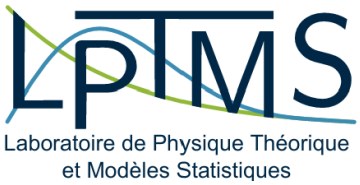

Internship report

A mean-field game approach for balancing intermittent power generation and consumption in electrical networks

Loumi Gatouillat

Master 2 Physics of Complex Systems
Université Paris-Saclay

Under the supervision of Denis Ullmo (LPTMS) and Cosimo Gorini (SPEC, CEA).

Abstract

This internship was conducted within the broader context of the energy transition, with a focus on decentralized electricity systems. As decentralized grids are becoming increasingly relevant in response to renewable energy integration and the need for local autonomy, existing studies, such as those by Schäfer [1], have highlighted both technical and economic challenges. The first part of the work involved an in-depth review of power grid dynamics, including stability, frequency control, and the role of distributed generation, in order to better understand the underlying physical constraints of such systems. In the second phase, the focus shifted toward the economic modeling of decentralized energy consumption. Building on Schäfer's framework, we adopted a Mean Field Game approach to model a large population of consumers who individually optimize their consumption profiles in response to a shared, uncontrollable, and potentially variable electricity production. The MFG formalism allows us to capture both the individual decision-making process and its impact on the collective behavior through an equilibrium framework. This modeling offers insights into how decentralized demand-side flexibility can contribute to grid stability and efficiency in future energy systems.

Contents

1 Introduction **4**
- 1.1 Context and motivations . . . 4
- 1.2 Mean Field Game theory . . . 4
- 1.3 Objectives . . . 5

2 Power grids and swing equation **5**

3 Stability Analysis of the Power Grid **6**
- 3.1 Synchronous steady state and global power balance . . . 7
  - 3.1.1 Average frequency . . . 7
- 3.2 Case of Constant Power Injections . . . 8
  - 3.2.1 Linearisation of the equations . . . 8
  - 3.2.2 Computation of the Synchronous Fixed Point . . . 8
  - 3.2.3 Linearized Dynamics and Local Stability Analysis . . . 9
- 3.3 Time dependant power injections: feedback with a price . . . 11
  - 3.3.1 Instantaneous adaptation $T = 0$ . . . 12
  - 3.3.2 Slow adaptation with discrete time steps . . . 13

4 First steps towards mean field games **13**
- 4.1 General Structure of a Mean field game . . . 13
- 4.2 A canonical example of MFG . . . 14
- 4.3 Simplified Model . . . 16

5 Conclusion **20**

Appendices **21**

A Swing equation **21**
- A.1 Equation of motion describing generator rotor dynamics: the swing equation . . . 21
- A.2 Derivation of the swing equation . . . 21

B Computation of the pseudo inverse **22**
- B.1 Reduction of the Laplacian . . . 22
- B.2 Moore Penrose pseudo-inverse . . . 23

C Reconstruction of the analytical solution **24**

1 Introduction

1.1 Context and motivations

The production and consumption of electricity are undergoing profound transformations, driven by two major trends. On the one hand, renewable energy sources like wind and solar are playing an increasingly central role in the energy mix. While essential for decarbonization, their intermittent and weather-dependent nature introduces significant variability into energy production. On the other hand, technological advances enable real-time monitoring and precise control of electricity use, even down to individual devices. Together, these developments bring opportunities—but also significant challenges—for the stability and management of power systems.

A key question is how to adapt a historically centralized and predictable grid to a more decentralized and variable environment, where both production and consumption are more distributed, dynamic, and uncertain. The goal remains the same: to ensure a real-time balance between supply and demand while maintaining stability in the face of fluctuations or localized failures. But the tools and strategies for achieving this are rapidly evolving.

In traditional power grids, energy production is typically centralized and controllable, coming from a small number of large facilities such as nuclear, hydro, or gas plants. These systems benefit from predictability and are generally operated using centralized control methods, with reserve capacities absorbing demand variability [2]. By contrast, consumption is diffuse, variable, and only partially predictable and controllable.

This balance is increasingly disrupted by the rise of renewable sources, which shift the uncertainty to the production side. Renewable output depends on environmental conditions and is often spatially dispersed and less controllable. As their share in the energy mix grows, relying on reserve generation becomes economically and environmentally unsustainable.

To cope with this, future systems may need to become more flexible on the consumption side—adapting demand to match volatile production, through price signals, automated responses, or coordinated strategies. However, managing the behavior of millions of devices and users—each with their own constraints and objectives—poses a fundamental coordination problem.

This motivates a shift in perspective: instead of relying on centralized control, one can model the system as a network of autonomous agents—producers, consumers, or both—who make decisions locally based on their own cost-benefit analysis. The interaction among these agents gives rise to emergent collective behavior. In this context, game theory provides a natural framework to describe how individual decisions shape and respond to global system dynamics.

However, when the number of agents becomes very large, classical game-theoretic approaches become analytically and computationally intractable. This is where Mean Field Game (MFG) theory enters the picture, offering a scalable alternative to model and analyze such large, decentralized systems.

1.2 Mean Field Game theory

To address large-scale coordination problems in decentralized systems like future power grids, Mean Field Game (MFG) theory offers a powerful modeling framework. It originates from the idea that in a population with many agents, each individual has a negligible impact on the whole, and thus reacts to an aggregate quantity—the mean field—rather than tracking the behavior of all others individually.

In this setting, each agent chooses a strategy, which is a time-dependent plan of action, aiming to minimize a cost function that reflects their preferences and constraints. This cost function may balance energy expenses, comfort levels, or operational constraints, and it often depends not only on the agent's own actions but also on the global behavior of the system, as summarized by the mean field.

The core idea of MFG theory is to study the equilibrium that arises when each agent optimally adapts their behavior in response to the evolving mean field, while the mean field itself is shaped

by the collective actions of all agents. This leads to a system of coupled equations: one governing the optimal control of a representative agent (often a Hamilton–Jacobi–Bellman equation), and the other describing the evolution of the distribution of agents over time (a Fokker–Planck or continuity equation).

MFG theory thus transforms a complex multi-agent optimization problem into a more tractable framework involving only a representative agent and an aggregate population effect. This makes it well-suited to studying decentralized systems where explicit modeling of all individual interactions is impossible.

Since their introduction by Lasry and Lions [3] and independently by Huang, Malhamé, and Caines [4], MFGs have been applied to a wide range of domains including economics, finance, engineering, traffic flow, and recently, energy systems. In particular, they offer a promising way to design scalable, distributed control mechanisms that align individual incentives with the collective goals of a power grid transitioning toward renewables and decentralization.

1.3 Objectives

This report focuses on modeling and understanding the dynamics of power grids, especially in the context of decentralized systems where production and consumption are more variable.

We start by studying the physical behavior of power networks using simplified dynamical models. In this first part, we assume that power generation and consumption are constant over time. This allows us to explore fundamental questions about grid stability—such as when and how the system reaches a synchronized state where all nodes operate at the same frequency.

In the second phase, we move to a more realistic setting where power injections change over time, reflecting the variability of renewable sources and changing demand. This is where price mechanisms come into play: agents (as consumers or producers) can adjust their behavior based on local electricity prices. We reproduce and analyze the results of [1], which show how decentralized price-based control can impact the overall stability of the system.

Finally, we introduce a Mean Field Game (MFG) approach to model the system when the number of agents is large. In this framework, each agent aims to minimize a personal cost that may include electricity price, comfort, or deviation from a target behavior. Rather than considering each interaction individually, MFG focuses on how agents respond to the average behavior of the population. This makes the problem more tractable while still capturing the key interactions between agents and the system as a whole.

2 Power grids and swing equation

A power grid is a system of electrically coupled devices designed to deliver power from generators to consumers. This system can be represented as a network, where each node corresponds to a location where power is either generated (producer), consumed (load), or redirected through transmission lines (branching point). To simplify the modeling of such networks, we often construct effective networks that preserve the essential structure of power transmission between generators or loads while eliminating internal nodes that do not generate or consume power. This reduction process is known as Kron reduction [5]. In what follows, we assume that each remaining node represents a rotating machine either producing or consuming power, and we neglect passive branching points. A link is a transmission line between two nodes, through which power is transmitted. Depending on the local balance between power generation and consumption, a node can act as a generator ($P_j^{\text{mech}} > 0$), a consumer ($P_j^{\text{mech}} < 0$), or remain balanced ($P_j^{\text{mech}} = 0$). The state of each rotating machine is determined by its phase $\theta_j(t)$ and its angular frequency $\omega_j(t)$, both are measured relative to the reference grid rotating at a nominal frequency $\Omega = 2\pi \times 50$ Hz. The absolute phase is given by $\Theta_j(t) = \Omega \times t + \theta_j(t)$, so that $\omega_j(t) = d\theta_j(t)/dt$ represents the deviation from the reference frequency. We are particularly interested in the phase differences $\theta_k - \theta_j$ between connected nodes, because they determine the direction and

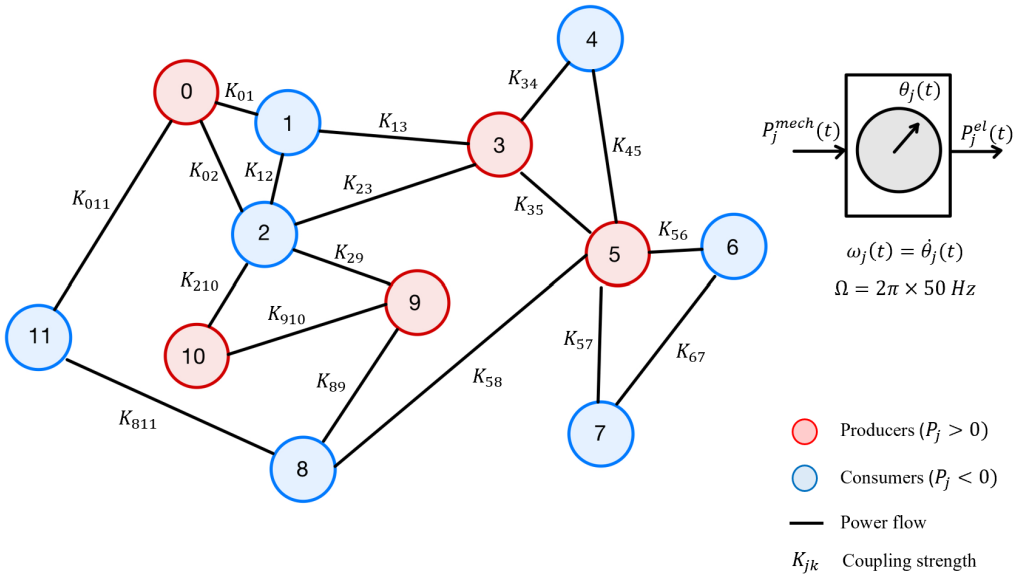

Figure 1: Schematic description of a power grid

magnitude of power flows in the network. The role of the electrical tension (voltage) is implicitly included in these flows and will later appear explicitly in the coupling terms of the model.
The dynamics of each node is given by the swing equation, which is derived from Newton's second law for rotational systems (see Appendix A.1).

$$M_j \frac{d^2\theta_j(t)}{dt^2} + \kappa_j \frac{d\theta_j(t)}{dt} = P_j^{mech}(t) - P_j^{el(t)} \tag{1}$$

where M_j is the moment of inertia, κ_j is a measure of the damping, and P_j^{el} is the electrical power that is transmitted to or from other nodes through transmission lines.

To understand how electrical power is redistributed across the network, and how this relates to the physical properties of the grid, we need to account for the underlying physics - in particular, Ohm's and Kirchhoff's laws. The relevant calculations are detailed in Appendix A.2. The reader can find more detailed derivations in [1] or [6].
The coupled swing equation we will use for the rest of this work writes

$$M_j \frac{d^2\theta_j(t)}{dt^2} + \kappa_j \frac{d\theta_j(t)}{dt} = P_j^{mech}(t) + \sum_{k=1}^{N} K_{jk} \sin(\theta_k(t) - \theta_j(t)) \tag{2}$$

with K_{jk} the coupling strength. This form assumes an ideal lossless network, where power transfer is determined only by the phase differences and not by resistive losses in the transmission lines.

3 Stability Analysis of the Power Grid

In this section, I will detail the stability analysis of the system, first in a simple case where the source terms are constant and then we will add a time dependency.

3.1 Synchronous steady state and global power balance

To ensure stable and steady operation of a power grid, all machines must rotate in synchrony, in other words, the frequency needs to be uniform across the entire network

$$\omega_j(t) = \omega_s \quad \forall j \in \{1, ...N\} \tag{3}$$

If the machines in the power grid do not rotate at the same frequency, that is, if $\omega_j(t) \neq \omega_s$ for some node j, then the phase differences between nodes will continuously change over time. As a result, the power flow between two nodes, given by $P_{jk}(t) = K_{jk} \sin(\theta_k(t) - \theta_j(t))$, will become time-dependent and oscillatory. These oscillations mean that the average power exchanged between nodes will tend to zero over time, preventing the establishment of a stable and efficient energy transfer. Therefore, maintaining a synchronous state, where all machines operate at the same frequency, is crucial for ensuring steady and reliable power flows in the grid. Substituting the condition $\omega_j(t) = \omega_s$ into equation of motion (2) shows that the synchronous state is determined by the equation

$$\kappa_j \omega_s = P_j^{mech}(t) + \sum_{k=1}^{N} K_{jk} \sin(\theta_k(t) - \theta_j(t)) \quad \forall j \in \{1, ...N\} \tag{4}$$

Summing up these equations for all j, for constant P_j and using that $K_{ij} = K_{ji}$ gives

$$\sum_j P_j^{mech} = \sum_j \kappa_j \omega_s \tag{5}$$

Since $\omega_j(t)$ represents the deviation from the nominal frequency Ω, we are particularly interested in the solution $\omega_j(t) = \omega_s = 0$, so the stability condition (5) becomes

$$\sum_j P_j^{mech} = 0 \tag{6}$$

That is equivalent to say that the total production must be equal to the overall consumption.

3.1.1 Average frequency

To understand the global behavior of the power grid, we analyze how the average frequency deviation $\langle\omega(t)\rangle$ evolves over time. We start from a stable solution $\langle\omega(0)\rangle$ and at $t = 0$ we change the power injections in the system, so it has to stabilize to a new equilibrium with a new characteristic time. For simplicity, we assume that all $M_j = M$ and all $\kappa_j = \kappa$ are identical. $\langle\omega\rangle = \frac{1}{N}\sum_j \omega_j$ is determined by

$$M\frac{d\langle\omega\rangle}{dt} + \kappa\langle\omega\rangle = \frac{1}{N}\sum_j P_j^{mech} \tag{7}$$

This equation (7) is analytically solvable, the solution for constant power injections is

$$\langle\omega(t)\rangle = \left(\langle\omega(t_0)\rangle - \frac{\sum_j P_j^{mech}}{N\kappa}\right) e^{-\kappa t/M} + \frac{\sum_j P_j^{mech}}{N\kappa} \tag{8}$$

In the large time limit, the system converges to $\langle\omega\rangle = \sum_j P_j^{mech}/N\kappa$.The average frequency deviation $\langle\omega\rangle$reflects the global power balance of the system. If we want the network to stabilize around the nominal frequency $\Omega = 50 \times 2\pi$ Hz, meaning $\langle\omega\rangle \to 0$ then the total injected power must equal the total consumed power. This is consistent with condition (6). However, $\langle\omega\rangle \to 0$ only ensures that the average behavior is stable. It does not guarantee that each machine operates synchronously. For full stability, we require that all individual frequencies ω_j converge to the same value, a stricter condition ensuring synchronization across the entire system.

Later in the analysis we will also need the solution for time dependent power injections which is given by:

$$\langle\omega(t)\rangle = (\langle\omega(t_0)\rangle + B(t))e^{\frac{-\kappa}{M}t} \quad \text{with} \quad B(t) = \frac{1}{M}\int_{t_0}^{t}\sum_j P_j^{mech}(\tau)e^{\frac{\kappa}{M}\tau}d\tau \tag{9}$$

Again, this result assumes that the M_j and κ_j are homogeneous across the system.

3.2 Case of Constant Power Injections

3.2.1 Linearisation of the equations

To investigate more general regimes and access the local behavior of the system, we linearize equation (2) assuming small angles θ_j. This step is necessary to express the system in a matrix form, which allows for a more tractable analysis of its stability properties.

The system can be linearized under the assumption that phase differences between connected nodes remain small. This assumption is physically motivated, as large phase differences would lead to excessive power losses along transmission lines and threaten the stability of the grid [6].
The linearized system can be written as

$$\underbrace{\begin{bmatrix}\dot\theta_1\\ \vdots\\ \dot\theta_N\\ \dot\omega_1\\ \vdots\\ \dot\omega_N\end{bmatrix}}_{\underline{\dot X}} = \underbrace{\begin{bmatrix}\mathbf{0} & \mathbf{I_N}\\ -\mathbf{M}^{-1}\mathbf{L} & -\mathbf{M}^{-1}\mathbf{D}\end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix}\theta_1\\ \vdots\\ \theta_N\\ \omega_1\\ \vdots\\ \omega_N\end{bmatrix}}_{\underline{X}} + \underbrace{\begin{bmatrix}0\\ \vdots\\ 0\\ P_1/M_1\\ \vdots\\ P_N/M_N\end{bmatrix}}_{\underline{B}} \tag{10}$$

where each block of the matrix $\mathbf{A}$ is size $N \times N$.
We define the matrices $\mathbf{M}$=diag(M_j) and $\mathbf{D}$=diag(κ_j) $\quad \forall j \in \{1, ..., N\}$. We call $\mathbf{L}$ the Laplacian matrix of the weighted graph K_{jk} and is defined by

$$L_{ij} = \begin{cases} -\sum_k K_{jk} & \text{if } i = j \\ K_{ij} & \text{if } i \neq j \end{cases}$$

3.2.2 Computation of the Synchronous Fixed Point

To study the stability of the system (10), we need first to determine the fixed points ie solve the equation $\underline{\dot X} = 0$. We already know that $\omega_j = 0$ for all j is the fixed point we are interested in, we still need to determine the associated θ's.
We solve

$$\underline{\dot X}^* = 0 = \mathbf{A}\underline{X}^* + \underline{B} \tag{11}$$

We will see that this equation for $\underline{X^*}$ cannot be solved with a usual matrix inversion.
Remember that we imposed for the fixed point $\underline{X}^*$ that its N last components, the ω, are 0, so we have

$$\begin{aligned} \forall j \quad 0 &= \frac{1}{M_j}\sum_k L_{jk}\theta_k^* - \frac{1}{M_j}\kappa_j\omega_j^* + \frac{P_j}{M_j} \\ 0 &= \sum_k L_{jk}\theta_k^* + P_j \end{aligned} \tag{12}$$

Solving equation (11) is equivalent to solving the $N \times N$ system

$$L \begin{bmatrix} \theta_1^* \\ \vdots \\ \theta_N^* \end{bmatrix} = - \begin{bmatrix} P_1 \\ \vdots \\ P_N \end{bmatrix} \Rightarrow L\underline{\theta}^* = -\underline{P} \tag{13}$$

The Laplacian matrix $\mathbf{L}$ is not invertible because its rank is strictly less than its dimension N. Indeed, each row (and column) of $\mathbf{L}$ sums to zero, which implies that $\mathbf{L}$ has a nontrivial kernel: at least one eigenvalue is zero. More precisely, the kernel of $\mathbf{L}$ contains at least the vector of ones $\underline{1} = (1, ..., 1)^T$ (for a connected graph Ker($\mathbf{L}$)=$\underline{1}$)), reflecting the invariance of the system under uniform phase shifts. Physically, this means that adding the same constant to all phase angles θ_j does not change the state of the system, since only phase differences appear in the coupling term $\sin(\theta_k - \theta_j)$ of the swing equation (2), and power flows depend on these differences rather than absolute angles. To handle this singularity, a common approach is to reduce the Laplacian by fixing a reference node (e.g., $\theta_1 = 0$) or by imposing the constraint $\sum_j \theta_j = 0$. This constraint effectively removes the redundant degree of freedom by centering the angles around zero, which is equivalent to choosing a global phase reference. This effectively removes the redundant degree of freedom, yielding a reduced Laplacian **L'** (obtained by removing the first row and column of L and subtracting this column to the remaining column), which is invertible if the graph is connected. The reduction is mathematically justified because in a connected graph, the Laplacian has rank $N-1$ and its kernel is the span of the vector of ones [7]. The computations to reduce the Laplacian can be found in Appendix B.
In numerical applications, a common alternative is to compute the Moore–Penrose pseudo-inverse of the Laplacian, denoted $\mathbf{L}^\dagger$. This approach avoids explicitly selecting a reference node and provides a generalized inverse that minimizes the norm of the solution [8].
The eigenvalues of the system matrix $\mathbf{A}$ depend directly on those of the Laplacian $\mathbf{L}$, since $\mathbf{L}$ encodes the network structure and coupling; thus, understanding the spectrum of $\mathbf{L}$ is key to analyzing the stability and dynamics of the full system governed by $\mathbf{A}$.
We compute the pseudo-inverse numerically and obtain the fixed point $\underline{X}^*$ that depends on the parameters of the system and we can perform a linear stability analysis around it.

3.2.3 Linearized Dynamics and Local Stability Analysis

We write $\underline{X} = \underline{X}^* + \delta\underline{X} \Rightarrow \dot{\underline{X}} = \delta\dot{\underline{X}} = \mathbf{A}\delta\underline{X}$. The stability of the solution around $\underline{X}^*$ is determined by the eigenvalues of $\mathbf{A}$. In the simulations, I used a random graph generated with a connection probability $p = 0.5$ (see Fig.2). The coupling strengths, damping coefficients, inertias, and power injections were all drawn randomly around values used in [7], to reflect heterogeneous parameters with a relevant order of magnitude.

Fig. 2 shows that there is one zero eigenvalue as expected, the remaining eigenvalues are complex with negative real part. This behavior was expected, as predicted by the theorem stating that the eigenvalues of the Laplacian are non-positive [7]. We expect to observe decreasing oscillations towards 0 : the fixed point is stable.
We can solve analytically the system of equations, we expand the solution over the eigenbasis of the matrix $\mathbf{A}$. This reduces the problem to solving a set of decoupled first-order differential equations for the time-dependent coefficients. These coefficients depend on the projection of the initial condition and of the source term onto the eigenvectors of $\mathbf{A}$. Since the matrix is not symmetric, we need to normalize the left and right eigenvectors to ensure a proper decomposition, for the two basis to be orthonormal (orthogonal is not sufficient). Once the coefficients are determined, the full solution is reconstructed as a linear combination of the right eigenvectors. Full details of the derivation are provided in Appendix C. The solution is shown in Fig.3

We can compare the analytical solution to the numerical one done for the non-linearized system, and observe that they have the same stable behavior, both converge to the fixed point $\langle\omega\rangle = 0$.

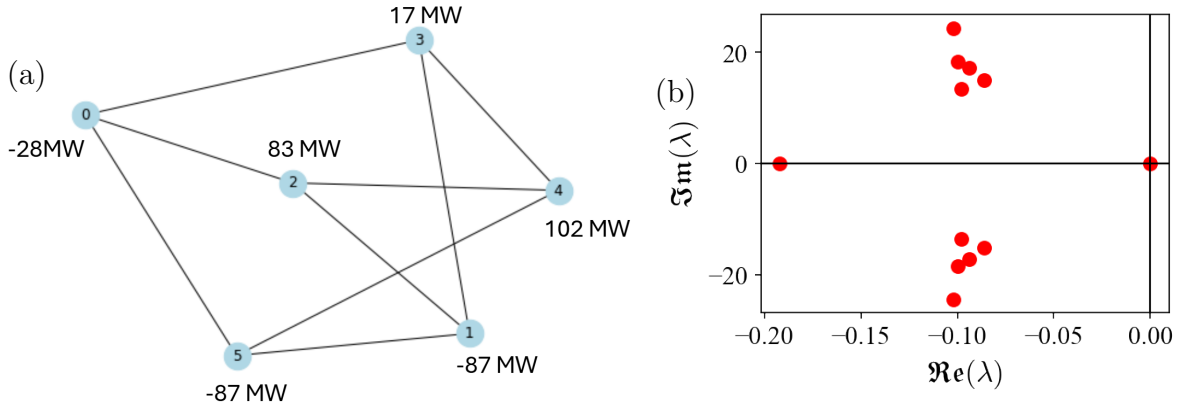

Figure 2: **(a)** Random graph used for the simulations. **(b)** Eigenvalues of $\mathbf{A}$ in the complex plane. We chose the power injections P_j^{mech} randomly such that their sum is zero. The coupling coefficients K_{jk} are randomly and uniformly distributed in the interval $[10, 500]$ MW. The inertias M_j are distributed around $10^4, \text{kg} \cdot \text{m}^2 \times \Omega$ with a $\pm$ 20% variation, and the damping coefficients κ_j around $0.2/s \times 10^4 \times \Omega$ also with a $\pm$ 20% variation [1].

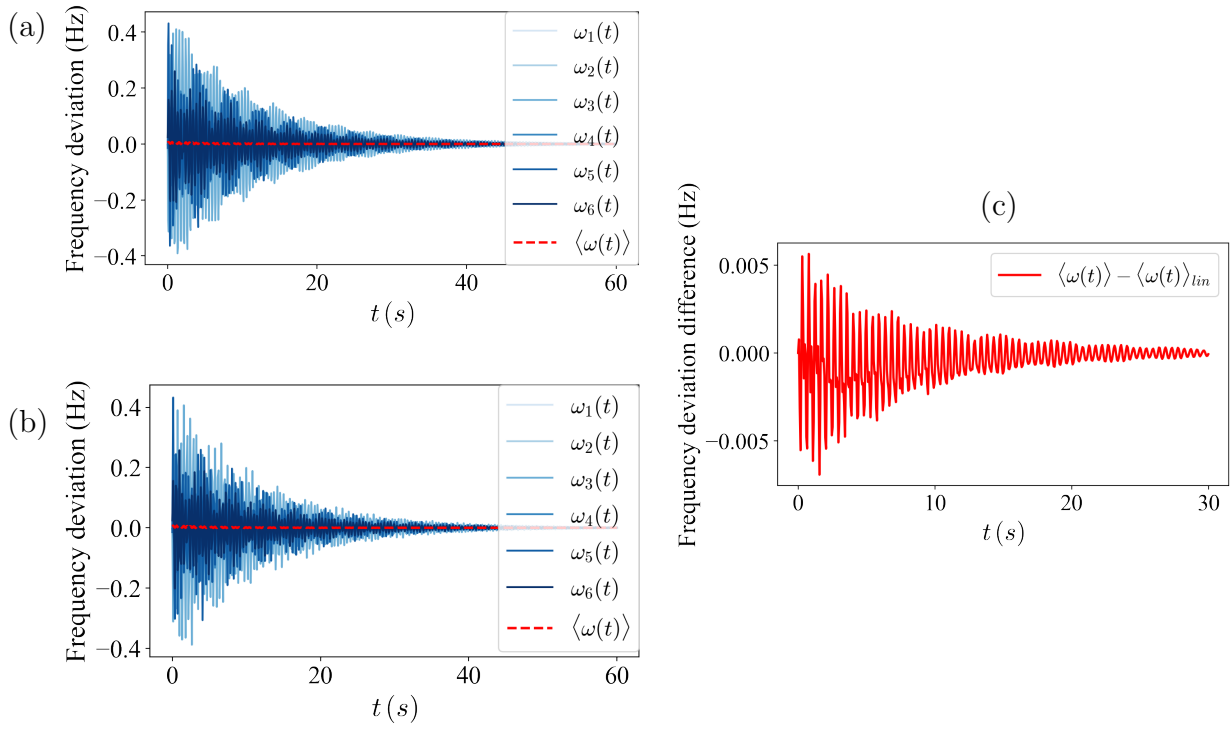

Figure 3: **(a)** Analytical solution in the linear case and **(b)** numerical solution of the non-linear equation.**(c)** Difference between the non linear solution and the linear solution. The numerical solution was obtained using the Runge-Kutta method. Initial conditions were randomly and uniformly distributed around zero, in the interval $[-0.1, 0.1]$.

The two curves exhibit extremely similar behavior, with their difference tending to zero. This confirms the validity of the linear approximation in this regime.

Solving the problem under the assumption of constant power injections is a useful first step, but it is not realistic in practical settings. In reality, both power production and consumption vary over time, particularly in the presence of renewable energy sources, which are inherently intermittent. The goal of the next section is to revisit the same problem while allowing time-dependent power injections, and to introduce a feedback mechanism based on dynamic pricing in order to influence and control the time-dependent behavior of both production and consumption.

3.3 Time dependant power injections: feedback with a price

A key missing element in implementing a Decentralized Smart Grid Control lies in establishing a clear relationship between the local angular frequency deviation ω_j of the grid and the corresponding electricity price $p_j(\omega_j)$. A simple and low-cost device could measure this local frequency and compute the real-time price according to a predefined function $p_j(\omega_j)$. Electric appliances with flexible on/off characteristics—such as washing machines, refrigerators, heat pumps, or electric vehicles—could automatically shift their power consumption to periods of high grid frequency, thereby reducing load during low-frequency periods. With a well-designed pricing function $p(\omega)$, this mechanism can be made economically attractive for both consumers and electricity providers. Grid operators would benefit from reduced reliance on primary, secondary, and tertiary frequency control services. A sharp price increase at low frequencies, coupled with cheap electricity at high frequencies, could also incentivize consumers to adjust their energy usage behavior. Schäffer proposes a Decentralized Smart Grid Control strategy in [1] that implements dynamic demand response within power grids, and investigate its key economic and dynamical implications. The mechanical power P_j^{mech} in the swing equation (2) represents the net power at node j, i.e., the difference between power generated and consumed. Both generation and consumption depend on the instantaneous electricity price $p_j(t)$, which is described through supply $S(p)$ and demand $D(p)$ curves:

$$P_j^{mech} = S_j(p_j(t)) - D_j(p_j(t)) \tag{14}$$

The supply function $S(p)$ reflects the amount of power producers are willing to generate at price p, while the demand function $D(p)$ describes how much consumers want to consume at that same price. Typically, $S(p)$ is an increasing function, whereas $D(p)$ is decreasing. These functions are external to the model and are shaped by generator strategies, consumer preferences, and external conditions such as weather.
Thus, the stability condition (5) becomes

$$\sum_j S_j(p_j) = \sum_j D_j(p_j) + \sum_j \kappa_j \langle\omega\rangle \Rightarrow \sum_j S_j(p_j) = \sum_j D_j(p_j) \quad \text{for } \langle\omega\rangle = 0 \tag{15}$$

In [1], Schäffer et al propose a control approach where the local price p_j is dynamically computed based on the local angular frequency deviation ω_j. Since frequency measurements and updates cannot be instantaneous in practice, the price is computed using a time-averaged frequency deviation ω_j. Assuming that the angular velocity is measured over an interval of fixed period length T, we define

$$\overline{\omega_j}(t) = \frac{1}{T}\int_{t-T}^{T} \omega_j(\tau)d\tau \tag{16}$$

Schäffer suggests in [1] to model the price as a decreasing linear function of the local frequency. In this formulation, when the frequency is high (indicating a surplus of energy), the price is low, which encourages consumption. Conversely, when the frequency is low, the price increases, thereby discouraging consumption and incentivizing generation.

$$p(\overline{\omega_j}(t)) = p_\Omega - \epsilon \times \overline{\omega_j}(t) \tag{17}$$

Here, p_Ω denotes the nominal price corresponding to the nominal frequency, and ϵ represents the sensitivity of the price to frequency deviations—larger values of ϵ lead to more responsive price adjustments with respect to frequency changes. Schäffer presents two technical ways to compute the price. First, a control where the system adapts only in discrete time steps of length T such that the local price is given by

$$p_j = p(\overline{\omega_j}(\lfloor t/T \rfloor T)) \tag{18}$$

where $\lfloor . \rfloor$ denotes the floor function. Second, it can take a certain delay τ for the price to adapt to the frequency deviation such that the price is given by

$$p_j = p(\overline{\omega_j}(t - \tau)) \tag{19}$$

I will only consider the first technical scenario in my analysis of the problem, my colleague Yacine studied the second one with delay during his internship.

3.3.1 Instantaneous adaptation $T = 0$

Consider the ideal case where the price and so the demand adapt instantaneously, i.e, $T = \tau = 0$. The source term in (2): $P_j^{mech}(S(t)) = S_j(S(t)) - D_j(S(t))$ becomes a function of the current angular frequency deviation $\omega_j(t)$. We linearize supply and demand around p_Ω, the reference price at which the system should stabilize (price corresponding to $\omega_j = 0$).

$$\begin{aligned} S_j(p_j) &= S_j(p_\Omega) + \underbrace{\left.\frac{dS_j}{dp}\right|_{p_\Omega}}_{=s_j} (p_j - p_\Omega) \\ D_j(p_j) &= D_j(p_\Omega) + \underbrace{\left.\frac{dD_j}{dp}\right|_{p_\Omega}}_{=d_j} (p_j - p_\Omega) \end{aligned} \tag{20}$$

In general, the supply increases with the price whereas the demand decreases so $s_j \geq 0$ and $d_j \leq 0$ such that

$$\forall j \quad s_j - d_j \geq 0 \tag{21}$$

To ensure global frequency stability in the synchronous regime ($\langle \omega \rangle = 0$), the total net injected power must vanish (see (6)). This imposes a consistency condition on the linearization: the sum of the linearized supply terms must equal the sum of the demand terms at p_Ω, i.e.,

$$\sum_j S_j(p_\Omega) = \sum_j D_j(p_{\Omega)}) \tag{22}$$

Substituting (20) in the Swing equation (2) yields

$$M_j \frac{d^2\theta_j(t)}{dt^2} + \kappa_j \frac{d\theta_j(t)}{dt} = S_j(p_\Omega) - D_j(p_\Omega) - \epsilon(s_j - d_j)\frac{d\theta_j}{dt} + \sum_{k=1}^{N} K_{jk} sin(\theta_k - \theta_j) \tag{23}$$

Using (21), we see that (23) is simply the swing equation with an effective damping coefficient

$$\kappa_j^{eff} = \kappa_j + \epsilon(s_j - d_j) \geq \kappa_j \tag{24}$$

An instantaneous economic response increases the effective damping coefficient, it shortens the convergence time to the fixed point after perturbations, see equation (8).

3.3.2 Slow adaptation with discrete time steps

In practice, the frequency measurements and updates cannot be instantaneous so the price is computed using a time average (see equation (16)). Here, I treat the case where the price is computed using the time average of the frequency but the update is instantaneous i.e, $\tau = 0$ and with discrete times steps (see equation (18)).
We consider a homogeneous network i.e, for all j $M_j = M$, and $\kappa_j = \kappa$. In this case, we can solve the equation for the average frequency $\langle\omega\rangle$ and we know that in the large time limit it relaxes to $\sum_j P_j/N\kappa$ (see equations (7) and (8)). Introducing this expression for $\langle\omega\rangle$ in the expression of the price (18), we obtain a new price as

$$\begin{aligned} p' - p_\Omega &= -\epsilon\langle\omega\rangle \\ &= -\epsilon\frac{\sum_j s_j - \sum_j d_j}{\sum_j \kappa_j}(p - p_\Omega) \quad \text{because } \sum_j S_j(p_\Omega) - \sum_j D_j(p_\Omega) = 0 \end{aligned} \tag{25}$$

The system is stable if and only if

$$\left|\epsilon \times \frac{\sum_j s_j - \sum_j d_j}{\sum_j \kappa_j}\right| < 1 \Rightarrow \epsilon < \epsilon_{crit} = \frac{\sum_j \kappa_j}{\sum_j s_j - d_j} \tag{26}$$

where ϵ determines the slope of the linear price–frequency relationship. It controls how strongly the price reacts to local frequency deviations: a larger ϵ means that even small changes in frequency cause larger adjustments in price. This enhances the feedback effect between physical dynamics and economic signals, but if ϵ is too large, it can destabilize the system by introducing or amplifying oscillatory behavior. To model this effect of the price on the frequency, I used the same parameters and network as [1]. We assume that the price and subsequently the supply and demand are adapted after a time step $T = 60s$. We chose the equilibrium price p_Ω=24ctkWh^{-1} and damping constants κ_j=0.2/s$\times M_j$ with $M_j = 10^4$kgm$^2 \times \Omega$. All transmission lines have the same capacity $K = 200$MW. The price elasticities are defined as $E_S = (dS_j/dp)/(S_j/p)|_{p_\Omega} = 0.3$ and $E_D = (dD_j/dp)/(D_j/p)|_{p_\Omega} = -0.3$. For these parameters, the system will be stable if and only if $\epsilon < \epsilon_{crit} \simeq 3$(ct/kWh)(2$\pi$Hz). An example of dynamics is shown in Fig. 4. For this simulation, I divided the time horizon into a finite number of equal intervals. On each interval, I solved the system of differential equations and computed both the instantaneous average frequency and the time-averaged frequency (as defined in equation (16)). This time-averaged frequency was then used to determine a price, which served as an input for the differential equations in the subsequent time interval.

Having understood the system's dynamics and how the price signal influences consumption behavior, we now move beyond this passive response to price. By introducing strategic agents within a game-theoretic framework, we allow each individual to actively choose their consumption time based on personal preferences and anticipated collective behavior—this leads us to the Mean Field Game approach.

4 First steps towards mean field games

4.1 General Structure of a Mean field game

Mean Field Games (MFGs) provide a mathematical framework for analyzing the strategic behavior of large populations of indistinguishable agents. Rather than modeling each agent individually, MFGs consider how agents interact indirectly through an aggregate quantity : the mean field. This leads to a tractable formulation of decentralized decision-making in large systems, as illustrated in Fig. 5. The idea, in the context of power grids, is that each agent optimizes its cost function. This function depends on its production S and consumption D. Each agent chooses a strategy to minimize its own

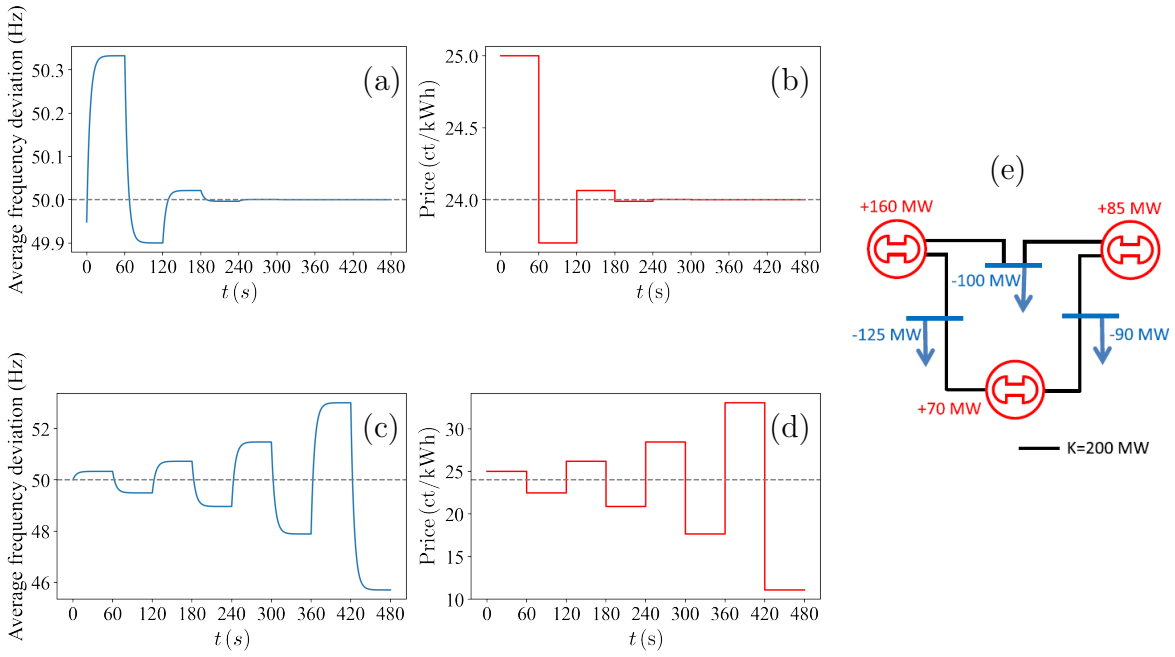

Figure 4: Dynamics of frequency and price under discrete-time updates with step size $T = 60$ s. System stability depends critically on the slope ϵ of the price-frequency curve. Panels **(a)** and **(b)** show the stable case with $\epsilon = 1$ (ct kWh^{-1}) $(2\pi$ Hz$)^{-1} < \epsilon_{\rm cr}$, while panels **(c)** and **(d)** show the unstable case with $\epsilon = 5$ (ct kWh^{-1}) $(2\pi$ Hz$)^{-1} > \epsilon_{\rm cr}$. We display the dynamics of **(a, c)** the local prices p_j and **(b, d)** the local frequencies $(\Omega + \omega_j)/2\pi$, starting from an initial price $p_0 = 25$ ct kWh^{-1}, slightly above the equilibrium price $p(\Omega) = 24$ ct kWh^{-1}. The angular frequency deviations ω_j vary only slightly across nodes, as illustrated in the inset of panel **(b)**. These residual oscillations decay over longer time scales, so that in the stable case (**a**, **b**), the system converges to a fixed point with $\omega_j = \langle\omega\rangle$ and $p_j = p(\langle\omega\rangle)$. The model grid is depicted in panel **(e)**, where generators are shown in red and consumers in blue. Additional parameter values are given in the main text.

cost, but the combination of all these strategies modifies the overall consumption and production in the system. As a result, this collective behavior affects the average frequency deviation of the grid, which is governed by the swing equation.

A price signal can then be defined based on this average frequency deviation. This price acts as an economic feedback that incentivizes agents to adjust their consumption: consuming less when the frequency drops (indicating high demand or low supply), and consuming more when the frequency rises. In doing so, the agents are economically encouraged to help stabilize the frequency around zero deviation, thereby maintaining the physical stability of the power system.

4.2 A canonical example of MFG

In this section, I will provide a general overview of the main ideas of MFGs.
As a starting point, consider N agents. Their individual states are described by a continuous variable $\underline{X}^i \in \mathbb{R}^d, i \in [1, N]$, which depending on the context, could represent the position of the agent or in our context could be the amount of energy the agent produces or consumes. We assume these state variables evolves according to Langevin equation

$$d\underline{X}^i_t = \underline{a}^i_t dt + \sigma d\underline{W}^i_t \quad \text{with initial conditions } \underline{X}^i_t = \underline{x}^i_0 \tag{27}$$

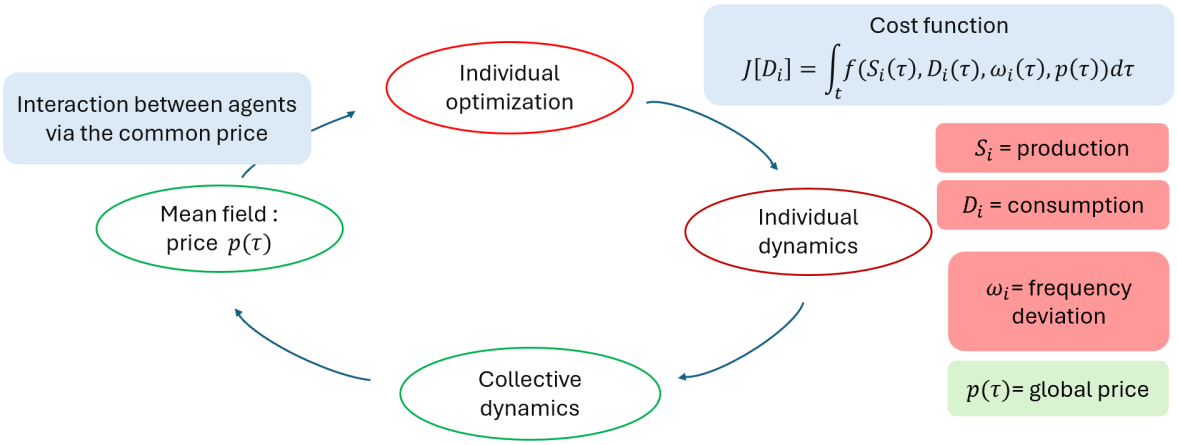

Figure 5: Schematic of a Mean Field Game (MFG) in a power system. Agents individually optimize their consumption $D_i(\tau)$ to minimize their cost J_i, which depends on production S_i, frequency deviation ω_i, and the global price $p(\tau)$. Their collective actions generate this price i.e the mean field, which in turn influences all agents' decisions. This creates a closed-loop equilibrium where no agent can unilaterally reduce their cost further called Nash equilibrium.

Each of the d components of $\underline{W^i}$ is an independent white noise with variance one and the control parameter is the velocity $\underline{a}_t^i$. This control parameter is adjusted in time in order to minimize a cost function over a time interval $[0, T]$.

$$J[\underline{a}_t^i](\underline{x}_t^i, t) = \mathbb{E}\left(\int_t^T (L(\underline{X}_\tau^i, \underline{a}_\tau^i) - V[m_\tau](\underline{X}_t^i))d\tau\right) + \mathbb{E}(J_T(\underline{X}_T^i)) \tag{28}$$

where $\mathbb{E}$ denotes the average over the realizations, so the cost function is averaged over the noise. The cost functional represents the expected cost for an agent who starts at $\underline{X}_t^i$ at time t, averaged over all possible realizations of the stochastic noise. It expresses the rational expectations of each player in the game. In the framework of differential games, this cost typically consists of two components: a running cost accumulated over time and a terminal cost incurred at the end of the optimization horizon at time T. The running cost takes the form $L(\underline{X}, \underline{a}) - V[m](\underline{X})$, where $L(\underline{X}, \underline{a})$ is the so-called free Lagrangian term, depending only on the individual state $\underline{X}$ and control $\underline{a}$, and $V[m_t](\underline{X})$ is a potential term that reflects interactions between agents. This potential is a functional of the population distribution m_t, meaning that it encodes how the collective behavior of all agents affects the cost perceived by an individual agent at position $\underline{X}$.

$$m_t(\underline{x}) = \frac{1}{N}\sum_i \delta(\underline{x} - \underline{X}_t^i) \tag{29}$$

The terminal cost, denoted $J_T(\underline{X})$, depends only on the final state of the agent and captures objectives or penalties associated with the end of the time horizon.

We have implicitly assumed that all agents exhibit identical behavior, meaning that they only differ by initial condition and the subsequent choices of control parameters.
Finally, as is standard in control theory, one defines the value function as the minimal cost achievable by an agent, optimized over all admissible control strategies, given an initial state $\underline{X}$ at time t. This function encapsulates the optimal expected cost and is independent of the specific agent index i, reflecting the generic behavior of any representative agent. It is expressed as:

$$u(\underline{x}, t) = \min_{\underline{a}} J[\underline{a}](\underline{x}, t) \tag{30}$$

Although the continuous differential game described above is already simplified, it remains intractable for many players. Mean Field Games offer a framework to study the system in the limit of an infinite number of agents. Instead of tracking each individual, the model considers the population density $m_t(\underline{X})$, which becomes deterministic as $N \to \infty$. This allows the problem to reduce to the optimization of a representative agent interacting with the mean field, whose value function satisfies a Hamilton–Jacobi–Bellman equation (HJB).

The derivations can be found in [9]

$$\begin{cases} \partial_t u_t(\underline{x}) + H(\underline{x}, \nabla u_t(\underline{x})) + \frac{\sigma^2}{2}\Delta u_t(\underline{x}) = V[m_t](\underline{x}) \\ u_T(\underline{x}) = J_T(\underline{x}) \end{cases} \tag{31}$$

where $H(\underline{x}, \underline{p}) = \inf_{\underline{\alpha}}(L(\underline{x}, \underline{\alpha}) + \underline{p}\underline{\alpha})$, with $\underline{p}$ being the conjugate momentum associated to the state variable $\underline{x}$. This HJB equation represents the mathematical formulation of the individual optimization process shown in the right part of the schematic. The optimal control corresponds to the value of $\underline{a}$ that minimizes the Hamiltonian, and is given by:

$$\underline{a_t^*}(\underline{x}) = \frac{\partial H}{\partial p}(\underline{x}, \nabla u_t(\underline{x})).$$

This optimal control governs the agent's actions, which in turn influence the collective dynamics shown in Fig. 5. By substituting this expression into the Langevin dynamics, the evolution of the optimized agent density m_t, starting from an initial distribution m_0, is governed by the following Fokker-Planck (or Kolmogorov) equation:

$$\partial_t m_t(\underline{x}) - \nu\Delta m_t(\underline{x}) - \nabla \cdot \left(m_t(\underline{x}) \frac{\partial H}{\partial p}(\underline{x}, \nabla u_t(\underline{x})) \right) = 0 \tag{32}$$

The solution of the HJB equation provides the optimal control for a given distribution m_t, which can then be used to update m_t through the Fokker–Planck (FP) equation, forming a self-consistent loop. The coupled system formed by the Hamilton–Jacobi–Bellman equation and the FP equation characterizes a Mean Field Nash Equilibrium. In this equilibrium, each agent optimally responds to the population distribution $m_t(\underline{x})$, assuming it is fixed, while the distribution itself evolves according to the collective behavior resulting from these optimal responses. In other words, no agent can reduce their expected cost by unilaterally changing their strategy, given the mean field induced by others. This defines a fixed point between individual optimization and population dynamics, completing the closed-loop structure of the game.

4.3 Simplified Model

The first idea was to model the optimization of electricity consumption in the presence of stochastic, variable renewable energy production. In such a setting, the production of agent i could be represented as a stochastic process $dS_i = \mu_0 dt + \sigma dW_i$.

However, in this first stage, we chose to simplify the model by working with deterministic and time-dependent production. That is, we assume that the supply $S_i(t)$ is a known, fixed function of time, which does not depend on the agent's actions nor include any stochastic fluctuations. This simplification allows us to focus more clearly on the impact of consumption decisions and to derive analytical results.

We consider a population of agents where each one consumes electricity once per day, at a chosen start time t_i, over a fixed interval of duration Δt. This is a rectangular function of amplitude D_0, starting at time t_i and lasting for duration Δt. For instance, it models taking a shower: the agent starts at t_i and lasts Δt. Although consumption is punctual, the fact that agents have different preferred times $\hat{t_i}$ spreads the aggregate demand across the day:

$$D_i(t, t_i) = D_0 \mathbb{1}_{[t_i, t_i + \Delta t]}(t) \tag{33}$$

Each agent has a preferred time of consumption $\hat{t_i}$, drawn from a probability distribution $\rho(\hat{t_i})$, and seeks to minimize a cost function that combines: an economic cost, reflecting the deviation between consumption and production weighted by the electricity price and a comfort (or social) cost, which penalizes deviation from the preferred time.
We define the total cost function at time t as:

$$J_i(t_i, \hat{t}_i, t) = \underbrace{\int_t^T p(\tau)\left(D_i(t_i,\tau) - S_i(\tau)\right) d\tau}_{\text{economic cost}} + \underbrace{\lambda(t_i - \hat{t}_i)^2}_{\text{social cost}} \tag{34}$$

This price is determined by the system's average frequency deviation:

$$p(\tau) = p_0 - \epsilon\langle\omega(\tau)\rangle \tag{35}$$

where the average frequency deviation evolves as:

$$\langle\omega(t)\rangle = \int_0^t e^{\frac{\kappa}{M}(\tau - t)} \frac{\langle S(t)\rangle - \langle D(t)\rangle}{M} d\tau \tag{36}$$

this expression is simply (9) imposing that $\langle\omega(t=0)\rangle = 0$ i.e, at the beginning of the day $\langle S(t)\rangle = \langle D(t)\rangle$. Here,

$$\begin{aligned} \langle S(t)\rangle &= \frac{1}{N}\sum_{i=1}^{N} S_i(t), \quad \text{mean production} \\ \langle D(t)\rangle &= \frac{1}{N}\sum_{i=1}^{N} D_i(t, t_i), \quad \text{mean consumption} \end{aligned} \tag{37}$$

describe the aggregate behavior of all agents. The price $p(\tau)$ is thus non-local, coupling all agents through their collective production and consumption. This interdependence creates a mean field game: each agent's optimal strategy depends on the population-wide averages $\langle P\rangle$ and $\langle C\rangle$, which in turn depend on all agents' actions. The price mechanism mediates this mean field interaction, linking individual decisions to global system dynamics via the frequency deviation $\langle\omega\rangle$.

While we do not implement the full MFG formalism, such as the HJB-FP system, our model captures the spirit of MFGs. Each agent's optimal decision t_i depends on the electricity price, which in turn is determined by the average frequency deviation of the grid, itself affected by the aggregate mismatch between consumption and production. This creates a mean-field-type coupling, where individual choices indirectly influence each other through global quantities. In this sense, our setting can be viewed as a simplified mean field game, centered around optimizing a scalar control variable per agent in a deterministic environment.
We aim to find the optimal time of consumption t_i^* for each agent. Since t_i is the control variable, we compute the derivative of the cost function with respect to t_i and set it to zero.
Let us denote the economic cost as:

$$I_i(t_i, t) = \int_t^T p(\tau)(D_i(t_i,\tau) - S_i(\tau))d\tau \tag{38}$$

We differentiate I_i with respect to t_i:

$$\begin{aligned} \frac{dI_i}{dt_i} &= \int_t^T \frac{d}{dt_i}\left[(p_0 - \epsilon\langle\omega(\tau)\rangle)(D_0 \mathbb{1}_{[t_i, t_i+\Delta t]}(\tau) - S_i(\tau))\right] d\tau \\ &= \int_t^T \Big[p_0 D_0(\delta(t_i - \tau) - \delta(t_i + \Delta t - \tau)) \\ &\quad - \epsilon\frac{d\langle\omega\rangle}{dt_i}(D_0 \mathbb{1}_{[t_i, t_i+\Delta t]}(\tau) - S_i(\tau)) \\ &\quad - \epsilon D_0(\delta(t_i - \tau) - \delta(t_i + \Delta t - \tau))\langle\omega\rangle(\tau)\Big] d\tau \end{aligned} \tag{39}$$

We can simplify this expression using the fundamental hypothesis of MFG. The mean field $\langle\omega\rangle$ does not depend on the individual strategies, so that the derivative $d\langle\omega\rangle/dt_i \simeq 0$. In fact, the average frequency depends on $\langle D(t)\rangle = \frac{D_0}{N}\sum_i \mathbb{1}_{[t_i^*, t_i^*+\Delta t]}(t) = \mathcal{O}(1/N)$, that is negligible for large N.
Evaluating the Dirac deltas gives:

$$\begin{aligned}\frac{dI_i}{dt_i} &= p_0 D_0(\mathbb{1}_{[t,T]}(t_i) - \mathbb{1}_{[t,T]}(t_i+\Delta t)) \\ &\quad - \epsilon D_0(\mathbb{1}_{[t,T]}(t_i)\langle\omega(t_i)\rangle - \mathbb{1}_{[t,T]}(t_i+\Delta t)\langle\omega(t_i+\Delta t)\rangle)\end{aligned} \tag{40}$$

Now we differentiate the full cost function:

$$\frac{d}{dt_i}J_i = \frac{dI_i}{dt_i} + 2\lambda(t_i - \hat{t}_i) \tag{41}$$

We evaluate this at the optimal point $t_i = t_i^*$, where the derivative vanishes:

$$\begin{aligned}0 = 2\lambda(t_i^* - \hat{t}_i) + \underbrace{p_0 D_0(\mathbb{1}_{[t,T]}(t_i^*) - \mathbb{1}_{[t,T]}(t_i^*+\Delta t))}_{A.1} \\ \underbrace{-\epsilon D_0(\mathbb{1}_{[t,T]}(t_i^*)\langle\omega(t_i^*)\rangle - \mathbb{1}_{[t,T]}(t_i^*+\Delta t)\langle\omega(t_i^*+\Delta t)\rangle)}_{A.2}\end{aligned} \tag{42}$$

We impose that all agents have to take their shower during the day i.e, $t_i \in [0, T-\Delta t]$ so that the shower starts and ends the same day, with T =24h. With this assumption we simplify each term:

- Term (A.1): Since $t_i^* \in [0, T-\Delta t]$, both t_i^* and $t_i^* + \Delta t$ are in $[t, T]$, hence:

  $$\mathbb{1}_{[t,T]}(t_i^*) = \mathbb{1}_{[t,T]}(t_i^*+\Delta t) = 1 \Rightarrow \text{(A.1)} = 0$$

  This term corresponds to the base energy cost at a fixed reference price. Since this price is constant and does not depend on the agent's decision, it vanishes in the derivative and does not influence the optimization.

- Term (A.2): Again, both indicator functions equal 1, so:

  $$\text{(A.2)} = -\epsilon D_0(\langle\omega(t_i^*)\rangle - \langle\omega(t_i^*+\Delta t)\rangle)$$

  This term reflects the sensitivity of the agent's cost to changes in the average frequency over time. Although the agent does not control the average frequency directly, this term encourages them to anticipate future variations and adapt their consumption accordingly.

Finally, the optimality condition becomes:

$$0 = 2\lambda(t_i^* - \hat{t}_i) - \epsilon D_0(\langle\omega(t_i^*)\rangle - \langle\omega(t_i^*+\Delta t)\rangle) \tag{43}$$

It represents quite well the competition in the optimization between the social term (in $\hat{t}_i$) and the economical term related to the consumption and the variation of the frequency of the system.
We have N equations (43), one for each agent and one equation (36) that defines the frequency and couples the N previous equations together. We have a system of self consistent equations that we solve numerically.

Numerical simulations

The idea to implement this optimization process is to proceed iteratively. First, we optimize the N equations independently starting from $t_i = \hat{t}_i$ for all i. Once the t_is are computed, we inject them into

$\langle D(t)\rangle$ to compute a new $\langle \omega(t)\rangle$ and to start a new optimization with the updated average frequency deviation. The average consumption is defined as

$$\langle D(t)\rangle = \int_{t-\Delta t}^{t} f(\tau)d\tau \tag{44}$$

where f is the density of t_i: time at which the consumption starts. To ensure global stability, we impose that the consumption over the day should equal to the daily production i.e:

$$\int_0^T \langle D(t)\rangle dt = \int_0^T \langle S(t)\rangle dt \tag{45}$$

To model a large number of agents in a tractable way, we group together agents who exhibit identical behavior. Specifically, two agents with the same preferred consumption time $\hat{t}_i$ and the same initial decision t_i share the same starting point and optimization objective. Since our model is fully deterministic, there is no randomness in production or in agent decisions, these agents will follow the same optimal strategy t_i^*.

To implement this grouping, we discretize the time interval $[0, T]$ into n equal subintervals of length δT, such that $n = T/\delta T$. Each group of agents is then characterized by a common preferred time $\hat{t}_i = i\delta T$, for some integer $k \in \{0, \ldots, n-1\}$. To preserve the information about the distribution of preferences across the population, we weight each group by the density of agents whose preferred times fall within the interval associated to that group. Formally, the aggregated distribution $\rho(\hat{t})$ of preferred consumption times is approximated by:

$$\rho(\hat{t}) = \sum_i \rho(i\delta T)\, \mathbb{1}_{\left[(i-\frac{1}{2})\delta T,\ (i+\frac{1}{2})\delta T\right]}(\hat{t}), \tag{46}$$

where $\rho(i\delta T)$ represents the weight (or density) of agents with preferred time around $i\delta T$. This discretization enables us to reduce the continuous distribution of preferences into a finite number of representative groups, making numerical treatment more feasible. We apply the same discretization approach to the density of strategies ff, which represents the distribution of actual consumption start times chosen by agents. Just like the preferred times $\hat{t}_i$, we assume that optimal strategies can be grouped into discrete bins. The strategy of the i^{th} group is denoted t_i, corresponding to the common consumption time selected by all agents in that group.

The distribution of preferred times ρ is assumed to be known. For each group, we iteratively determine the optimal strategy t_i^* that minimizes the individual cost function:

$$t_i^* = \arg\min_{t_i} J_i(t_i, \hat{t}_i, t) \quad \text{with} \quad 0 = 2\lambda(t_i^* - \hat{t}_i) - \epsilon D_0(\langle\omega\rangle(t_i^*) - \langle\omega\rangle(t_i^* + \Delta t)) \tag{47}$$

With these optimal times, we compute the strategy distribution f at each iteration. However, directly updating f using only the newly computed t_i^* often leads to oscillatory behavior. To address this, we introduce a relaxation mechanism: at each iteration, we update the strategy distribution as a combination of the previous distribution and the new optimal one. Specifically, only a proportion $\alpha \in [0, 1]$ of agents is assumed to adopt the newly computed optimal strategy at each step, while the remaining $1 - \alpha$ retain their previous strategy. This results in the following update rule:

$$f_{n+1} = (1 - \alpha)f_n + \alpha f^* \tag{48}$$

where f^* is the distribution corresponding to the optimal strategies t_i^* computed at iteration n, and f_n is the current distribution. This update slightly alters the initial assumption that all agents sharing a common $\hat{t}_i$ adopt the same final strategy t_i^*, but the relaxation helps ensure smoother convergence of the fixed-point algorithm.

5 Conclusion

Throughout this work, I first studied the structure and dynamics of power grids to better understand the physical system underlying our study. I then investigated the stability of electrical networks, initially under constant power injections and later under time-dependent inputs, incorporating a feedback mechanism based on electricity price. Once the system's behavior was well understood, the focus shifted to an optimization problem, still outside the Mean Field Game framework. Under the mean field assumption, I explored how agents with fixed production could optimally adapt their consumption. This adaptation balances two competing effects: an economic cost, which encourages stability by depending on system frequency, and a social cost, which penalizes changes in consumption habits. This sets the stage for a more complete Mean Field Game formulation in future work.

I would like to thank the LPTMS and the CEA for welcoming me and giving me the opportunity to complete this internship. I would like to sincerely thank Denis for his guidance and support throughout my internship, as well as for his warm welcome. I am also grateful to Cosimo for assisting me with the administrative procedures at the CEA. A special mention goes to Guillaume Roux, who, although not officially my supervisor, was always available to discuss ideas and provide valuable insights. Thanks also to the entire project team for their support throughout my stay. Finally, I want to thank Yacine, with whom I worked closely on this project; our collaboration was both productive and enjoyable.

Appendices

A Swing equation

A.1 Equation of motion describing generator rotor dynamics: the swing equation

A synchronous generator consists of a rotor that is driven by a mechanical torque $\tilde{T_m}$ to produce electrical power through electromagnetic induction. The electrical power demanded by the grid exerts an opposing torque on the rotor, known as the electrical torque T_e, which typically acts as a decelerating torque due to Lenz's law — the system naturally opposes changes in magnetic flux.

To describe the rotor's dynamics, we apply Newton's second law for rotational systems, which states that the rate of change of angular momentum is equal to the net torque acting on the rotor. This leads to:

$$\frac{d}{dt}\left(J\dot{\theta}\right) = \tilde{T_m} - D_m\omega - \frac{1}{R}\theta\omega - D_e\theta\omega - T_e \tag{49}$$

where J is the moment of inertia, θ is the angle of the rotor relative to to a frame rotating at the reference frequency ω_R, ω is the the angular frequency of the rotor and so $\dot{\theta} = \omega - \omega_R = \theta\omega$ the frequency deviation. D_m is the mechanical damping coefficient due to mechanical friction and D_e is the electrical damping coefficient arising from the electrical losses from the generator and R is the regulation parameter that characterizes the governor's frequency-sensitive feedback control, where it adjusts the generator output power proportionally to measured frequency deviations.
In [5], T.Nishikawa and al rewrite the equation as

$$J\ddot{\theta} + \tilde{D}\dot{\theta} = T_m - T_e \tag{50}$$

where $\tilde{D} := D_m + D_e + 1/R$ and $T_m = \tilde{T_m} - D_m\omega$ is the net mechanical torque. Multiplying both sides by the reference frequency and using the fact that the torque (in N.m) multiplied by a frequency (in rad/s) gives a power in watts. We get an equation in terms of power:

$$J\omega_R\ddot{\theta} + \tilde{D}\omega_R\dot{\theta} = \frac{\omega_R}{\omega}(T_m\omega - T_e\omega) \approx \tilde{Pm} - \tilde{Pe} \tag{51}$$

We assumed that the ration ω_R/ω is close to one. This approximation is valid in our context of studying the stability of the power system.
We then normalize the previous equation by the reference power P_R to make $\tilde{Pm}$ and $\tilde{Pe}$ *per unit* (p.u) quantities. The factor $J\omega_R$ becomes $2H/\omega_R$ with the inertia constant $H = W/P_R$ and $W := J\omega_R^2/2$ the kinetic energy of the rotor. The term $\tilde{D}\omega_R$ becomes D/ω_R with $D := \tilde{D}\omega_R/P_R$. We then obtain the swing equation:

$$\frac{2H}{\omega_R}\ddot{\theta} + \frac{D}{\omega_R}\dot{\theta} = P_m - P_e \tag{52}$$

The term P_m represents the net mechanical power input to the rotor and P_e the electrical power demanded by the network.
Redefining the terms in the equation we get (1).

A.2 Derivation of the swing equation

In power grids, the voltage at each node k is represented as a phasor:

$$\underline{\mathbf{V}}_k = V_k e^{i\theta_k} \tag{53}$$

where V_k is the voltage magnitude and θ_k is the phase angle. The transmission line between nodes k and l is characterized by its impedance z_{kl} and its admittance y_{kl}:

$$z_{kl} = r_{kl} + ix_{kl} \quad , y_{kl} = g_{kl} + ib_{kl} = \frac{1}{z_{kl}} \tag{54}$$

Using Ohm's law, the current from node k to l is:

$$\underline{\mathbf{I}}_{kl} = y_{kl}(\underline{\mathbf{V}}_k - \underline{\mathbf{V}}_l) \tag{55}$$

The total current injection at node k sums contributions from all connected lines:

$$\underline{\mathbf{I}}_k = \sum_l y_{kl}(\underline{\mathbf{V}}_k - \underline{\mathbf{V}}_l) = \left(\sum_l y_{kl}\right)\underline{\mathbf{V}}_k - \sum_l y_{kl}\underline{\mathbf{V}}_l \tag{56}$$

This can be expressed using the admittance matrix $\mathbf{Y}$ with elements:

$$Y_{kl} = \begin{cases} \sum_n y_{kn} & k = l \\ -y_{kl} & k \neq l \end{cases} \tag{57}$$

yielding the network equation in matrix form:

$$\underline{\mathbf{I}} = \mathbf{Y}\underline{\mathbf{V}} \tag{58}$$

The complex power at node k is:

$$\underline{S}_k = \underline{\mathbf{V}}_k \underline{\mathbf{I}}_k^* \tag{59}$$

$$= \underline{\mathbf{V}}_k \sum_l Y_{kl}^* \underline{\mathbf{V}}_l^* \tag{60}$$

$$= V_k^2 Y_{kk}^* - \sum_{l \neq k} V_k V_l Y_{kl}^* e^{i(\theta_k - \theta_l)} \tag{61}$$

For an ideal power grid with no loss in transmission $r_{kl} = 0$ (for high-voltage lines $r_{kl} \ll x_{kl}$) [1], so we consider:

$$Y_{kl} \approx iB_{kl} \tag{62}$$

The active power becomes:

$$P_k = \Re\{\underline{S}_k\} \approx \sum_l V_k V_l B_{kl} \sin(\theta_k - \theta_l) \tag{63}$$

The power balance at each generator node gives:

$$M_k \ddot{\theta}_k + \kappa_k \dot{\theta}_k = P_k^{mech} - P_k^{elect} \tag{64}$$

Substituting the active power yields:

$$M_k \frac{d^2\theta_k}{dt^2} + \kappa_k \frac{d\theta_k}{dt} = P_k^{mech} + \sum_l K_{kl} \sin(\theta_l - \theta_k) \tag{65}$$

where the coupling coefficients are:

$$K_{kl} = V_k V_l |B_{kl}| \tag{66}$$

B Computation of the pseudo inverse

B.1 Reduction of the Laplacian

To reduce the Laplacian $\mathbf{L}$, we impose a constraint on the degrees of freedom. A common choice is to fix the sum of all angles to zero:

$$\sum_{j=1}^{N} \theta_j = 0$$

This eliminates the global rotational invariance of the system. In particular, we can express θ_1 in terms of the other angles:

$$\theta_1 = -\sum_{j=2}^{N} \theta_j$$

Substituting this relation into the fixed point equation (13),

$$\sum_{k=1}^{N} L_{jk}\theta_k^* = -P_j,$$

we obtain, for each row j:

$$\begin{aligned} L_{j1}\theta_1^* + \sum_{k=2}^{N} L_{jk}\theta_k^* &= -P_j \\ L_{j1}\left(-\sum_{l=2}^{N}\theta_l^*\right) + \sum_{k=2}^{N} L_{jk}\theta_k^* &= -P_j \\ \sum_{k=2}^{N} \left(L_{jk} - L_{j1}\right)\theta_k^* &= -P_j \end{aligned}$$

This procedure effectively removes one degree of freedom and yields a new system of dimension $(N-1)\times(N-1)$. In matrix terms: we remove the first row and column of $\mathbf{L}$ (corresponding to θ_1), and subtract the first column from all remaining columns, as shown by the computation above.

To remain consistent, we also remove the first component of the power vector $\underline{P}$. This is valid because $P_1 = -\sum_{j=2}^{N} P_j$, and similarly, θ_1 can be recovered from the reduced solution via:

$$\theta_1 = -\sum_{j=2}^{N} \theta_j.$$

B.2 Moore Penrose pseudo-inverse

An alternative to reducing the Laplacian is to use the Moore–Penrose pseudo-inverse, denoted by $\mathbf{L}^\dagger$. The fixed point equation to solve is:

$$\mathbf{L}\underline{\theta}^* = -\underline{P}$$

Since $\mathbf{L}$ is not invertible, there are infinitely many solutions. The Moore–Penrose pseudo-inverse $\mathbf{L}^\dagger$ gives the minimum-norm solution among them:

$$\underline{\theta}^* = -\mathbf{L}^\dagger \underline{P}$$

This solution minimizes $\|\underline{\theta}^*\|_2^2$ subject to the constraint $\mathbf{L}\underline{\theta}^* = -\underline{P}$.
To compute $\mathbf{L}^\dagger$, one typically performs the spectral decomposition of $\mathbf{L}$:

$$\mathbf{L} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$$

where $\mathbf{\Lambda} = \mathrm{diag}(\lambda_1, \ldots, \lambda_N)$ contains the eigenvalues of $\mathbf{L}$, with $\lambda_1 = 0$ corresponding to the eigenvector $\mathbf{u}_1 = \frac{1}{\sqrt{N}}\mathbf{1}$. The matrix $\mathbf{U}$ is an orthogonal matrix whose columns $\mathbf{u}_1, \ldots, \mathbf{u}_N$ form an orthonormal basis of eigenvectors of $\mathbf{L}$. The pseudo-inverse is then given by:

$$\mathbf{L}^\dagger = \mathbf{U}\mathbf{\Lambda}^\dagger\mathbf{U}^\top$$

where $\mathbf{\Lambda}^\dagger = \mathrm{diag}(0, \lambda_2^{-1}, \ldots, \lambda_N^{-1})$, i.e., we invert only the nonzero eigenvalues and set the inverse of the zero eigenvalue to zero.

The key point is that $\mathbf{L}$ has a one-dimensional kernel spanned by $\mathbf{1}$, reflecting invariance under uniform phase shifts. Rather than explicitly removing this degree of freedom (e.g., by fixing a reference angle or imposing $\sum_j \theta_j = 0$), the pseudo-inverse implicitly projects the solution onto the subspace orthogonal to Ker($\mathbf{L}$).

This projection effectively reduces the problem to one of dimension $N-1$, ensuring a unique solution that satisfies the physical constraint of phase consistency, without the need to arbitrarily fix one of the phases.

The key point is that $\mathbf{L}$ has a one-dimensional kernel spanned by $\mathbf{1}$, reflecting invariance under uniform phase shifts. Since there are infinitely many solutions to $\mathbf{L}\underline{\theta} = -\underline{P}$, selecting the solution with minimal Euclidean norm corresponds to projecting onto the subspace orthogonal to Ker($\mathbf{L}$). In other words, the Moore–Penrose pseudo-inverse selects, among all possible solutions, the unique one with minimal norm, which necessarily lies in the orthogonal complement of the kernel of $\mathbf{L}$. This projection implicitly reduces the problem to one of dimension $N-1$ and enforces a phase consistency condition (such as $\sum_j \theta_j = 0$) without requiring it to be fixed explicitly.

C Reconstruction of the analytical solution

The equation we want to solve is

$$\partial_t \underline{X} = \mathbf{A}\underline{X} + B \quad \text{with} \quad \underline{X} = \sum_k \alpha_k(t)\underline{R_k} \tag{67}$$

where $\underline{R_k}$ is the k^{th} right eigenvector of $\mathbf{A}$. The goal is to reconstruct $\underline{X}$ so to find the coefficients $\alpha_k(t)$. We have

$$\sum_k \dot{\alpha_k}(t)\underline{R_k} = \sum_k \alpha_k(t)\lambda_k\underline{R_k} + \sum_k \beta_k\underline{R_k} \quad \text{with} \quad \lambda_k \text{ the } k^{th} \text{ eigenvalue} \tag{68}$$

The coeffcients are detemined by the folowing differential equation

$$\dot{\alpha_k}(t) = \lambda_k\alpha_k(t) + \beta_k \tag{69}$$

And its solution is

$$\alpha_k(t) = \begin{cases} \alpha_k^0 + \beta_k t, & \text{if } \lambda_k = 0 \\ \left(\alpha_k^0 + \frac{\beta_k}{\lambda_k}\right) e^{\lambda_k t} - \frac{\beta_k}{\lambda_k}, & \text{otherwise} \end{cases}$$

To reconstruct the solution, we need to determine the β_k and α_k^0 coefficients.

1. **Initial coefficients α_k^0**
   We denote $\underline{L_k}$ the k^{th} left eigenvector of the matrix $\mathbf{A}$. To find the initial coefficients α_k^0, we project our initial state onto the left eigenvector:

   $$\underline{L_k}^\top \underline{X}^0 = \sum_{k'} \alpha_{k'}^0 \underbrace{\underline{L_k}^\top \underline{R_{k'}}}_{=\mu\delta_{kk'}} \tag{70}$$

   Since $\mathbf{A}$ is not symmetric, the basis of left and right eigenvectors is orthogonal but not orthonormal (this is why a factor μ appears in the scalar product above). We need to normalize the eigenvectors so that $\underline{L_k}^\top \underline{R_{k'}} = \delta_{kk'}$, allowing direct projection of the initial state onto the eigenvectors to obtain the coefficients. For all k, we impose

   $$\mathbf{L}_k^\top \mathbf{R}_k = 1 \quad \Rightarrow \quad \mathbf{R}_k \leftarrow \frac{\mathbf{R}_k}{\sqrt{\mathbf{L}_k^\top \mathbf{R}_k}}, \quad \mathbf{L}_k \leftarrow \frac{\mathbf{L}_k}{\sqrt{\mathbf{L}_k^\top \mathbf{R}_k}^*}$$

   Now that the basis of eigenvectors are orthonormal, we can easily extract the coefficients α_k^0 needed to reconstruct the solution.

2. **Projection of the source term on the eigenvectors β_k**

$$\begin{aligned} & \underline{P} \in \mathbb{R}^N, \\ & \underline{B} = \begin{bmatrix} \underline{0}_N \\ \mathbf{M}^{-1}\,\underline{P} \end{bmatrix} = \sum_k \beta_k\,\underline{R_k} \\ & \underline{B} = \mathbf{R}\,\underline{\beta} \quad \text{with } \mathbf{R} \text{ the matrix of right eigenvectors} \\ \Rightarrow \quad & \underline{\beta} = \mathbf{R}^{-1}\,\underline{B} \end{aligned} \tag{71}$$

We extracted the α_k^0 and the β_k, so we have all that we need to express $\alpha_k(t)$ solution of (69). And the complete solution is

$$\underline{X}(t) = \sum_{k=1}^{2N} \alpha_k(t)\underline{R}_k \tag{72}$$

References

[1] B. Schäfer, "Decentral smartt grid control," *New Journal of Physics*, vol. 17, no. 015002, 2015.

[2] D. Witthaut, "Collective nonlinear dynamics and self-organization in decentralized power grids," *Reviews of Modern Physics*, vol. 94, no. 1, p. 015005, 2022.

[3] J.-M. Lasry and P.-L. Lions, "Jeux à champ moyen. i - le cas stationnaire," *Comptes Rendus Mathematique*, vol. 343, no. 9, pp. 619–625, 2006. C. R. Acad. Sci. Paris, Sér. I.

[4] M. Huang, R. P. Malhamé, and P. E. Caines, "Large population stochastic dynamic games: closed-loop mckean-vlasov systems and the nash certainty equivalence principle," *Communications in Information & Systems*, vol. 6, no. 3, pp. 221–252, 2006.

[5] T. Nishikawa and A. E. Motter, "Comparative analysis of existing models for power-grid synchronization," *New Journal of Physics*, vol. 17, no. 015012, 2015.

[6] P. Bastard, M. Hennebel, J. Dai, M. Petit, and T. D. Le, "Fonctionnement et réglages des systèmes de transport et de distribution d'électricité en régime permanent," 2020. Support de cours, CentraleSupélec.

[7] C. Godsil and G. Royle, *Algebraic Graph Theory*, vol. 207 of *Graduate Texts in Mathematics*. Springer, 2001.

[8] A. Ben-Israel and T. N. Greville, *Generalized Inverses: Theory and Applications*. CMS Books in Mathematics, Springer, 2nd ed., 2003.

[9] D. Ullmo, I. Swiecicki, and T. Gobron, "Quadratic mean field games," *Physics Reports*, vol. 799, pp. 1–35, 2019.