(\mathbf{v}, \tilde{\varpi}^N)},$$

Recurrent neural networks:

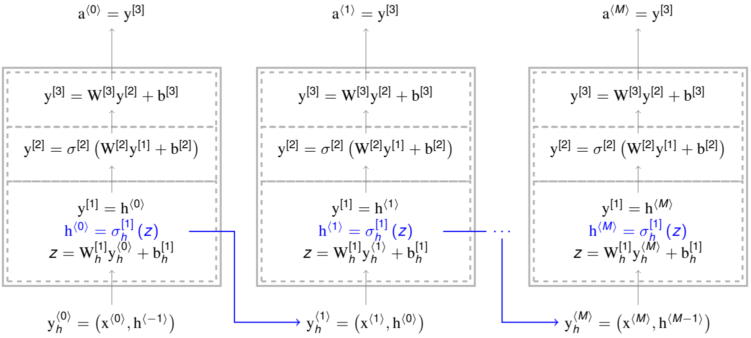

$$\text{Input } x = \left(x^{\langle 0\rangle}, \dots, x^{\langle M\rangle}\right). \text{ Output } a = \left(a^{\langle 0\rangle}, \dots, a^{\langle M\rangle}\right).$$

$$\Theta = \left(W^{[1]}, b^{[1]}, \dots, W^{[L]}, b^{[L]}\right). \quad \text{Training = Optimize } \mathcal{L}(\Theta).$$

Why recurrent networks?
Consider,

$$\overline{X}(t) = \int_{\mathbb{R}} x m(t,x) dx, \quad t \in [0,T],$$

where m is the push-forward of m_0 under v^*. Then, $\frac{d}{dt}\overline{X} = Q$; that is,

$$\overline{X}(t) = \int_{\mathbb{R}} x m_0(x) dx + \int_0^t Q(s) ds, \quad t \in [0,T].$$

Although m is infinite-dimensional, we recover its mean from the path of Q.

Adversarial-like training:

$$\mathrm{v}(\Theta_v) \quad \text{and} \quad \varpi(\Theta_\varpi).$$

Hyper-parameters:
$L_\varpi = L_v = 3$ (shallow),
$\sigma^{[1]} = \sigma^{[2]} =$ sigmoid, $\sigma^{[3]} =$ Id.
The dimension of y is 32 and that of h is 16.

$$\mathcal{L}\left(\Theta_v, \Theta_\varpi\right) = \frac{1}{N}\sum_{i=1}^{N}\left(\sum_{k=0}^{M-1} h\Big(L(X^{(i)\langle k\rangle}, \mathrm{v}^{(i)\langle k\rangle}(\Theta_v)) + \varpi^{\langle k\rangle}(\Theta_\varpi)\left(\mathrm{v}^{(i)\langle k\rangle}(\Theta_v) - Q^{\langle k\rangle}\right)\Big) + u_T(X^{(i)\langle M\rangle})\right)$$

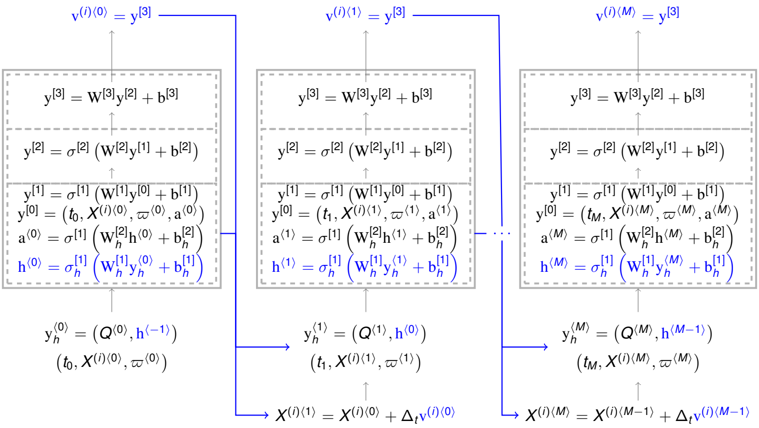

Iteration of the recurrent neural network $v(\Theta_v)$.

Inputs $t = (t_0, \ldots, t_M)$, $X^{(i)\langle 0\rangle}$, $\varpi = \left(\varpi^{\langle 0\rangle}, \ldots, \varpi^{\langle M\rangle}\right)$, $Q = \left(Q^{\langle 0\rangle}, \ldots, Q^{\langle M\rangle}\right)$.
Outputs $v^{(i)} = \left(v^{(i)\langle 0\rangle}, \ldots, v^{(i)\langle M\rangle}\right)$.

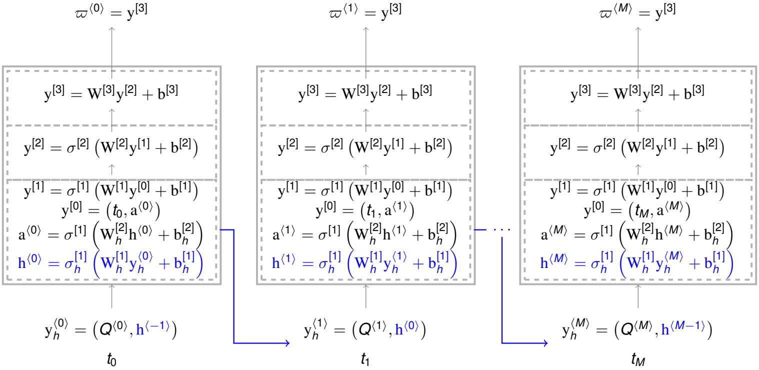

Iteration of the recurrent neural network $\varpi(\Theta_\varpi)$.

Inputs $t = (t_0, \ldots, t_M)$, $Q = \left(Q^{\langle 0 \rangle}, \ldots, Q^{\langle M \rangle} \right)$.
Outputs $\varpi = \left(\varpi^{\langle 0 \rangle}, \ldots, \varpi^{\langle M \rangle} \right)$.

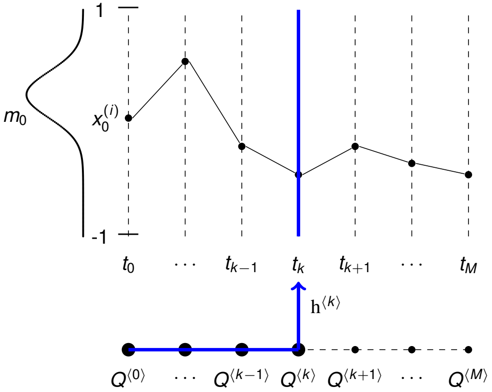

Space-time grid computation diagram.

Notice: The time grid is discrete, while the state grid is continuous.

Neural networks approach

Input : initial density m_0, number of samples N, number of time steps M, supply $Q = (Q(t_k))_{k=0}^{M}$, training steps J

Initialize $\Theta_v^1, \Theta_\varpi^1$;
for $j = 1, \dots, J$ **do**
  sample $x_0^1, \dots, x_0^N$ according to m_0;
  compute $\left(\varpi^{\langle 0 \rangle}, \dots, \varpi^{\langle M \rangle}\right)$ using $\varpi(\Theta_\varpi^j)$;
  for $i = 1, \dots, N$ **do**
    compute $\left(\mathrm{v}^{(i)\langle 0 \rangle}, \dots, \mathrm{v}^{(i)\langle M \rangle}\right)$ using $v(\Theta_v^j)$;
  end
  compute $\mathcal{L}(\Theta_v^j, \Theta_\varpi^j)$;
  compute Θ_v^{j+1} by updating Θ_v^j in the descent direction $\mathcal{L}_{\Theta_v}(\Theta_v^j, \Theta_\varpi^j)$;
  for $i = 1, \dots, N$ **do**
    compute $\left(\mathrm{v}^{(i)\langle 0 \rangle}, \dots, \mathrm{v}^{(i)\langle M \rangle}\right)$ using $v(\Theta_v^{j+1})$;
  end
  compute $\mathcal{L}(\Theta_v^{j+1}, \Theta_\varpi^j)$;
  compute Θ_ϖ^{j+1} by updating Θ_ϖ^j in the ascent direction $\mathcal{L}_{\Theta_\varpi}(\Theta_v^{j+1}, \Theta_\varpi^j)$;
end

Output: Θ_ϖ^J, Θ_v^J

$$L(x, v) = \frac{\eta}{2}(x-\kappa)^2 + \frac{c}{2}v^2, \quad \text{and} \quad u_T(x) = \frac{\gamma}{2}(x-\zeta)^2,$$

where $\kappa, \zeta \in \mathbb{R}$, $\eta, \gamma \geqslant 0$, and $c > 0$. We set $T = 1$, $M = 30$,

$$Q'(t) = \left(\overline{Q}(t) - Q(t)\right), \quad Q(0) = q_0, \quad \overline{Q}(t) = 7e^{-t}\sin(3\pi t),$$

$N = 10$, $m_0 \sim \mathcal{N}(0, 0.4)$.

Euler-Lagrange is tested over 100 trajectories every epoch (500 iterations).
Training for 50.000 steps.

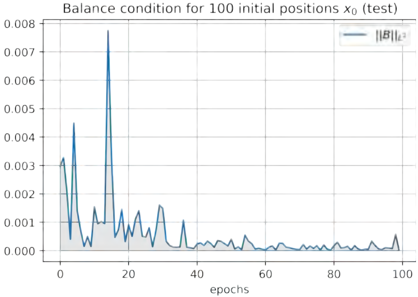

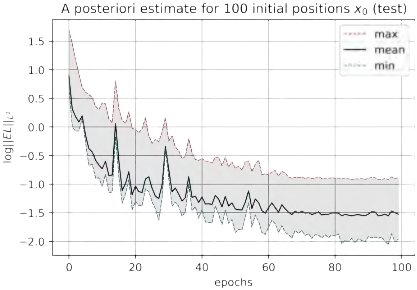

Balance condition and Euler-Lagrange residual.

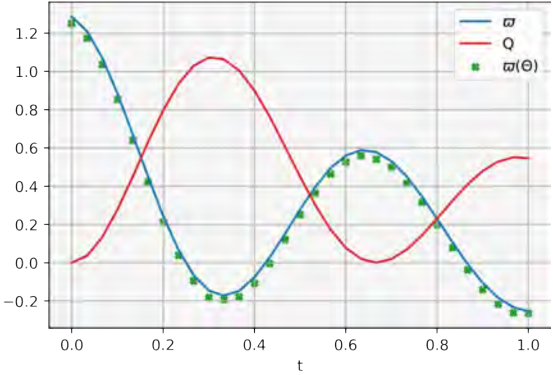

Supply and price.

Numerical results

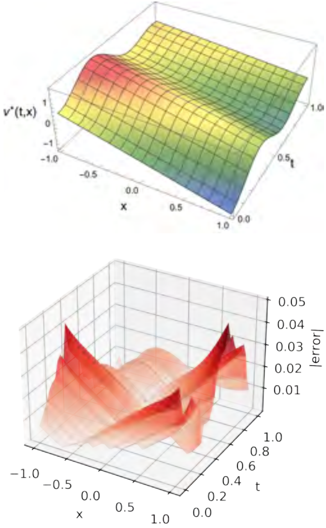

v^* and error.

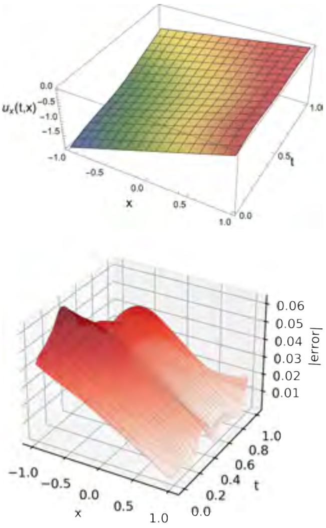

u_x and error.

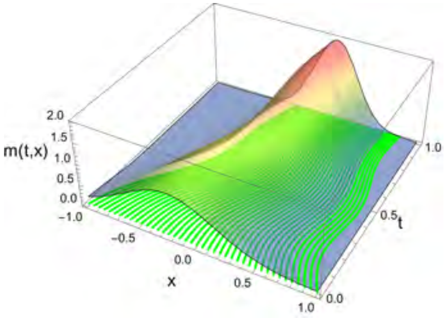

m and v^* trajectories.

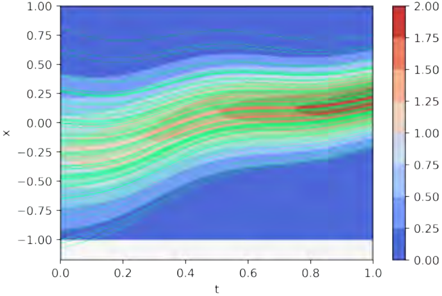

m and $v(\Theta)$ trajectories.

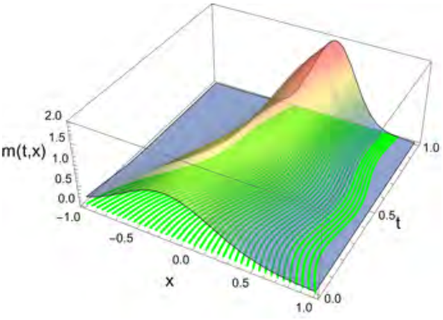

m and v^* trajectories.

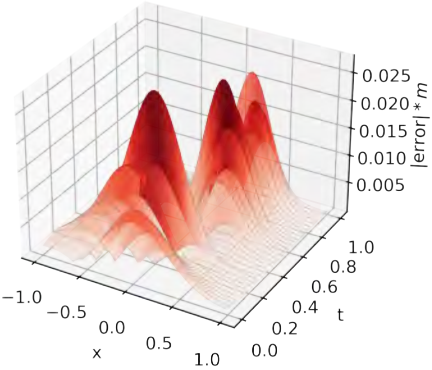

m-weighted error (v^*).

(Gomes et al. (2021)) We consider $dQ(t) = b^s(Q, \varpi)dt + \sigma^s(Q, \varpi)dW(t)$.
Solution: assume

$$d\varpi(t) = b^p(Q, \varpi)dt + \sigma^p(Q, \varpi)dW(t).$$

The balance condition determines b^p and σ^p from b^s and σ^s.

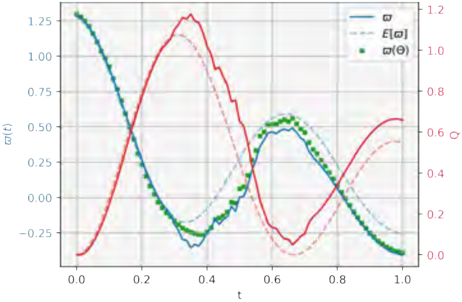

Exact price and approximation (one realization).

Contributions:

- Lagrange multiplier formulation
- Optimization framework (min-max problem)
- Use of neural networks in an optimization setting
- Optimality conditions as a tool to assess convergence of training.
- Decoupling the mean-field games system enables numerical schemes both for the Hamilton-Jacobi and transport equations.

References I

René Aid, Andrea Cosso, and Huyên Pham. Equilibrium price in intraday electricity markets, 2020.

René Aïd, Roxana Dumitrescu, and Peter Tankov. The entry and exit game in the electricity markets: A mean-field game approach. *Journal of Dynamics & Games*, 8(4):331–358, 2021.

C. Alasseur, I. Ben Taher, and A. Matoussi. An extended mean field game for storage in smart grids. *Journal of Optimization Theory and Applications*, 184(2):644–670, 2020. URL `www.scopus.com`.

Clémence Alasseur, Luciano Campi, Roxana Dumitrescu, and Jia Zeng. Mfg model with a long-lived penalty at random jump times: application to demand side management for electricity contracts, 2021.

Y. Ashrafyan, T. Bakaryan, D. Gomes, and J. Gutierrez. A variational approach for price formation models in one dimension, 2022.

T. Basar and R. Srikant. Revenue-maximizing pricing and capacity expansion in a many-users regime. In *Proceedings.Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 1, pages 294–301 vol.1, 2002. doi: 10.1109/INFCOM.2002.1019271.

References II

B. Djehiche, J. Barreiro-Gomez, and H. Tembine. Price Dynamics for Electricity in Smart Grid Via Mean-Field-Type Games. *Dynamic Games and Applications*, 10(4):798–818, December 2020. doi: 10.1007/s13235-020-00367-. URL https://ideas.repec.org/a/spr/dyngam/v10y2020i4d10.1007_s13235-020-00367-8.html.

O. Féron, P. Tankov, and L. Tinsi. Price Formation and Optimal Trading in Intraday Electricity Markets with a Major Player. *Risks*, 8(4):1–1, December 2020. URL https://ideas.repec.org/a/gam/jrisks/v8y2020i4p133-d457902.html.

Olivier Féron, Peter Tankov, and Laura Tinsi. Price formation and optimal trading in intraday electricity markets, 2021.

M. Fujii and A. Takahashi. A Mean Field Game Approach to Equilibrium Pricing with Market Clearing Condition. Papers 2003.03035, arXiv.org, March 2020. URL https://ideas.repec.org/p/arx/papers/2003.03035.html.

M. Fujii and A. Takahashi. Equilibrium price formation with a major player and its mean field limit, 2021.

References III

D. Gomes and J. Saúde. A Mean-Field Game Approach to Price Formation. *Dyn. Games Appl.*, 11(1):29–53, 2021. ISSN 2153-0785. doi: 10.1007/s13235-020-00348-x. URL https://doi.org/10.1007/s13235-020-00348-x.

D. Gomes, J. Gutierrez, and R. Ribeiro. A mean field game price model with noise. *Math. Eng.*, 3(4):Paper No. 028, 14, 2021. doi: 10.3934/mine.2021028. URL https://urldefense.com/v3/__https://doi.org/10.3934/mine.2021028__;!!Nmw4Hv0!glu4jrj9qpvrd4SxjKTZRrM0dtKQgydRLdDbzi_hci3DOX4uA3x59ZtzthwYIur-F_b0og-EZBYScg$.

Ernest K. Ryu and Stephen P. Boyd. A primer on monotone operator methods. 2015.

H. Shen and T. Basar. Pricing under information asymmetry for a large population of users. *Telecommun. Syst.*, 47(1-2):123–136, 2011. doi: 10.1007/s11235-010-9306-2. URL https://doi.org/10.1007/s11235-010-9306-2.

References IV

A. Shrivats, D. Firoozi, and S. Jaimungal. A Mean-Field Game Approach to Equilibrium Pricing, Optimal Generation, and Trading in Solar Renewable Energy Certificate Markets. Papers 2003.04938, arXiv.org, March 2020. URL https://ideas.repec.org/p/arx/papers/2003.04938.html.