NBER WORKING PAPER SERIES

NUCLEAR OPERATIONS WITH A HIGH PENETRATION OF RENEWABLES: THE CASE OF FRANCE

Nicolas Astier
Frank A. Wolak

Working Paper 34662
http://www.nber.org/papers/w34662

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2026

Thanks to Louis Cialti for his assistance in collecting environmental data associated with nuclear plant operations. We are grateful for discussions with Tim Schittekatte, Charlotte De Cannière, Ömer Karaduman, Leticia Pieraerts, Claire Lim and Aleksandar Zaklan, as well as to seminar and conference participants at the Toulouse School of Economics, the IFN Stockholm Conference on Sustainable Energy Transition, the 11th FAERE Conference, the Helsinki Graduate School of Economics, the IAEE Session at the 2025 ASSA Meetings, the UC Louvain Workshop on Electricity Market Design, CEA Saclay, the 32nd NBER-TCER-CEPR (TRIO) Conference, the 13th Mannheim Conference on Energy and the Environment, the 46th IAEE International Conference, the 1st Alpine Workshop on "Energy Economics and Policy", the Paris School of Economics and the 2025 CEPR Paris Symposium for their feedback and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Nuclear Operations with a High Penetration of Renewables: The Case of France
Nicolas Astier and Frank A. Wolak
NBER Working Paper No. 34662
January 2026
JEL No. L94, Q2, Q42

ABSTRACT

Nuclear and intermittent renewables (wind and solar) are generally regarded as the only scalable technologies producing low-carbon electricity. However, the extent to which these technologies can co-exist in a reliable power system depends on whether nuclear units can adjust their operations to renewable output fluctuations. Using hourly data from the French power system, we find that nuclear units are operated quite flexibly, and that the foregone energy production due to "load following" actions (relative to the counterfactual of operating at full capacity during load following events) is currently limited. However, we find that an additional load following event is associated with a slightly higher likelihood of a unit failure. We also find that unit-level minimum output constraints are binding more frequently as system-wide renewable generation increases, especially so for units most exposed to solar generation. In 2024, hours during which available nuclear flexibility was exhausted are associated with non-positive hourly day-ahead prices.

Nicolas Astier
Paris School of Economics
nicolas.astier@psemail.eu

Frank A. Wolak
Stanford University
Department of Economics
and NBER
wolak@zia.stanford.edu

1 Introduction

Nuclear power and intermittent renewables (wind and solar) are generally regarded as the only mature technologies that have the potential to generate the amounts of electricity required to decarbonize global energy consumption.[1] While many industrialized countries have chosen to focus on the deployment of wind and solar generation resources, especially after the 2011 Fukushima disaster (Davis, 2012; Joskow and Parsons, 2012), we currently observe a renewed interest in nuclear power.[2] Therefore, future low-carbon power systems may, at least in some jurisdictions, include large installed capacities of both intermittent renewables and nuclear power.

Because of their technical complexity and economic characteristics (high construction costs, low variable operating costs), most nuclear reactors around the world are used to produce a steady output of electricity, fueling the perception that they are inherently inflexible assets. In contrast, wind and solar produce electricity intermittently, frequently requiring other generation technologies to ramp up and down for short periods of time whenever renewable output vanishes or booms. From a purely operational standpoint, it is therefore unclear whether nuclear, wind and solar can jointly provide reliable low-carbon electricity.

This paper explores this issue empirically, leveraging the fact that France is a jurisdiction whose electricity supply is already heavily relying on both nuclear and renewables.[3] While the lion share (70 to 80%) of France's low-carbon electricity comes from nuclear power, the penetration of wind and solar has been steadily growing, reaching 14% of total generation in 2024 (RTE, 2025). The two technologies are now comparable in terms of potential maximum instantaneous output (63 GW for nuclear and about 50 GW for wind and solar), implying that one may expect them to strongly interact during some hours.

We address three main research questions. First, to what extent are nuclear units able to quickly decrease (and increase) their instantaneous power output, an operational mode called "load-following," as a response to short-term fluctuations in wind and solar output? Second,

[1] Industrialized countries are usually considered to have exhausted the main locations suitable for large-scale hydropower. In addition, the cost-effectiveness and/or technological feasibility of alternative low-carbon power generation technologies, such as carbon capture or nuclear fusion, remains to be shown.

[2] For example, countries such as France, the United Kingdom or Poland have announced plans to build new nuclear reactors.

[3] Over the past two decades, domestic generation from low-carbon power plants (nuclear, hydro, wind and solar) has on average exceeded domestic electricity demand by 2% in France (Astier, 2025).

what are the impacts of load-following in terms of loss of capacity factor, environmental externalities and safety associated with the production of nuclear energy? Third, what are the relevant operating constraints that limit the ability nuclear units to load-follow and how do they interact with an increase in renewable output?

These research questions are of first-order importance for the European power system, as recently highlighted in a *Financial Times* article (Millard, 2025). The operating constraints of power plants, and in particular start-up costs for thermal power plants, have been shown to have significant implications for short-term electricity market design (Reguant, 2014), and to strongly interact with the growing penetration of intermittent renewables (Jha and Leslie, 2025). In this paper, we also highlight that one type of non-convexity, namely "minimum output constraints," may exacerbate the degree of substitution between nuclear and renewables, possibly limiting their joint ability to provide a reliable supply of low-carbon electricity (see Section 2).

Our empirical analysis relies on detailed public data on the French power system over the past decade (see Section 3). This includes the characteristics of French nuclear units, hourly aggregate outcomes of the French power system (2012–2024), hourly generation and outage information for nuclear units (2015–2024), dynamic information on minimum output constraints (2021–2024) and yearly plant-level environmental reports (2019–2024). We employ different empirical strategies to address our three research questions.

First, we use system-level hourly data to assess the prevalence of nuclear load-following in the French power system (see Section 4). We regress hourly aggregate nuclear output on hourly wind and solar output, as well as hourly residual demand,[4] for the sample period 2012–2024. In contrast to the perceived "inflexibility" of nuclear power plants, we estimate a large coefficient of substitution between renewables and nuclear. Most of the effect remains when controlling for day-of-sample fixed effects, suggesting that most of this substitution is due to "load-following", that is, short-duration decreases in the output of a subset of nuclear units. Specifically, the day-of-sample fixed effects regression finds that an additional MWh of renewable generation predicts a 0.66 MWh decrease in nuclear output during the hour.

Second, we explore the consequences of the flexible operation of French nuclear units

[4]We define residual demand as gross domestic consumption plus net exports minus net hydropower generation (because pumped-storage hydropower plants can both consume and produce electricity) and must-take biothermal (see Section 4).

using hourly unit-level data for 2015 to 2024. To do this, we define two quantitative measures of flexible unit operation: (i) a count of load-following events, and (ii) the amount of (full-load hours of) "lost energy", which we define as the additional energy that would have been produced if a load-following nuclear unit had instead produced at its maximum output during each load-following event. An obvious implication of load-following events is a decrease in how intensively a nuclear unit is used during a fuel cycle. This usage intensity is measured by the capacity factor of a generation unit, defined as the ratio between the average and maximum possible power generation during a pre-specified time period. We find the (mechanical) direct reduction in the capacity factor of nuclear units due to load-following to be small, amounting to at most a few percentage points.[5] In addition, we are not able to detect any impact of nuclear load-following on environmental outcomes, at least based on an empirical analysis of monthly plant-level observations. We also investigate indirect channels that may have an impact on the capacity factor of nuclear units, namely fuel efficiency, the duration of maintenance, and planned/forced outages. Only the latter is found to be significantly associated with load-following. This relationship is, however, noisy, due to the large heterogeneity in the cause and duration of outages. In addition, the occurrence of forced outages may be imperfectly observed since it is measured in a context of asymmetric information (Hausman, 2014; Bizet et al., 2022). We therefore next restrict attention to "emergency" forced outages, namely automatic shutdowns and manual emergency interventions, which are unlikely to be concealed to the regulator. We estimate proportional hazards models (see Section 5) and find that an increase of five full-load hours (GWh/GW) of "lost energy" (or, equivalently, one additional load-following event) is associated with a 1% increase in the hazard rate of an "emergency" forced outage, which implies a higher probability of such an outage occurring during the unit's fuel cycle.

Third, we explore the interaction between intermittent renewables and the operating constraints that limit the ability of nuclear units to load-follow (see Section 6). Most importantly, like other thermal power plants, a nuclear unit must produce at or above a minimum level of output to maintain stable operations of the generation unit and transmission network. Most of the French nuclear units face a "nominal" minimum output constraint

[5] Of course, the average capacity factor differs from the average economic value of the electricity produced. This caveat is for example well-known for renewables, for which the decrease in the levelized-cost of energy may over-estimate the increase in the social value of wind or solar (Joskow, 2011).

(MOC) of 20 percent of the unit's nameplate capacity and the remainder have a MOC equal to 25 percent of the unit's nameplate capacity. Using panel regressions with unit-hour observations for 2015–2024, we find that wind and solar output are associated with increasingly-binding nominal MOCs. Specifically, we estimate that, on average, the nominal MOC of one additional nuclear unit starts binding in a given hour for every 3 GWh increase in wind and solar output. We next compute a plant-by-day-of-sample metric capturing the extent to which each nuclear unit is "exposed" to the installed capacity of other generating technologies (wind, solar, hydro and thermal). Using this wind and solar exposure metric, we find empirical evidence that nuclear units located closer to solar power plants are more likely to reach their nominal MOC. This result suggests that, with increasing solar generation, grid constraints may place additional restrictions on the choice of which nuclear unit(s) the system operator must ramp down. Conversely, when renewable output is low, grid constraints may force the system operator to temporarily set the minimum output constraint of a nuclear unit above its nominal MOC to maintain grid stability. Such "dynamic" MOCs, which are publicly disclosed since June 2021, may also arise for other considerations than grid stability. We observe dynamic MOCs to be binding much more frequently than nominal MOCs, in large parts due to the fuel management constraints faced by nuclear units. When accounting for both nominal and dynamic MOCs, we show that, in 2024, hours during which the available downward flexibility of the nuclear fleet was exhausted are strongly associated with zero or negative day-ahead prices. In other words, we find that the very large increase in the occurrence of zero or negative day-ahead prices that France experienced in 2024 (such prices occurred about 6% of the time or more than 500 hours during the year), results from many hours showing a combination of both high intermittent renewable output and depleted load-following capability of the nuclear fleet.

Our empirical results paint a nuanced answer to the question of whether nuclear, wind and solar can jointly provide reliable low-carbon electricity. On the one hand, in contrast to widely held views, we show that nuclear units can be operated quite flexibly, and that the associated lost energy sales seem to be currently limited. Of course, France may represent an optimistic case study, with a large nuclear fleet operated by a single operator with decades of experience with load-following operations. Yet, the observed amount of load following is well within the specifications of the so-called European Utility Requirements that all modern nuclear reactors meet. On the other hand, even in this favorable setting, we find evidence suggesting that the nuclear fleet has exhausted its load-following potential for a non-negligible number of hours in 2024. During such hours, the interaction between high renewable generation and a non-convexity in the supply function of nuclear units, known as the minimum output constraint, induces non-positive spot prices, and thus, in the absence of sufficient alternative sources of flexibility such as large-scale storage or demand response, a strong substitution between nuclear and renewables.[6]

While the importance of nuclear load-following has been widely discussed in industry and institutional reports (e.g. Bruynooghe et al. (2010); Lokhov (2011); Grünwald and Caviezel (2017); Morilhat et al. (2019); OECD (2021)), and explored using model simulations in engineering studies (e.g. Jenkins et al. (2018); Loisel et al. (2018); Lynch et al. (2022); Blanchard and Massol (2025)), a detailed *ex post* analysis using data from actual operations has, to the best of our knowledge, only been attempted in the on-going work by Johannsen et al. (2025). This latter study, however, focuses on the United States electricity supply industry, where nuclear load-following represents an exceptional operating mode because the share of nuclear in total generation in all regions of the United States is significantly smaller than in France. The economics literature has otherwise studied nuclear energy from a number of perspectives, including the incentives of plant operators regarding cost-efficiency (Davis and Wolfram, 2012) or safety (Hausman, 2014), unilateral market power in offer-based markets (Davis and Hausman, 2016; Liski and Vehviläinen, 2018; Lundin, 2021), avoided environmental externalities (Severnini, 2017; Adler et al., 2020; Jarvis et al., 2022), and green industrial policy (Andersson and Finnegan, 2024; Makarin et al., 2024). We complement this body of work by (i) showing that nuclear units have the ability to load-follow, similar to other large steam turbine generation units, and (ii) quantifying the main trade-offs associated with unit owners engaging in load-following actions. In addition, we further document the importance of non-convexities in the aggregate "supply" function of energy in electricity markets. In particular, start-up and/or ramping costs have been shown to raise challenges for

[6]Note that, although nuclear load-following is of first-order importance to assess the extent to which nuclear and intermittent renewables may co-exist at a large scale in a power system, our results should not be interpreted as definitive evidence that nuclear power must, or must not, represent a large share of future low-carbon electricity mixes. Indeed, the answer to this question also critically depends on the construction costs of new nuclear units, a highly-debated topic (Grubler, 2010; Boccard, 2014; Rangel and Lévêque, 2015).

market design (Reguant, 2014) and market power investigations (Mansur, 2008), especially so with the increasing penetration of intermittent renewables (Jha and Leslie, 2025). We further show that minimum output constraints, yet another supply-side non-convexity in electricity markets, can have a major influence on market outcomes in a low-carbon power system given the intermittency of renewables.

The rest of the paper is organized as follows. Section 2 provides relevant background on nuclear operations and outlines the economic intuition behind the importance of the (in)flexibility of nuclear units. Section 3 describes our empirical setting and data sources. Section 4 defines load-following operating behavior by nuclear units and discusses alternative ways to quantify its intensity. Section 5 estimates the impacts of load-following in terms of loss of capacity factor, environmental externalities and safety. Section 6 turns to the operating constraints limiting the ability to load-follow, and investigates how they interact with intermittent renewables and market outcomes. Finally, Section 7 discusses the external validity and potential policy implications of our results. Section 8 concludes.

2 Nuclear Operations and Renewables: Technical Background and Economic Relevance

This Section first introduces useful background on nuclear power technology that underlies subsequent analyses, with a particular emphasis on minimum output constraints. We next discuss how such constraints may interact with intermittent renewables. In particular, we explain why this interaction has first-order implications for the economics of reliable low-carbon power systems.

2.1 Nuclear Operations

A nuclear power plant is a facility composed of one or several nuclear units, that is, individual installations able to produce electricity independently from each other. Each unit hosts a reactor where neutrons are used to break large radioactive atoms (e.g. uranium 235), a reaction known as nuclear fission. This reaction releases neutrons and heat. The former fuels a chain reaction by which the process is self-sustained. The latter is used, as in any other steam turbine power plant, to boil water,[7] and produce pressurized steam that is sent to a turbine to generate electricity. Nuclear units can reach an installed capacity of 1,000 MW and above, which is orders of magnitude larger than the typical size of a single wind turbine (about 2 MW) or solar panel (about 300 W).

Given their size, nuclear plants involve high construction and fixed costs, whose recovery is typically only possible through high utilization rates. In addition, a stable chain reaction requires a complex real-time monitoring of the net flow of neutrons (using chemicals and control rods) to prevent meltdowns. As a result, in the vast majority of jurisdictions, online nuclear units consistently produce at (or close to) their maximum capacity, an operating mode often referred to as "baseload" (see left panel on Figure 1).

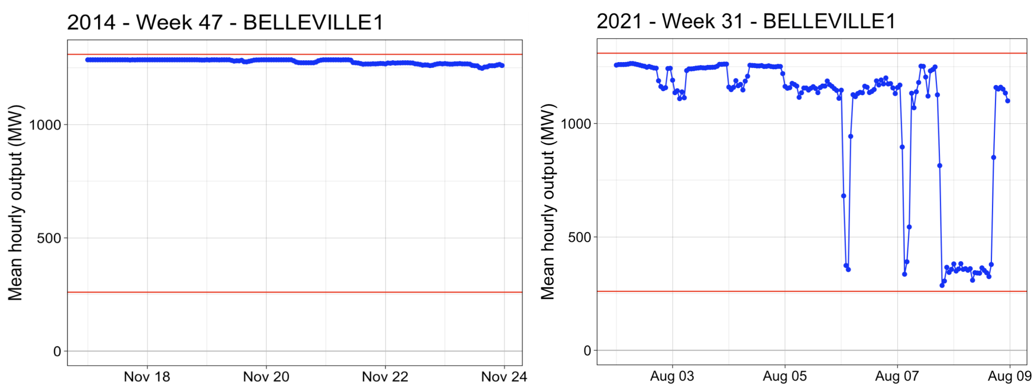

Figure 1: Hourly output of the French nuclear unit "Belleville 1" for two different weeks. Left panel: example of baseload operation. Right panel: example of load-following operation. The upper red horizontal line represents nominal capacity, and the lower red horizontal line the level of the nominal minimum output constraint.

Similarly to other thermal power plants, nuclear units have, nonetheless, the ability to ramp up and down, an operating mode known as "load-following" (see right panel on Figure 1). This operating mode is subject to two main constraints. First, ramping constraints mandate that the absolute rate of variation in output, called "ramp", should not exceed an engineering-determined upper bound. This upper bound is typically set to 3 to 5% of installed capacity per minute for nuclear units (Lokhov, 2011). As a result, ramping constraints allow for variations (the sum of the absolute value of output changes) exceeding 100% of installed capacity over the course of an hour and, as apparent from Figure 1, cannot be reliably inferred from hourly output data. They are therefore not studied in this paper. Second,

[7]Either directly in Boiling Water Reactor (BWR) designs or indirectly in Pressurized Water Reactor (PWR) designs. All commercial reactors in France are PWRs.

while a nuclear unit can ramp down when switching from "baseload" to "load-following" operations, its output must stay above an engineering-determined lower bound to maintain stable operations. We call this static lower bound the "nominal minimum output constraint" (nominal MOC) of the unit. For most French nuclear units, the nominal MOC is equal to 20% of the installed capacity of the unit (see Figure A.12).

Unlike other thermal (e.g. coal- or gas-fired) power plants, however, a nuclear plant operator cannot "top-up" the fuel contained in its reactor. Instead, the nuclear unit must be shut down for refueling, which has two main implications. First, the relevant timescale to study the economics of nuclear units corresponds to "fuel cycles", that is, sequences of a period of production followed by a refueling outage (see Section 5). Second, the chemical composition of the nuclear reactor changes over time as the fuel gets depleted. This evolution in turn limits the ability of a nuclear unit to load-follow when it reaches the end of its fuel cycle (Lynch et al., 2022). More generally, in any given hour, the enforced lower bound on the output of a given nuclear unit may be higher than the nominal MOC, either for unit-specific (e.g. balancing the chemical composition of the reactor) or system-specific (e.g. grid stability) reasons. When such situations arise, we call the enforced lower bound on output the "dynamic minimum output constraint" (dynamic MOC).

In economics, the minimum output constraint, either nominal or dynamic, represents a non-convexity in the supply function of a nuclear unit. The next paragraph highlights its critical role in shaping the economics of a power system that predominantly relies on nuclear and intermittent renewables.

2.2 Substitution between Nuclear and Renewables

Consider a power system whose electricity generation mix is exclusively composed of wind and solar on the one hand, with a zero marginal cost of production, and of nuclear on the other hand, with a positive marginal cost of production. Figure 2 illustrates the corresponding supply function for a given hour, also known as the "merit order" curve in the electricity industry. The horizontal green segment represents the total output from wind and solar for the considered hour, and the horizontal orange segment the total capacity of nuclear units that are up and running during that hour. Further assume that electricity demand (in blue) is lower that the sum of available renewable output and nuclear capacity, a situation that is expected to frequently occur with such an electricity mix (Mallapragada et al., 2021).

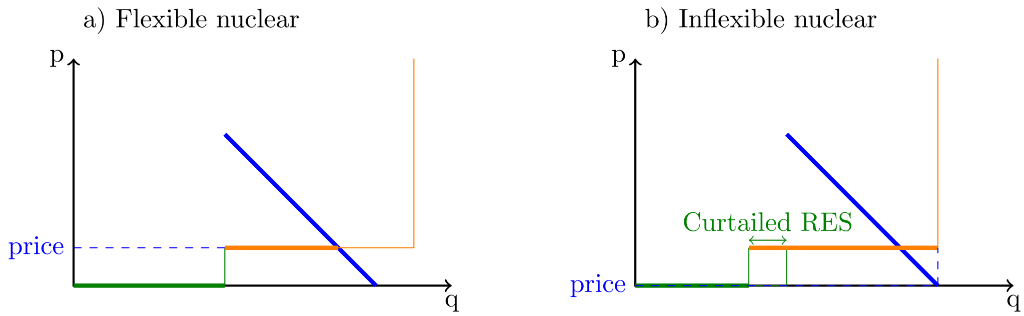

Figure 2: Market-clearing equilibrium prices when the nuclear fleet in either flexible (left panel) or inflexible (right panel) in the short run.

Two very different situations may then arise. On the one hand, if the nuclear fleet is flexible enough, meaning it can decrease its output while meeting the MOCs of all nuclear units, it will adjust its supply to demand (left panel on Figure 2). The competitive market price will then be the marginal cost of nuclear and no renewable output will be curtailed. On the other hand, if the nuclear fleet is inflexible, meaning it is unable sufficiently decrease output due to binding MOCs (right panel on Figure 2), renewables will then need to be curtailed and the competitive market price collapses to a non-positive value.[8]

In summary, the extent to which nuclear operations can be flexible has first-order implications for equilibrium market prices in a low-carbon power system, and thus for the long-term profitability of the different production technologies, as well as of other investments (batteries, electric vehicles, etc.). In addition, in the above example, nuclear load-following increases total surplus by avoiding renewable output curtailments. On the flip side, however, operating nuclear units in a flexible manner may increase the probability of outages and thus decrease overall reliability. Empirical evidence from actual operations is therefore needed to explore the trade-offs associated with nuclear load-following.

[8]Negative prices may arise if renewable support schemes imply that renewable producers incur a positive (private) opportunity cost when they do not generate electricity.

3 Empirical Setting: the French Power System

Our empirical evidence is drawn from the French power system, which we now briefly describe along with our main data sources, all of which are public.

3.1 Background on the French Power System

Electricity Mix

Electricity generating technologies in France may be divided into five broad categories: nuclear, hydropower, fossil-fueled thermal plants, wind and solar. Over the past decade, the installed capacities of these technologies have exhibited very different trends.

First, nuclear power represents the largest share of both installed capacity and generated electricity. In the 2010s, the nuclear fleet consisted of 19 power plants, hosting between two and six units, for a total of 58 units (see Figure A.13). The majority of units have a nameplate capacity of around 900 MW, while most remaining units have a nameplate capacity of around 1300 MW. Only the 4 most recent units, which were commissioned in the early 2000s, have a higher nameplate capacity of around 1500 MW. During our period of interest, two 900 MW units (Fessenheim power plant) were closed in 2020 and one 1600 MW unit (Flamanville 3) was commissioned in December 2024.[9] The current fleet therefore consists of 18 plants composed of 57 units, for a total capacity of 63 GW.

Similarly to nuclear, the installed capacities of hydro and fossil-fueled thermal power plants has remained fairly flat over the past decade.[10] In contrast, the installed capacities of wind and solar have been steadily growing. Installed wind capacity increased from 6.7 GW as of 31 December 2011 to about 24 GW as of 31 December 2024. Solar photovoltaic capacity increased from about 2.5 GW to 24 GW over the same time window. In addition, transfer capacities with neighbor countries have also increased with the commissioning of new DC interconnectors with Spain (2015), Great Britain (2021-22) and Italy (2022), as well as grid and operational upgrades in the AC transmission network.

Figure 3 shows the resulting monthly time series in terms of energy, along with domestic

[9] We do not include Flamanville 3 in our dataset because it had not started commercial operations as of 31 December 2024.

[10] Hydropower nonetheless shows a mildly increasing trend driven by the commissioning of small-scale run-of-the-river units. For fossil-fueled technologies, the closure of coal power plants has been compensated by the commissioning of gas-fired power plants.

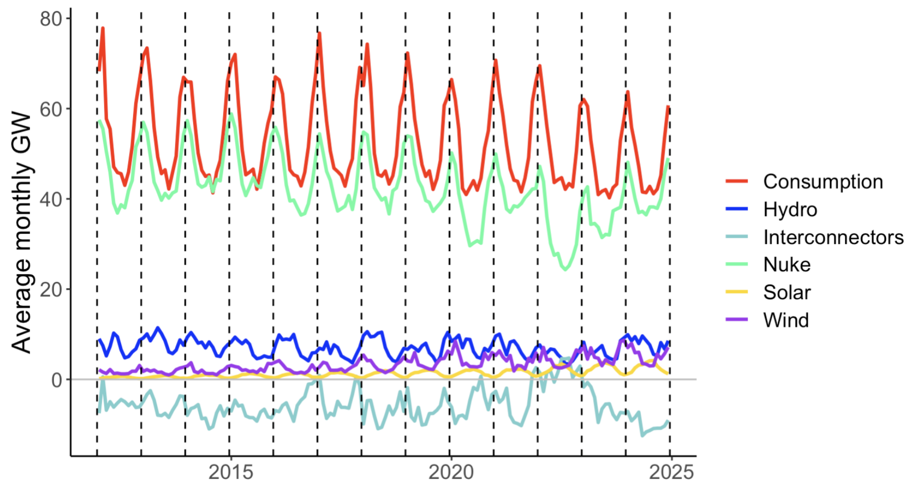

Figure 3: Monthly electricity consumption, generation by main technology and commercial exchanges over 2012–2024. Dashed vertical lines indicate the beginning of each year.

electricity consumption. Several observations are worth making. First, France is a so-called "winter peak" system, meaning that electricity consumption is highest during cold winter months (December, January and February). Indeed, given the high penetration of electric heating and the lack of widespread adoption of air conditioning, typical electricity demand during the winter is 1.5 to 2 times higher than during the summer. Second, nuclear is by far the dominant technology in the electricity mix, representing three quarter of domestic generation and an even larger fraction of domestic consumption. Third, as a result of the two previous observations, the fleet-level output of nuclear power plants is far from exhibiting a flat "baseload" pattern. Indeed, because of the very large share of nuclear in the electricity mix, nuclear load-following operations in France started decades before the emergence of intermittent renewables (Commission de régulation de l'énergie, 2025). Finally, the generation from wind and solar has been steadily growing from negligible amounts in 2012, to a combined output comparable to hydropower towards the end of the sample period. While this increase in average renewable output may seem relatively small, it masks considerable hourly variations. Indeed, the joint output from wind and solar can represent a large fraction of domestic generation in some hours. In addition, with the recent growth of wind and solar, the current French power system could, at least in theory, exclusively rely on nuclear and renewables.[11] Therefore, France is a particularly relevant case study to explore

[11] See Figure A.14 in the Appendix.

how power systems may behave as they approach an electricity mix relying exclusively on nuclear and renewables.

Institutional Background

As in the rest of Europe, wholesale electricity markets have been launched in France in the early 2000s. However, the incumbent utility, "Electricité de France" (EDF), still owns and operates the entire nuclear fleet, along with many hydro, thermal and renewable power plants. Electricity generation is thus very concentrated, with an HHI index exceeding 6,000 (Astier, 2025).

Strategic behavior and unilateral market power therefore represent legitimate concerns. Our research questions, however, relate to the short-term substitution between nuclear and intermittent renewables. There are a number of good reasons to believe that market power is unlikely to be of first-order importance in our setting. First, incentives to engage in (short-term) capacity withholding are mitigated by vertical integration (EDF remaining the dominant retailer) and the fact that the nuclear fleet was subject, for our period of interest, to a (partial) financial divesture mechanism.[12] Second, there is significant regulatory oversight of market players' bidding behaviors, so that bids significantly departing from (declared) opportunity costs would likely be detected.[13] Third, because renewables benefit from generous pay-as-produced support mechanisms, cross-ownership of nuclear and renewables seems unlikely to induce preventive curtailment of renewables in order to prevent market prices from collapsing. Consistently, over 2012–2024, renewable curtailment levels have been negligible or small: the annual technical reports of the system operator do not mention renewable curtailments until 2024. The report for the year 2024 (RTE, 2025) states that 2024 was the first year with non-negligible curtailed volumes, with 1.7 TWh of wind and solar output curtailed because of negative prices, that is, about 2.4% of total wind and solar production. This year indeed also coincides with a spike in non-positive hourly day-ahead prices (see Figure 4), which we will return to below.

In summary, while strategic responses to the growing penetration of intermittent renew-

[12] The mechanism was called the "Accès régulé à l'électricité nucléaire historique" (ARENH). See: `https://www.cre.fr/electricite/marche-de-gros-de-lelectricite/acces-regule-a-lelectricite-nucleaire-historique-arenh.html`.

[13] In addition, even if declared opportunity costs were to be mis-reported or strategically set, the relative ranking between renewables and nuclear in the merit order, which triggers load-following (see Figure 2), would very likely remain the same.

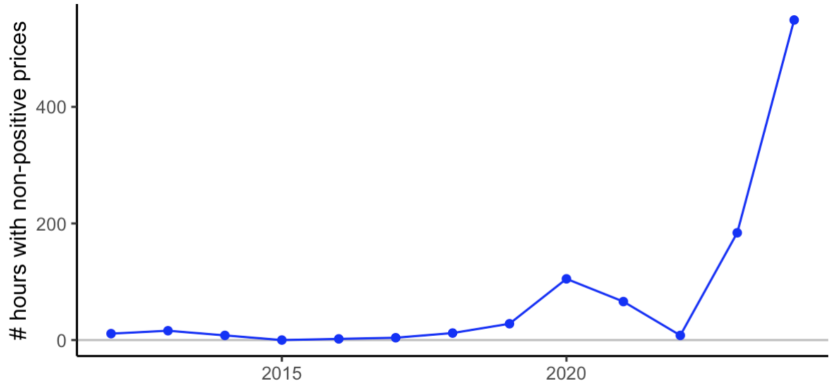

Figure 4: Number of non-positive hourly day-ahead prices per year in France.

ables represent an interesting an area of investigation, they may not play a critical role in our setting. We therefore leave their exploration for further research.

3.2 Data Sources

Our analyses exclusively rely on public data sources. The technical characteristics of each unit (nominal capacity, commissioning date, vintage, nominal minimum output level, etc.) are retrieved from EDF. Country-level information about the power system is publicly available at a (sub)hourly time scale from both the transmission system operator[14] and Entso-e.[15] We retrieve from these data sources hourly electricity consumption, aggregate generation by technology and net imports/exports for the period 2012–2024.

Unit-level hourly output of nuclear units is available from the transmission system operator (TSO) from 2012 onward.[16] Over the period 2012–2024, the hourly output levels of 56 (after 2020) to 58 units (before 2020) represent over 6 million observations. Episodes of missing or erroneous data are therefore a relevant concern. Because our later analysis predominantly relies on hourly unit-level output time series, we perform extensive data cleaning and sanity checks.[17]

[14] From 2012 onward at `https://www.rte-france.com/eco2mix`, last accessed on 9 April 2025.

[15] From 2015 onward at `https://transparency.entsoe.eu/`, last accessed on 9 April 2025.

[16] `https://www.services-rte.com/fr/visualisez-les-donnees-publiees-par-rte/production-realisee-par-groupe.html`, last accessed on 9 April 2025.

[17] First, we flag obvious outliers. Second, we cross-validate the data from the French TSO with the same data from Entso-e available from 2017 onward. Specifically, we rely on Entso-e data for days when we detect data quality issues in the French TSO data. Third, we aggregate unit-level output at the fleet level and compare it with the reported hourly aggregate nuclear output. Finally, we plotted and visually inspected all 58 time series, leading to a small number of manual corrections. Besides consistency with fleet-level output, these manual changes were informed, when relevant, by public information on unit-level outages.

Next, we leverage four sources of data on unit-level outages. First, the nuclear safety agency publishes on its website a notification whenever a nuclear unit is switched off and needs the authorization of the agency to be turned on again.[18] These outages usually correspond to planned outages for refueling and maintenance. We retrieve the dates of these outages, and whether or not refueling operations were performed.[19] Second, transparency regulations require the incumbent utility, which operates all nuclear units, to declare planned and forced outages. These notifications are published by both the TSO[20] and EDF.[21] The final dataset covers the period 2015–2024.[22] Since June 2021, the outage data published by EDF also include information about unit-level dynamic MOCs.[23] Finally, we obtained from the nuclear safety agency the date and unit of all the automatic shutdowns and manual emergency interventions that occurred in 2015–2024.[24]

Information on environmental outcomes at the plant-month level are manually retrieved for 2019–2024 from yearly environmental reports on nuclear installations (pursuant article 4.4.4. of the decree of 7 February 2012). Finally, hourly day-ahead prices are for example available from Entso-e's transparency platform.

4 Load-Following: Evidence and Measurement

In this Section, we first provide quantitative evidence of the widespread use of nuclear load-following in the French power system. We next propose discrete and continuous metrics that we will subsequently use to estimate associated impacts.

[18] `https://www.asn.fr/l-asn-controle/actualites-du-controle/installations-nucleaires/arret-de-reacteurs-de-centrales-nucleaires`, last accessed on 9 April 2025.

[19] We cross-validate and complement the information on refueling events by investigating manually, for each unit, all shut-down periods lasting more than a couple of weeks. A handful of (very likely) refueling events (based on typical fuel cycle lengths) could not be found on the nuclear safety agency website and have been added manually.

[20] `https://www.services-rte.com/fr/telechargez-les-donnees-publiees-par-rte.html?category=generation&type=unavailabilities`, last accessed on 9 April 2025.

[21] `https://www.edf.fr/groupe-edf/ambition-neutralite-co2-pour-edf-a-l-horizon-2050/optimisation-et-trading/listes-des-indisponibilites-et-des-messages/liste-des-indisponibilites`, last accessed on 9 April 2025.

[22] Besides an identifier of the outage, each notification consists of a publication date, a version (i.e., the number of updates about this outage that have been published to date), and information on the event (start, end, available capacity and, possibly, some minimal description). Although both data sources agree most of the time, their overlap is not perfect. We therefore build a merged dataset that keeps track of the latest notification about each event from either data source.

[23] Specifically, due to transparency obligations, the utility must disclose the actual MOC enforced at a given unit at a given point in time whenever its value is at least 100 MW larger than the nominal MOC. We leverage these disclosure messages to build unit-level hourly time series of the dynamic MOCs from June 2021 to December 2024.

[24] We are very grateful to the Autorité de Sûreté Nucléaire et de Radioprotection (ASNR) for sharing this dataset with us.

4.1 Aggregate Evidence of Load-Following

We use country-level hourly data for 2012–2024 to estimate several models of the form:

$$N_h = \alpha + \beta RES_h + \gamma RD_h + \sum_T \delta_{h\in T} + \epsilon_h \tag{1}$$

and:

$$N_h = \sum_{b\in\{\text{bins RD-RES}\}} (\alpha_b + \beta_b RES_h + \gamma_b RD_h) + \sum_T \delta_{h\in T} + \epsilon_h \tag{2}$$

where, for a given hour h, N_h denotes the output of the nuclear fleet, RES_h the combined output of wind and solar[25] and RD_h the residual demand except for wind and solar, that is, gross consumption (including pumped hydro) plus net exports minus hydropower and must-run biomass.[26] Specification (2) allows for non-linear effects by estimating different coefficients for different bins of system residual demand $RD_h - RES_h \in \{< 25, 25-40, 40-63, > 63\}$ (in GWh/h). We estimate different models that differ by the set of time fixed effects $\delta_{h\in T}$ included in the specification: none, week-of-sample or day-of-sample.

	Dependent variable: N_h (period 2012-2024)					
	(1)	(2)	(3)	(4)	(5)	(6)
RES_h	-0.78 (0.0064)	-0.62 (0.0065)	-0.66 (0.0075)			
RD_h	0.774 (0.0035)	0.626 (0.0057)	0.63 (0.0067)			
RES_h:bin<25GW				-0.3 (0.26)	-0.6 (0.28)	-0.8 (0.15)
RD_h:bin<25GW				0.4 (0.25)	0.6 (0.29)	0.7 (0.15)
RES_h:bin25-40GW				-1.04 (0.011)	-0.81 (0.011)	-0.85 (0.0098)
RD_h:bin25-40GW				1.09 (0.011)	0.82 (0.01)	0.85 (0.0095)
RES_h:bin40-63GW				-0.70 (0.0095)	-0.53 (0.0073)	-0.56 (0.0085)
RD_h:bin40-63GW				0.709 (0.0054)	0.56 (0.0067)	0.56 (0.008)
RES_h:bin>63GW				-0.43 (0.048)	-0.25 (0.022)	-0.24 (0.014)
RD_h:bin>63GW				0.44 (0.028)	0.29 (0.014)	0.25 (0.0091)
Week-of-sample FE	N	Y	N	N	Y	N
Day-of-sample FE	N	N	Y	N	N	Y
Observations	113,973	113,973	113,973	113,973	113,973	113,973
R2	0.919	0.984	0.994	0.930	0.988	0.995

Table 1: Obtained estimates for different specifications of Equations (1) and (2). Robust standard errors are clustered by day-of-sample.

[25] Because renewable curtailments have been negligible over the considered period (see above), using realized or forecasted output is unlikely to affect the obtained results.

[26] While hydro supply and imports/exports may raise endogeneity concerns, taking them as exogenous is a standard approximation in the literature (e.g. Borenstein et al. (2002)). The symmetry of the estimated coefficients for RES_h and RD_h suggests that this approximation is likely to be reasonable in our application as well. If anything, it may under-estimate the short-term substitution between nuclear and renewables. Indeed, by taking hydro output as exogenous (when computing hourly residual demand), we shut down the possibility that hydro reservoirs may endogenously reduce output during high renewable hours and substitute for nuclear output in a latter hour.

Table 1 reports the our estimation results. Depending on specifications, we find that an additional MWh of wind or solar generation is associated with a 0.6 to 0.8 MWh decrease in nuclear output. Because they do not include any time fixed effects, the estimates from specifications (1) and (4) capture the sum of both medium-term (scheduling of planned outages) and short-term (load-following) responses of the nuclear fleet to intermittent renewables. In contrast, by including week- or day-of-sample fixed effects, the remaining specifications control for the contemporaneous availability of the fleet (i.e., how many nuclear units are online), and therefore isolate load-following. Consistently, the estimate from specification (1) is larger (in absolute value) than estimates from specifications (2) and (3). Yet, the comparison of the different estimates suggests that the bulk of the response of the nuclear fleet to intermittent renewables is associated to load-following.[27] Non-linear specifications align with the "merit order" intuition of Figure 2: load-following predominantly occurs when residual demand is low relative to installed nuclear capacity. In addition, consistently with Figure 3, maintenance and refueling are scheduled in priority during periods with low residual demand (specification (4)).

Because most of the response of the nuclear fleet to the variability of wind and solar output consists of load-following,[28] we next discuss how to measure the intensity of load-following at the unit-level in order to subsequently investigate the potential opportunity costs associated with this operating mode.

4.2 Measuring Unit-level Load-Following

In order to empirically assess the consequences of nuclear load-following, it is necessary to construct metrics that capture its intensity. We rely on unit-level hourly output and outages to build such metrics.

Specifically, we define a "load-following event" as a set of contiguous hours during which a unit (i) is up and running (i.e., we exclude start-up, stretching and forced/planned outage events); and (ii) produces less than 75% of its installed capacity. While this latter threshold

[27] An heterogeneity analysis estimating (for specification (3)) separate substitution coefficients by year-of-sample suggests that its value remained fairly consistent across years, ranging from -0.6 to -0.76.

[28] Nuclear and renewables also interact at a sub-hourly time scales through ancillary services provision. Because of the small amplitude and short duration (seconds to minutes) of the corresponding variations in output, we cannot reliably measure this interaction from hourly data and therefore leave it to further research.

is somewhat arbitrary, it represents a sensible value to avoid false positives[29] while capturing load-following operations with a very high probability.[30] Counting the number of load-following events then provides a discrete metric of the intensity of load-following. Over 2015–2024, we detect 11,000+ load-following events, with a slightly decreasing trend over 2015–2022 and a sharp increase in 2023–2024, the latter year experiencing twice as many load-following events as the historical average (see Figure A.15).

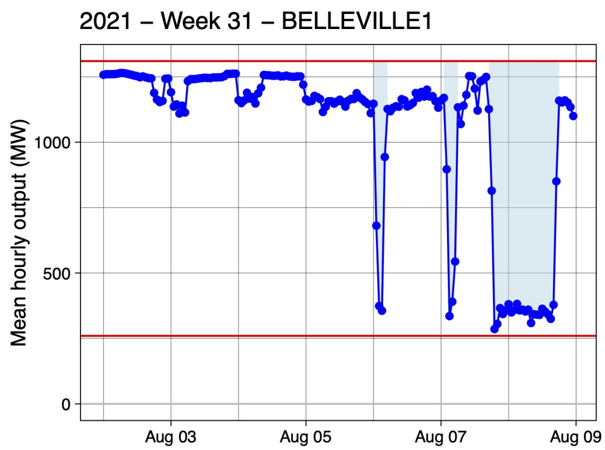

Figure 5: Illustration of "load-following events" and "lost energy".

Load-following events can be, however, very heterogeneous in terms of duration (median of 4 hours, with an inter-quartile range spanning from 3 to 7 hours) and depth of load reduction (median minimum output of 29% of installed capacity, with an inter-quartile range spanning from 24 to 51%). To account for this heterogeneity, we also introduce the concept of "lost energy" (in GWh), defined as the additional electricity that would have been generated if the unit had produced at its nominal capacity. Figure 5 shows the hourly output of the unit "Belleville 1" for one week in 2021. During that week, we detect that the unit performed three "load-following events", with a corresponding "lost energy" depicted as the shaded blue area. Normalizing lost energy (in GWh) by installed capacity (in GW) provides a measure of load-following that is commensurate to "full load hours" of lost energy, that is, the number of hours the unit would need to run at full capacity to produce the

[29] Output variations in the range 80-100% are frequent and exhibit very heterogeneous patterns, possibly due to the provision of ancillary services.

[30] Because load-following requires careful and costly monitoring by plant operators, a given amount of output reduction is usually achieved by significantly decreasing the output of handful of units rather than small reductions spread across all units.

energy output that was foregone due to load-following.

Of course, alternative measures of "load-following" can be defined. For example, Blanchard and Massol (2025) instead use the sum of absolute ramps. We found, however, the different measures of load-following to be highly correlated with each other, so that the results of subsequent analyses are very similar across the various possible metrics.

5 Opportunity Costs of Load-Following

This Section uses our previously defined "load-following" metrics to estimate the associated impacts in terms of loss of capacity factor, environmental externalities and safety.

5.1 Loss of Capacity Factor

Load-following has both direct and indirect impacts on the utilization rate a nuclear unit. This utilization rate, called the "capacity factor", is defined as the ratio between average and maximum output for a specified time period. As discussed in Section 2, the specific features of nuclear operations imply that the relevant timescale to measure average output is a fuel cycle, that is, the sequence consisting of a production period and a refueling outage.

The direct impact is a mechanical consequence of the fact that, during load-following events, the unit is producing below its installed capacity. Indirect impacts may also arise if load-following interacts with fuel efficiency, maintenance and/or outages.

Direct impact

Figure 6 illustrates our definition of the "direct" impact of load-following on the capacity factor of a nuclear unit. It represents a stylized fuel cycle, with the production period in blue and the refueling outage in orange. The shaded blue area materializes the lost-energy L due to load-following. If we denote with K the installed capacity of the unit, the number full-load hours of lost energy is, by definition, $H \equiv L/K$.

Let D denote the observed duration (in hours) of the whole fuel cycle (production and refueling). As a first approximation, the amount of electricity E generated during a cycle may be considered as a constant determined by the amount of fuel loaded during the previous refueling outage. Therefore, in the absence of load-following, the duration of the fuel cycle

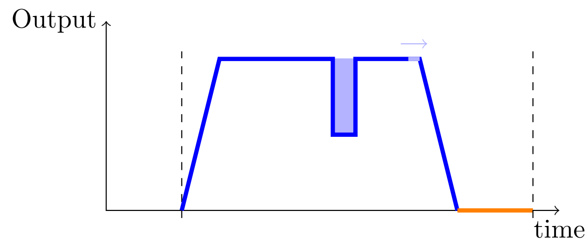

Figure 6: Direct impact of load-following on the average capacity factor of a unit.

would have been $D - H$. Using these notations, the (mechanical) decrease Δ in capacity factor due to load-following is equal to:

$$\Delta \equiv 100 \times \left(\frac{E}{K(D-H)} - \frac{E}{KD} \right) \simeq 100 \times \frac{E}{KD}\frac{H}{D} \tag{3}$$

Figure A.16 shows, across the fuel cycles that we fully observe in our sample, the relationship between the intensity of load-following and the decrease in capacity factor. Consistently with industry reports,[31] we find this decrease to be lower than 1 p.p. for the vast majority of fuel cycles, and to roughly amount to a 1 p.p. decrease in capacity factor for each 180 full-load hours of "lost energy". In addition, because load-following predominantly occurs during hours with low electricity prices, the per-MWh social value of "lost energy" is, by construction, lower than the average social value of electricity. As an illustration, we find the lost-energy-weighted average price of electricity to be 26€/MWh. In contrast, the nuclear-output-weighted price is 65€/MWh.

Indirect impacts

Load-following may also have an impact on the average capacity factor of the nuclear fleet through several indirect channels. First, fuel efficiency (i.e., the amount of electricity produced per ton of radioactive fuel) may increase or decrease. Second, the duration of refueling outages may increase if more maintenance is performed when the unit is shut down. Third, the occurrence of outages (planned or forced) may increase.

To explore these channels, we run regressions of the form:

$$Y_{if} = \alpha_i + \beta H_{if} + \epsilon_{if} \tag{4}$$

[31] For example, (Bruynooghe et al., 2010) mentions a 1.2 percentage point (p.p.) decrease in capacity factor.

where Y_{if} is the outcome of interest for unit i during fuel cycle f ((i) total electricity produced, (ii) duration of the refueling outage, and (iii) number of hours planned/forced outages, either in logs or in levels), α_i are unit fixed-effects and H_{if} denotes the number of full-load hours of lost energy (either in logs or in levels) for unit i during fuel cycle f.

Tables A.5 and A.6 in the Appendix report the obtained results. If anything, we estimate a positive relationship between load-following and fuel efficiency, although economically negligible. In addition, no significant relationship is found with the duration of subsequent refueling outages. Note that this null result does not necessarily imply that load-following does not induce additional maintenance, but rather that maintenance schedules seem to be set independently of the amount of load-following performed during a given fuel cycle. Finally, we estimate a positive and significant association between load-following and planned/forced outages. These coefficient estimates should, however, be interpreted with care. First, they are somewhat imprecise given how heterogeneous outages may be in terms of cause and duration,[32] not to mention the risk of strategic disclosure. Second, load-following may be less likely during fuel cycles that experienced outages early on, inducing a downward bias. Although one approach to mitigate this latter concern could be to instrument for H_{if},[33] we rather adopt in what follows a different empirical strategy (see below).

5.2 Environmental Impacts

Load-following may also have an impact on the environmental externalities associated with nuclear power. First, nuclear units use vast amounts of water for cooling. Load-following may thus have an impact on water intakes, water consumption,[34] or the temperature and pH of released cooling water. Second, nuclear units also release chemicals, radioactive (e.g. tritium) or not (e.g. boron), as part of the monitoring of the chain reaction.

We retrieved data from public environmental reports to assess whether we could detect an impact of load-following on these externalities. This empirical exercise, however, suffers from two important caveats. First, our dataset on environmental outcomes consists of month-plant

[32]For example, in 2023, strikes were the main cause of planned outages, inducing about 16 TWh of lost energy (Commission de régulation de l'énergie, 2025).

[33]Consistently with the above intuition, using the amount of load-following predicted by hourly system residual demand and day-of-sample fixed effects to instrument for H_{if} yields a non-significant coefficient estimate for planned outages and increases the magnitude of the coefficient estimate for forced outages.

[34]For air-cooled units, some water evaporates during cooling and therefore is not returned to the river.

observations for 2019–2024, that is, a pretty coarse panel. Second, liquid waste need not be released contemporaneously with nuclear production. Indeed, liquid chemicals can be stored in pools in order to be released only when circumstances meet environmental regulations.

With these shortcomings in mind, we are not able to detect any significant and/or economically meaningful association between load-following and environmental externalities (see Tables A.7 and A.8 in the Appendix).

5.3 Safety

The above evidence of an association between load-following and outages suffered from two main caveats. First, outages can be very heterogeneous events and, given asymmetric information, their disclosure may be endogenous (Bizet et al., 2022). Second, using fuel cycles as the unit of observation also raises endogeneity concerns, since a major outage at the beginning of the cycle may lead to subsequently more conservative operations.

To address the first concern, we follow Hausman (2014) and restrict attention to "emergency" outages, that is, automatic shutdowns and manual emergency interventions. We assume that such outages cannot be concealed to the nuclear safety agency, who provided us the date and units of these events.[35] Over the period 2015–2024, we observe 600+ emergency outages, about half of which occur as the unit is ramping up after a refueling outage, likely due to warming-up testing programs.

To address the second concern, we leverage unit-level hourly data and focus attention on the time window between the actual start of a fuel cycle[36] and either (i) the first emergency outage (if any) or (ii) the next refueling outage. Restricting attention to this time window indeed allows to credibly satisfy the underlying assumptions of survival models (Cox and Oakes, 1984). First, units that just completed a refueling outage, during which maintenance and repairs are performed, may be considered to be in statistically comparable states in terms of wear and tear. Second, there exists an unambiguous measure of time, which we will

[35]While we observe automatic shutdowns and manual emergency interventions separately, our main analysis treats both types of events as similar "emergency" outages. Indeed, when an incident occurs, it is sometimes the case that the unit operator will first try to perform manual emergency interventions to avoid an automatic shutdown. Our main results are robust, however, to restricting the sample to automatic shutdowns only.

[36]Given the numerous outages occurring during warming up testing programs, we define the "actual start of a fuel cycle" as the first time the unit reaches 85% of its installed capacity. While somewhat arbitrary, this threshold is chosen based on the start-up program of the most recent nuclear unit, Flamanville 3, whose last testing plateau was 80% of installed capacity (see Figure 18 of Commission de régulation de l'énergie (2025)).

denote by t. In what follows, time will be measured as “full-load hours” (cumulative output divided by installed capacity), but similar results are obtained using clock hours. Third, we perfectly observe both emergency events and fuel cycle ends. Finally, we observe our explanatory variable of interest (load-following) at every point in time for every unit.

Denote with T the date at which the first emergency event occurs for a given unit. Survival models treat T as a random variable with cdf $F(.)$ and pdf $f(.)$. However, instead of studying these latter functions, survival models focus attention on survival and hazard functions, which represent equivalent alternative ways of summarizing the distribution of T. The survival function $S(.)$ is defined as:

$$S(t) \equiv Pr(T \geq t) = 1 - F(t) \tag{5}$$

It represents the probability that the unit will “survive”, that is, will not experience an emergency event, at least until time t. The hazard function $h(.)$ is defined as:

$$h(t) \equiv \frac{f(t)}{S(t)} = \frac{f(t)}{1 - F(t)} \tag{6}$$

Loosely speaking, it captures the instantaneous rate of failure. More precisely, the probability of experiencing an emergency event between t and $t + dt$ (conditional on having survived until time t) is given by $h(t)dt$.

The Cox proportional hazards model assumes that the hazard function $h_i(.)$ of unit i may be parametrized as:

$$h_i(t, \theta) \equiv \exp(\theta^T x_i(t)) h_0(t) \tag{7}$$

where θ is a vector of parameters to be estimated, $x_i(t)$ is a vector of (possibly time-varying) explanatory variables, and $h_0(t)$ is an unspecified baseline hazard function. An appealing feature of this model is that θ can be consistently (although not efficiently) estimated by maximizing the following partial likelihood function (see Cox and Oakes (1984) p.117):

$$L_1(\theta) \equiv \Pi_{i=1}^{n} \frac{\exp(\theta^T x_i(\tau_i))}{\sum_{l \in R(\tau_i)} \exp(\theta^T x_l(\tau_i))} \tag{8}$$

where i indexes all observed dates τ_i of emergency events, and $R(\tau_i)$ denotes the “risk set” at τ_i, that is, all observations for which an emergency event either occurs at a later date or does not happen before the end of the fuel cycle (censoring). The underlying intuition for this partial maximum likelihood approach is that, since $h_0(t)$ is shared across all units, most of the identification power for θ comes from "which" unit experiences an event (conditionally on an event occurring) rather than "when" the event occurs. From this perspective, at each date τ_i, a survival model needs to be able to predict which individuals among the set $R(\tau_i)$ of remaining individuals will experience an emergency event. Equation (8) precisely corresponds to the likelihood function of a logit model applied to this problem.

	(1)	(2)	(3)	(4)
Full-load hours of lost energy	0.0018	0.0020		
	(0.00073)	(0.00078)		
# load following events			0.009	0.010
			(0.0033)	(0.0035)
vintage CP1		-0.3		-0.4
		(0.25)		(0.25)
vintage CP2		-0.5		-0.5
		(0.28)		(0.28)
vintage P4		0.0		0.0
		(0.287)		(0.288)
vintage P'4		-0.4		-0.4
		(0.28)		(0.28)
vintage N4		-0.6		-0.6
		(0.37)		(0.37)

Table 2: Estimated impact of cumulative load-following on the hazard rate.

Table 2 shows the estimation results under four different specifications,[37] which measure cumulative load-following either as full-load hours of lost energy (specifications (1) and (2)) or as the count of load-following events to date (specifications (3) and (4)). In addition, specifications (2) and (4) include dummy variables for the vintage of the units which are, from oldest to newest, CP0 (used as the reference), CP1, CP2, P4, P'4 and N4 (see Figure A.23).

Across all specifications, we find that one additional load-following event increases, on average, the hazard rates by about 1%. While this effect is fairly small for a single event, it is worth remembering that we detect 11,000+ load following events over 2015–2024, and that dozens of load-following events may be performed during a single fuel cycle.

To put this number into perspective using back-of-the-envelope calculations, note that, as a baseline, we observe at least one emergency event for about 60% of fuel cycles. Therefore,

[37] We estimate proportional hazards models using the R package *survival* (Therneau, 2024).

when the hazard rate increases by 0.002 as a result of 1 additional full-load hour of load-following, the increase in the probability to experience one additional emergency outage is (roughly) $0.002 \times 40\%$. Because the observed median (mean) number of full-load hours of lost energy during such events is 60 hours (165 hours), this indirect effect may amount to at least 5% to 15% of the direct effect in terms of capacity factor loss. This relative effect represents, however, a lower bound. First, it does not account for the possibility of repeated events. Second, in contrast to load-following, which occurs at low spot prices, forced outages are more random, and thus happen at higher average prices.

6 Limits to Load Following

While technically possible at a seemingly non-prohibitive opportunity cost, nuclear load-following is limited, however, by the operational constraints described in Section 2. This Section explores how wind and solar interact with such operating constraints.

6.1 Nominal and Dynamic Minimum Output Constraints

Inferring binding MOC constraints

Although we observe unit-level nominal MOCs (static values) and dynamic MOCs (time series for June 2021-2024), we do not directly observe whether the constraints are actually binding at any given point in time. Indeed, minimum output constraints are instantaneous *power* constraints. In contrast, our data consist of hourly *energy* output, that is, average power over the course of one hour. As a result, the MOC may bind within a given hour even though the average power during this hour is higher than this MOC.

In order to infer whether a given MOC was (likely) binding within a given hour, we proceed as follows. First, because we are interested in operating constraints that bind in the context of normal operations, we focus attention on non-outage hours. In addition, to avoid picking up hours when the unit enters the minimum output level region as it is ramping down/up before/after an outage, we exclude from the sample the set of contiguous hours before/after of an outage where output remains below 85% of nominal capacity.

Restricting attention to such hours (when a unit is experiencing "normal" operations), we

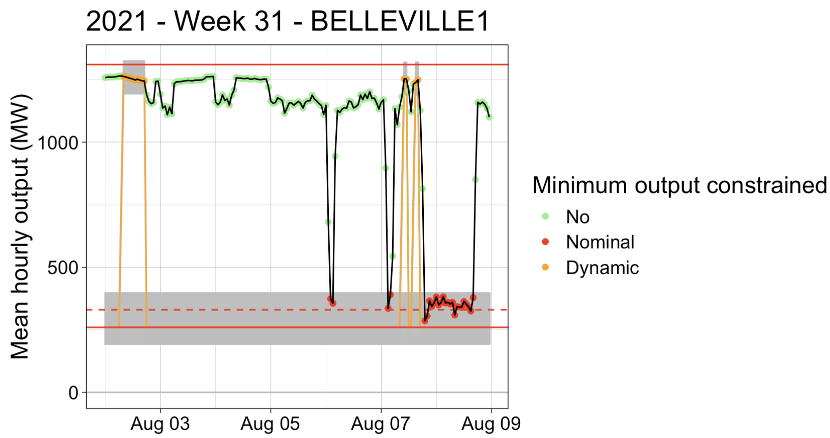

Figure 7: Illustration of the labeling (non-outage) of hours as "minimum output constrained", either due to the nominal MOC (red) or a dynamic MOC (orange).

Note: *The bottom red horizontal line represents the nominal MOC and the dashed red horizontal line this minimum output level plus the maximum capacity reserved for ancillary services. The nominal MOC is assumed to bind in a given hour if hourly output is within a buffer of 5% of nominal capacity around those two lines. Similarly, the dynamic MOC is assumed to bind if observed hourly output is within 5% of nominal capacity around this value.*

define the nominal MOC region as a buffer (with a width equal to 5% of the nominal capacity) around the nominal MOC and the nominal MOC plus the maximum capacity committed to ancillary services (2 to 5% of nominal capacity depending on the unit). Similarly, for dynamic MOCs, we consider that the minimum output constraint is binding within a given hour whenever the observed hourly output is within a buffer of 5% of nominal capacity around the other MOC in that hour. Figure 7 illustrates this imputation.

Occurrence of MOC constraints

In order to build intuition about whether and how intermittent renewables may interact with MOC constraints, we start with an exploration of the circumstances under which these constraints are observed to bind.

First, Figure 8 restricts the sample to unit-hour observations during which a unit is running (i.e., not outaged) and plots, for each month in 2015–2024, the fraction of observations for which each type of MOC is inferred to bind. We observe binding nominal MOCs to be relatively rare (about 1% of observations) and volatile events. In particular, the occurrence of binding nominal MOCs peaked during the Covid-19 lock down (spring 2020), a period when gross electricity consumption decreased sharply (Buechler et al., 2022). Similarly, the

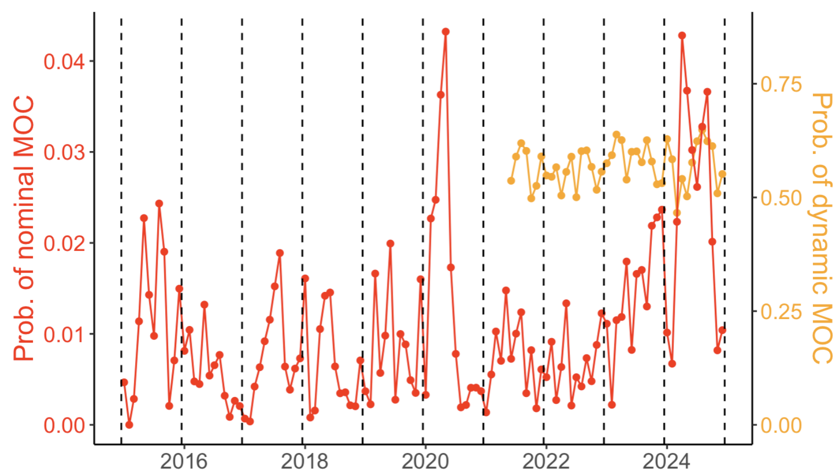

Figure 8: Fraction of hours in each month during which the nominal MOC (resp. dynamic MOC) was binding (conditionally on the unit being online and not outaged).

years 2022 to 2024, during which consumption was lower than usual (due to energy savings and mild winters) and the installed capacity of renewables increased significantly, exhibits a sharp increasing trend in the occurrence of nominal MOCs. In contrast, dynamic MOCs are found to bind very frequently (over 50% of the time) with a relatively flat time trend.

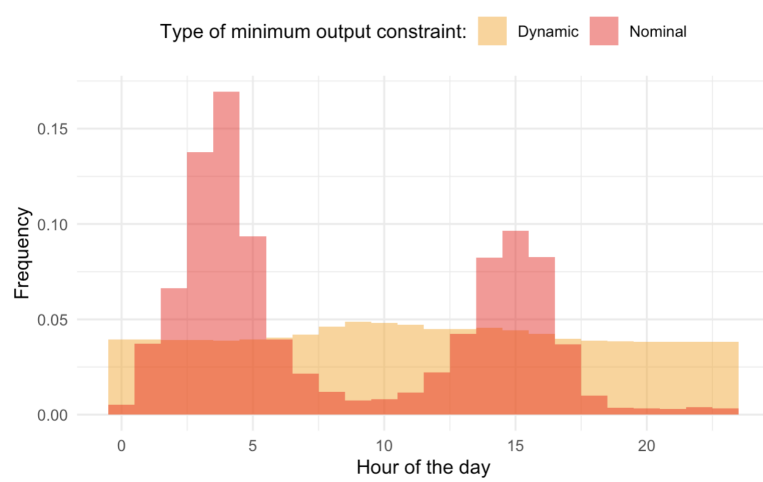

Figure 9: Histogram of the hours of the day when the minimum output constraints (either nominal or dynamic) are inferred to bind.

Next, Figure 9 shows the hours of the day during which either type of MOC most frequently binds. For nominal MOCs, the obtained distribution exhibits two modes: one during the night (2-6 am), when gross demand is low and wind generation can be high, and one during the afternoon (2-5 pm), when solar generation is high.[38] Figure A.18 further reveals that the relative magnitude of these modes has evolved over time, with a sharp increase in 2024, especially for the afternoon mode. In contrast, the occurrence of binding dynamic MOCs is distributed almost uniformly across hours of the day.

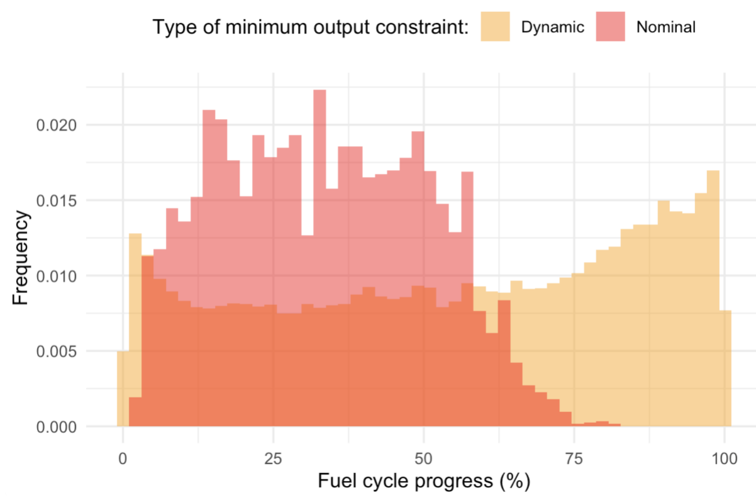

Figure 10: Occurrences of minimum output constraints (unit-hour observations) as a function of the position of the hour within the fuel cycle of the unit.

Finally, Figure 10 shows the observed frequency of nominal and dynamic MOCs as a function of the position of the hour within the fuel cycle of the unit. We observe that virtually all hours during which the nominal MOC binds take place during the first two thirds of the fuel cycle. Indeed, because of the physics and chemistry of nuclear reactors, French nuclear units can no longer operate very flexibly when they are reaching the end of a fuel cycle (OECD, 2021; Lynch et al., 2022). Consistently, dynamic MOCs bind relatively more frequently at the very beginning and during the last third of the fuel cycle, suggesting that they frequently arise because of fuel management constraints.

Several takeaways emerge from the previous descriptive graphs. First, despite being very flexible at the fleet level, nuclear units very frequently face minimum output constraints, especially dynamic MOCs. Second, the frequency of occurrence of these latter constraints does not exhibit a significant time trend, seasonality or intraday pattern. Instead, dynamic MOCs seem to relate, at least to a significant extent, to fuel management constraints. Third, the frequency of occurrence of nominal MOCs shows, in contrast, a strong intraday pattern,

[38] Figure A.17 shows the average aggregate generation of each technology for each hour of the day.

as well as monthly spikes, that suggest a strong sensitivity to residual demand, and thus wind and solar generation. In what follows, we quantitatively explore this latter relationship.

6.2 Renewables and Minimum Output Constraints

To assess the relationship between intermittent renewables and binding MOCs, we replace in Equation (1) the dependent variable with the number of units for which a minimum output constraint is inferred to bind in a given hour. Table 3 reports the obtained results. We also include one additional specification where the output from renewables is split into solar generation on the one hand, and wind generation on the other hand.

	Dependent variable: # of constrained units in hour h								
	Nominal MOC (2015-2024)			Dynamic MOC (Jun 2021-2024)			Either MOC (Jun 2021-2024)		
	(1)	(2)	(3)	(4)	(5)	(6)	(6)	(7)	(8)
RES_h (GWh)	0.081	0.30		-0.14	0.00		-0.03	0.34	
	(0.004)	(0.0078)		(0.023)	(0.0172)		(0.023)	(0.023)	
RD_h (GWh)	-0.026	-0.28	-0.28	0.44	0.35	0.34	0.40	0.02	0.01
	(0.0014)	(0.0071)	(0.0072)	(0.013)	(0.019)	(0.018)	(0.014)	(0.026)	(0.024)
$Solar_h$ (GWh)			0.31			0.06			0.41
			(0.0082)			(0.017)			(0.022)
$Wind_h$ (GWh)			0.27			-0.32			0.00
			(0.0076)			(0.029)			(0.034)
Day-of-sample FE	N	Y	Y	N	Y	Y	N	Y	Y
Num.Obs.	87,662	87,662	87,662	31,437	31,437	31,437	31,437	31,437	31,437
R2	0.108	0.656	0.658	0.358	0.838	0.845	0.321	0.822	0.830

Table 3: Relationship between wind and solar output and the occurrence of minimum output constraints. Robust standard errors are clustered by day-of-sample.

The obtained results are consistent with the previous stylized facts. First, when controlling for the level of residual demand (except for wind and solar) and day-of-sample fixed effects, a higher generation from wind and solar in a given hour is associated with more units reaching their nominal MOC during this hour. On average in a given hour, 1 additional nuclear unit is reaching its nominal MOC for every 3.3 GWh of additional wind and solar output. Estimating separate coefficients for wind and solar yields similar estimates. In contrast, the occurrence of dynamic MOCs is not associated with intraday variations in intermittent renewable generation. If anything, wind generation seems to be associated with a decrease in the occurrence of dynamic MOCs, perhaps when some steady minimum level of generation is needed locally for grid stability reasons.

6.3 Grid Location

We further explore the relationship between intermittent renewables and binding nominal MOCs locationally. While Table 3 reveals a strong relationship between the two variables, wind and solar power plants are not spread uniformly across space. In this paragraph, we thus explore whether spatial considerations may have a measurable impact on which nuclear units are ramped down when wind and solar generation is high.

Consider a given renewable power plant that produces 1 MWh of electricity. The "closer" this power plant is to a nuclear unit, the more substitutable this MWh may be to a MWh generated by the nuclear unit. Ideally, the proximity between the nuclear unit and the renewable facility should be measured by their "electrical distance", that is, by modeling power flows in a way that accounts for the topology of the power grid and Kirchhoff laws. This exercise, however, requires very detailed information about the power system at every point in time. Such data is unfortunately not publicly available.

Instead, we compute an "exposure" metric that may be derived from publicly available data (see Appendix B for more detail). Specifically, we rely on spatial rather than electrical proximity. Consider a given date t and a given nuclear unit n. Let $i \in I_{t,\tau}$ index all the power plants of technology $\tau \in \{wind, solar, hydro, thermal\}$ that are in service at date t, and denote with K_i the installed capacity of unit i. We denote with $d(i,n)$ the as-the-crow-flies distance between unit i and nuclear unit n. We then define the "exposure" $X_{n,t,\tau}$ of nuclear unit n to technology τ at date t as:

$$X_{n,t,\tau} \equiv \sum_{i \in I_{t,\tau}} (\frac{\max(d(i,n), d_0)}{d_0})^{-\gamma} K_i \tag{9}$$

where d_0 and γ are (positive) tuning parameters. The parameter d_0 plays two roles. First, it defines a buffer around the nuclear unit, to avoid giving an infinite weight to units located in the same municipality. In practice, nuclear power plants tend to be isolated (the 1% quantile of $d(i,n)$ is 30 km or more for all installations), so that the exact size of this buffer has little influence on the results. Second, it defines the unit for distances. However, because we later rely on relative exposures (see Equation (10)), this dimension plays no role in our results. In what follows, we therefore simply set $d_0 = 1$ km. The parameter $\gamma \in (0,1)$

controls the speed of decay. Specifically, if an installation is located twice further away from a given nuclear plant than another installation, its weight will be lower by a factor $2^{-\gamma}$. As a illustration, the obtained relative weights for $\gamma \in \{0.25, 0.5, 0.75\}$ are $\{0.84, 0.71, 0.59\}$. We use $\gamma = 0.5$ in what follows, but similar results are obtained for $\gamma = 0.25$ and $\gamma = 0.75$ (see below). Figure B.26 shows the obtained "exposure" metric (in MW) to wind and solar for each nuclear plant. We observe a significant amount of both temporal and cross-sectional variation. Quite intuitively, nuclear plants located in the Southern part of the country are most exposed to solar, those in the Northern part most exposed to wind and those in the center exposed to both wind and solar.[39]

For wind and solar, we then compute the "share" $W_{\tau,h,n}$ (in GWh) of the aggregate national output $S_{\tau,h}$ of technology τ in hour h that is "assigned" to unit n as follows. First, we retrieve, for hour h, the set M of units that are not experiencing an outage. Second, we use the vector of exposures as a sharing key for $S_{\tau,h}$ and define:

$$W_{\tau,h,n} \equiv \frac{X_{\tau,h,n}}{\sum_{m \in M} X_{\tau,h,m}} S_{\tau,h} \tag{10}$$

Our empirical strategy is then to estimate, on the sample of non-outaged unit-hour observations, linear probability models of the form:[40]

$$Y_{h,n} = \alpha + \beta_1^{solar} Solar_h + \beta_1^{wind} Wind_h + \beta_2^{solar} W_{solar,h,n} + \beta_2^{wind} W_{wind,h,n} + \eta RD_h + \delta_{h \in D} + \lambda_n + \epsilon_{h,n} \tag{11}$$

where $Y_{h,n}$ is a dummy variable taking the value 1 if the nominal minimum output constraint is assessed to be binding in hour h for unit n, λ_n denotes unit fixed effects, and $\delta_{h \in D}$ refers to day-of-sample fixed effects.

The rationale behind Equation 11 is to nest the specification whose results are reported in Table 3. Specifically, if we remove the terms $W_{\tau,h,n}$ from Equation 11, we obtain the same specification as before, except that it is estimated on an unbalanced panel of unit-hour observations, rather than on a time series that aggregates the cross-sectional dimension.

[39] Note that, for this spatial analysis, "wind" only refers to on-shore installations. Indeed, as of 2025, France only hosts 1.5 GW of off-shore wind (vs. 22.5 GW of on shore wind), and the corresponding plants were commissioned during the very end of our sample period (the above country-level analysis does, however, include off-shore wind output in aggregate wind generation). Adding off-shore wind as an additional technology does not alter the obtained results.

[40] We also estimated logit models and obtained marginal effects of similar magnitudes. We therefore report the results of linear probability models to simplify the interpretation of results.

Adding the terms $W_{\tau,h,n}$ then asks the question: controlling for aggregate residual demand and renewable output, does the "exposure" of a nuclear unit to a given renewable technology nonetheless still positively correlates with the occurrence of binding nominal MOCs?

Dependent Variable:	Is nominal MOC binding? ($\gamma = 0.5$)				
	(1)	(2)	(3)	(4)	(5)
RES_h (GWh)	0.0078	0.0077			
	(6.4e-05)	(0.00012)			
RD_h (GWh)	-0.00707	-0.00707	-0.00716	-0.00715	-0.00730
	(5.5e-05)	(5.5e-05)	(5.6e-05)	(5.6e-05)	(5.7e-05)
$W_{RES,h,n}$ (GWh)		0.004			
		(0.0038)			
Solar_h (GWh)			0.0082	0.0080	0.0081
			(0.00007)	(0.00015)	(0.00015)
Wind_h (GWh)			0.0069	0.0071	0.0069
			(6.8e-05)	(0.00017)	(0.00017)
$W_{solar,h,n}$ (GWh)				0.008	0.011
				(0.0048)	(0.0049)
$W_{wind,h,n}$ (GWh)				-0.006	0.005
				(0.0058)	(0.0059)
$X_{hydro,h,n}$					-0.42
					(0.036)
$X_{therm,h,n}$					-1.0
					(0.09)
Num.Obs.	3,370,555	3,370,555	3,370,555	3,370,555	3,370,555
R2	0.064	0.064	0.064	0.064	0.064

Table 4: Obtained results when estimating Equation (11). Robust standard errors (HC1) clustered by unit and day-of-sample are reported in parenthesis. Exposure metrics are computed using $\gamma = 0.5$.

Importantly, because residual demand includes net exports, this approach allows to control (at least in parts) for any "exposure" to foreign generation units (whose location is not observed). Table 4 reports the obtained estimation results. The first and third columns correspond to the counterpart of the results shown in Table 3: one additional GWh of renewable output increases the probability that a unit available to ramp down reaches its nominal MOC by 0.007 to 0.008. Because there are on average 38 such units in a given hour, this estimate closely matches the previous estimate from aggregated data ($0.008 \times 38 = 0.3$), a mechanical result. Adding the combined exposure to renewable output ($W_{RES,h,n} \equiv W_{solar,h,n} + W_{wind,h,n}$) as a co-variate (specification (2)) yields a statistically insignificant coefficient. However, when exposures to wind and solar enter separately, a statistically significant relationship between exposure to solar generation and binding nominal MOCs is found. In contrast, no robust effect is estimated for wind, and exposure to hydro and thermal technologies (measured with $X_{\tau,h,n}$ for simplicity) is, quite intuitively, negatively associated to binding nominal MOCs. Tables A.9 and A.10 in the Appendix report the obtained results for alternative values of the parameter γ (0.25 and 0.75 respectively). Consistently with the hypothesis that solar generation may induce local grid constraints, a larger coefficient for $W_{solar,h,n}$ is estimated when closer solar units are given a higher weight ($\gamma = 0.75$) than when they are given a lower weight ($\gamma = 0.25$).

Overall, our analysis supports the hypothesis that the increasing penetration of solar generation may induce grid constraints that further limit the ability of the nuclear fleet to accommodate renewable generation. One possible explanation for the absence of a similar result for wind is that, in contrast to solar generation, which exhibits a very high degree of contemporaneous correlation (Wolak, 2016), wind generation is more diversified across locations (see Figure A.17). Therefore, our exposure metric $W_{wind,h,n}$ for wind may be less precise (i.e., less correlated with actual power flows) than our exposure metric for solar.

6.4 Remaining Flexibility and Spot Prices

In 2024, France experienced a sharp increase in the occurrence of non-positive prices (Figure 4), a trend that is persisting in 2025. As discussed in Section 2, an exhausted ability to perform nuclear load-following combined with low residual demand levels may rationalize such episodes. In this paragraph, we empirically explore this possibility.

More precisely, we note that, at any given point in time, a given unit must be in one of three possible states. First, the unit may be in an "outage" state (planned/forced outages, refueling). In such situations, it cannot provide any downward flexibility. Second, the unit may be producing at its minimum output level, either nominal or dynamic. Similarly, such units cannot decrease their output further. Third, the unit may be running above its minimum output constraint. We denote with $\mathcal{F}_h$ (for "flexible") the set of units that fall within this latter category in hour h.

We then denote with RF_h the remaining (downward) flexibility of the nuclear fleet in hour h, defined as:

$$RF_h \equiv \sum_{n \in \mathcal{F}_h} (P_{n,h} - MOC_{n,h}) \tag{12}$$

where $P_{n,h}$ is the observed output of unit n during hour h, and $MOC_{n,h}$ is the minimum output constraint of this unit for this hour. Figure A.19 shows the evolution of the (monthly) distribution of RF_h between June 2021 and December 2024. Although its median value does not exhibit any obvious pattern, the lower tail of the distribution (first percentile and below) has been significantly shifting to the left.

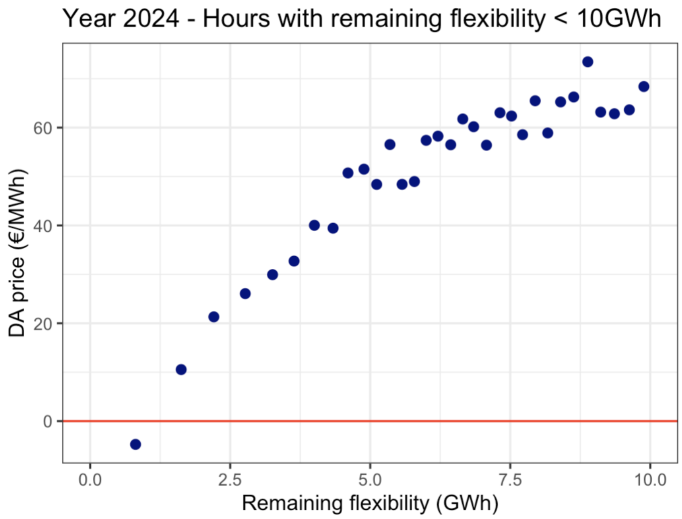

Figure 11: Binscatter plot of hourly day-ahead prices against remaining nuclear flexibility (drawn with the R package binsreg by Cattaneo et al. (2024)). The sample is restricted to hours in 2024 with a remaining flexibility lower than 10 GWh.

Figure 11 shows a bin-scatter plot, for the year 2024, of the remaining flexibility RF_h against hourly day-ahead prices, censoring the sample to hours with 10 GWh or less of remaining flexibility. We find that hours with non-positive prices typically correspond to situations where the short-run flexibility of the nuclear fleet is exhausted.

Overall, this observation suggests that, since 2024, the French power system has encountered novel operating conditions where nuclear load-following is reaching its limits for a non-negligible fraction of hours. This situation results in more frequent episodes of non-positive prices, as well as a growing reliance on renewable curtailments (RTE, 2025). Should this trend persist, it may end up compromising the profitability of both nuclear and intermittent renewables, and thus their ability to jointly provide reliable low-carbon electricity supply.

7 External Validity in Space and Time

Arguably, France may represent an optimistic case study for studying the ability of nuclear power to perform load-following. Indeed, the French nuclear fleet is operated by a single operator who has, given the very high share of nuclear in the electricity mix, decades of experience with load-following operations (Commission de régulation de l'énergie, 2025). In particular, the design of reactors was adjusted early on to improve their load-following abilities (Bruynooghe et al., 2010). However, all modern reactors supposedly now share this ability. For example, prior to the nuclear phase out, load-following was also performed in Germany (Lokhov, 2011; Grünwald and Caviezel, 2017) and the European Utility Requirements specify that a "*nuclear unit should be able to go through the following number of load variations: 2 per day, 5 per week and 200 per year.*"

Figures A.20–A.22 in the Appendix explore how often these thresholds are reached in practice. At a daily time scale, at most 4% of unit-day observations in a month exhibit 2+ load-following events. At the weekly time scale, although the number of unit-week observations with 5+ load-following events has increased sharply in recent years, they still represent only about 4% of unit-week observations in 2024. Finally, none of the unit is observed to perform more than 200 load-following events per year.[41]

Overall, the external validity of our findings to other jurisdictions rather hinges on regulatory and institutional feasibility. For example, the United States has explicit regulations restricting the use of nuclear load-following (Lokhov, 2011). Given the small but positive association between load-following and emergency forced outages, concerns about nuclear safety seem likely to affect the political feasibility of lifting up such regulations.

Looking forward, the evidence that the nuclear fleet has exhausted its load-following potential for a non-negligible number of hours in 2024 suggests that, while nuclear has the ability to operate quite flexibly, additional flexible assets are likely to be required once intermittent renewables reach very high penetration levels. In a first-best environment, the increasing occurrence of non-positive prices should attract investments in such assets (pumped hydro, batteries, demand response, etc.). Whether this prediction will materialize in

[41] Besides load-following, "load variations" may also refer to the provision of ancillary services. However, such provision is increasingly coming from other flexible assets such as batteries.

practice is an open question given the severe market failures inherent to European electricity markets (Graf, 2025). For example, the installed capacity of battery storage, while growing, remains relatively small in France (1.1 GW as of 31 December 2024). Similarly, French retail rates remain heavily regulated (Astier, 2025) and generally fail to pass-through episodes of non-positive prices to end-consumers.

8 Conclusion

This paper studies empirically the behavior of a low-carbon electricity system relying on high amounts of both renewables and nuclear power. We focus on the case of France, a country whose generation fleet includes, as of 31 December 2024, 63 GW of nuclear capacity and about 50 GW of cumulated wind and solar capacity.

We first explore how the French nuclear fleet adjusts its operations when wind and solar generation is high. We find that, in contrast to widely held beliefs, nuclear units are in practice operated very flexibly: an additional 1 MWh of domestic wind and solar generation is on average associated to a 0.6 MWh decrease in nuclear output. This substitution is made possible by so-called "load-following" operations, that is, short-duration decreases in the output of a handful of units. We next seek to assess the potential economic, environmental and safety costs associated with load-following. Although the costs we are able to estimate do not stand out as prohibitive, they are nonetheless non-negligible. In particular, load-following is found to slightly but significantly increase the hazard rate of automatic shutdowns and manual emergency interventions.

Finally, we study the main operational constraints limiting the technical feasibility of load-following, namely minimum output constraints (MOCs). Despite the significant flexibility of nuclear operations, we find that MOCs, especially dynamic MOCs raised by fuel management constraints and, possibly, grid congestion, are frequently binding, limiting load-following potential. In addition, units that are available for load-following are increasingly constrained by their nominal MOC as wind and solar generation increases.

Overall, the interaction between high renewable generation and non-convexities in the supply function of nuclear units (MOCs) can ultimately result in non-positive spot prices and renewable curtailments, exacerbating the substitution between nuclear and renewables.

Consistently, in 2024, the large increase in the occurrence of non-positive prices is found to be strongly associated with situations where the remaining downward short-run flexibility of the nuclear fleet was exhausted.

References

Adler, David B, Akshaya Jha, and Edson Severnini, "Considering the nuclear option: Hidden benefits and social costs of nuclear power in the US since 1970," *Resource and Energy Economics*, 2020, *59*, 101127.

Andersson, Julius J and Jared J Finnegan, "Industrial Policy and Decarbonization: The Case of Nuclear Energy in France," *Mimeo*, 2024.

Astier, Nicolas, "French market design in practice: Some lessons from the 2022 energy crisis," *The Electricity Journal*, 2025, *38* (3), 107483.

_ , Ram Rajagopal, and Frank A Wolak, "Can distributed intermittent renewable generation reduce future grid investments? Evidence from france," *Journal of the European Economic Association*, 2023, *21* (1), 367–412.

Bizet, Romain, Petyo Bonev, and François Lévêque, "The effect of local monitoring on nuclear safety and compliance: Evidence from France," *Journal of Environmental Economics and Management*, 2022, *113*, 102651.

Blanchard, Ange and Olivier Massol, "The value of nuclear power plants' flexibility: A multistage stochastic dynamic programming approach," *European Journal of Operational Research*, 2025.

Boccard, Nicolas, "The cost of nuclear electricity: France after Fukushima," *Energy Policy*, 2014, *66*, 450–461.

Borenstein, Severin, James B Bushnell, and Frank A Wolak, "Measuring market inefficiencies in California's restructured wholesale electricity market," *American Economic Review*, 2002, *92* (5), 1376–1405.

Bruynooghe, Christiane, Arne Eriksson, and Gianluca Fulli, *Load-following operating mode at Nuclear Power Plants (NPPs) and incidence on Operation and Maintenance (O&M) costs-Compatibility with wind power variability*, Joint Research Center, 2010.

Buechler, Elizabeth, Siobhan Powell, Tao Sun, Nicolas Astier, Chad Zanocco, Jose Bolorinos, June Flora, Hilary Boudet, and Ram Rajagopal, "Global changes in electricity consumption during COVID-19," *Iscience*, 2022, *25* (1).

Cattaneo, Matias D, Richard K Crump, Max H Farrell, and Yingjie Feng, "On binscatter," *American Economic Review*, 2024, *114* (5), 1488–1514.

Commission de régulation de l'énergie, "Evaluation des coûts complets de production de l'électricité au moyen des centrales électronucléaires historiques pour la période 2026-2028," Technical Report 2025.

Cox, David Roxbee and D Oakes, *Analysis of survival data*, Chapman and Hall/CRC, 1984.

Davis, Lucas, "Prospects for nuclear power," *Journal of Economic perspectives*, 2012, *26* (1), 49–66.

— **and Catherine Hausman**, "Market impacts of a nuclear power plant closure," *American Economic Journal: Applied Economics*, 2016, *8* (2), 92–122.

— **and Catherine Wolfram**, "Deregulation, consolidation, and efficiency: Evidence from US nuclear power," *American Economic Journal: Applied Economics*, 2012, *4* (4), 194–225.

Graf, Christoph, "Simplified short-term electricity market designs: Evidence from Europe," *The Electricity Journal*, 2025, p. 107486.

Grubler, Arnulf, "The costs of the French nuclear scale-up: A case of negative learning by doing," *Energy Policy*, 2010, *38* (9), 5174–5188.

Grünwald, Reinhard and Claudio Caviezel, "Load-following capability of German nuclear power plants," *Büro für Technikfolgen-Abschätzung beim Deutschen Bundestag (TAB), Berlin*, 2017.

Hausman, Catherine, "Corporate incentives and nuclear safety," *American Economic Journal: Economic Policy*, 2014, *6* (3), 178–206.

Jarvis, Stephen, Olivier Deschenes, and Akshaya Jha, "The private and external costs of Germany's nuclear phase-out," *Journal of the European Economic Association*, 2022, *20* (3), 1311–1346.

Jenkins, Jesse D, Zhi Zhou, Roberto Ponciroli, Richard B Vilim, Francesco Ganda, Fernando de Sisternes, and Audun Botterud, "The benefits of nuclear flexibility in power system operations with renewable energy," *Applied energy*, 2018, *222*, 872–884.

Jha, Akshaya and Gordon Leslie, "Start-up costs and market power: Lessons from the renewable energy transition," *American Economic Review*, 2025, *115* (2), 690–726.

Johannsen, Alida, Chloé Le Coq, and Sebastian Schwenen, "Nuclear and Renewables - Friends or Foes?," *Mimeo*, 2025.

Joskow, Paul L, "Comparing the costs of intermittent and dispatchable electricity generating technologies," *American Economic Review*, 2011, *101* (3), 238–241.

— **and John E Parsons**, "The future of nuclear power after Fukushima," *Economics of Energy & Environmental Policy*, 2012, *1* (2), 99–114.

Liski, Matti and Iivo Vehviläinen, "Ownership and collusive exit: Theory and a case of nuclear phase-out," *WorkingPaper, MIT Center for Energy and Environmental Policy Research*, 2018, pp. 212–230.

Loisel, Rodica, Victoria Alexeeva, Andreas Zucker, and David Shropshire, "Load-following with nuclear power: Market effects and welfare implications," *Progress in Nuclear Energy*, 2018, *109*, 280–292.

Lokhov, Alexy, "Load-following with nuclear power plants," *NEA news*, 2011, *29* (2), 18–20.

Lundin, Erik, "Market power and joint ownership: Evidence from nuclear plants in Sweden," *The Journal of Industrial Economics*, 2021, *69* (3), 485–536.

Lynch, Arthur, Yannick Perez, Sophie Gabriel, and Gilles Mathonniere, "Nuclear fleet flexibility: Modeling and impacts on power systems with renewable energy," *Applied Energy*, 2022, *314*, 118903.

Makarin, Alexey, Nancy Qian, and Shaoda Wang, "The Political Economic Determinants of Nuclear Power: Evidence from Chernobyl," *Mimeo*, 2024.

Mallapragada, Dharik S, Cristian Junge, Cathy Xun Wang, Johannes Pfeifenberger, Paul L Joskow, and Richard Schmalensee, "Electricity price distributions in future renewables-dominant power grids and policy implications," Technical Report, National Bureau of Economic Research 2021.

Mansur, Erin T, "Measuring welfare in restructured electricity markets," *The Review of Economics and Statistics*, 2008, *90* (2), 369–386.

Millard, Rachel, "France 'flexes' nuclear output as solar reshapes European energy," *Financial Times*, 2025.

Morilhat, Patrick, Stéphane Feutry, Christelle Le Maitre, and Jean Melaine Favennec, "Nuclear power plant flexibility at EDF," *Mimeo*, 2019.

OECD, *Technical and economic aspects of load following with nuclear power plants*, OECD Publishing, 2021.

Rangel, Lina Escobar and Francçois Lévêque, "Revisiting the cost escalation curse of nuclear power: New lessons from the French experience," *Economics of Energy & Environmental Policy*, 2015, *4* (2), 103–126.

Reguant, Mar, "Complementary bidding mechanisms and startup costs in electricity markets," *The Review of Economic Studies*, 2014, *81* (4), 1708–1742.

RTE, "Bilan électrique 2024," Technical Report 2025.

Severnini, Edson, "Impacts of nuclear plant shutdown on coal-fired power generation and infant health in the Tennessee Valley in the 1980s," *Nature Energy*, 2017, *2* (4), 1–9.

Therneau, Terry M, *A Package for Survival Analysis in R* 2024. R package version 3.8-3.

Wolak, Frank A, "Level versus variability trade-offs in wind and solar generation investments: The case of California," *The Energy Journal*, 2016, *37* (2_suppl), 185–220.

Appendices

A Supplementary Figures and Tables

A.1 Additional Figures

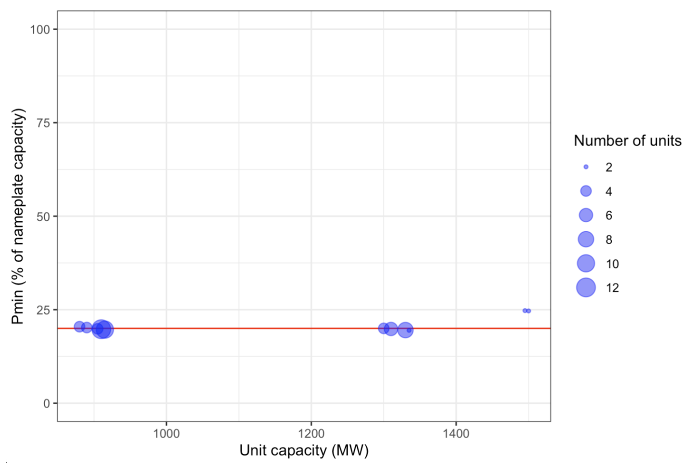

Figure A.12: Minimum output level (expressed as a percentage of nominal capacity) of French nuclear units (source: EDF). The red horizontal line corresponds to 20% of nominal capacity.

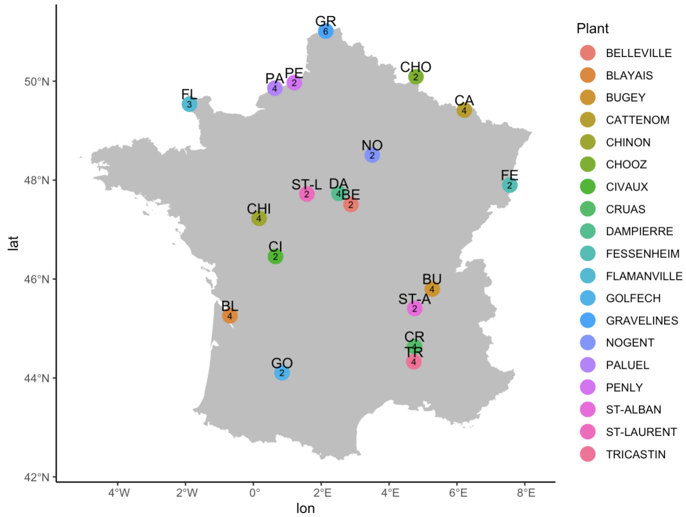

Figure A.13: Location of the 19 nuclear power plants. The numbers in the circles indicate the number of units that compose each power plant. The two units of Fessenheim (in the North East) have been permanently shut down in 2020 and are no longer in operations. The third unit of Flamanville was commissioned in 2024 has not started commercial operations in that year. It is therefore not included in our dataset.

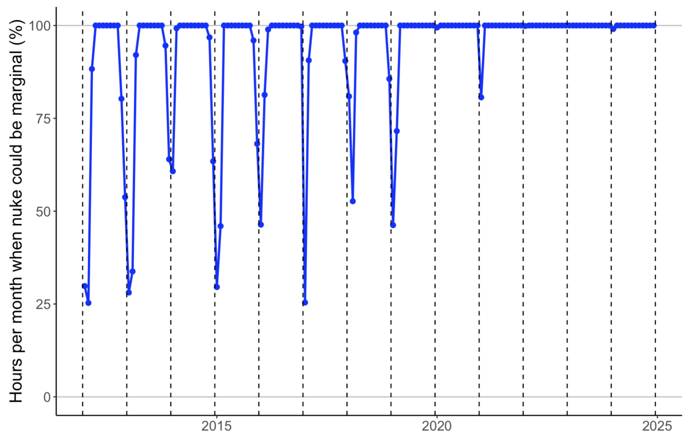

Figure A.14: Percentage of hours in each month when nuclear would be marginal in an hypothetical situation where all the nuclear units would be available all the time.

Note: *to draw this Figure, we compute, for each hour, the "residual demand" defined as gross electricity consumption (including water pumping) plus (net) exports minus renewable generation (wind, solar, hydro and biomass). In other words, the residual demand represents the remaining electricity load that must be supplied by "dispatchable" assets, such as nuclear or fossil-fueled power plants. For each hour, we then compare the level of residual demand to installed nuclear capacity. If the latter exceeds the former, the nuclear fleet is sufficient to balance supply and demand under the hypothetical scenarios where all nuclear units would be available. The Figure plots, for each month in our sample, the percentage of hours where such a "theoretical marginality" of nuclear occurs. We observe that, over the past three years, the increase in wind and solar (combined with the decrease in gross consumption and changes in imports/exports patterns) has been sufficient to make nuclear "theoretically marginal" in all hours.*

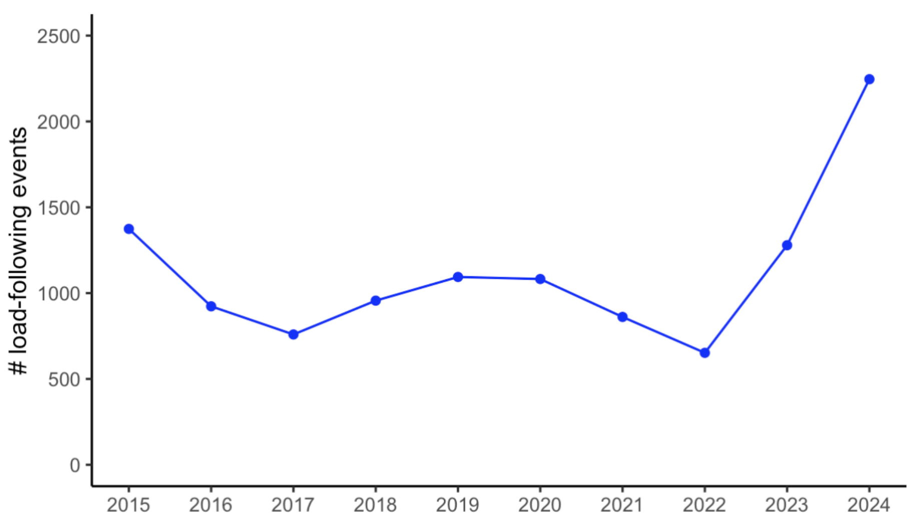

Figure A.15: Number of detected load-following events per year (2015-2024).

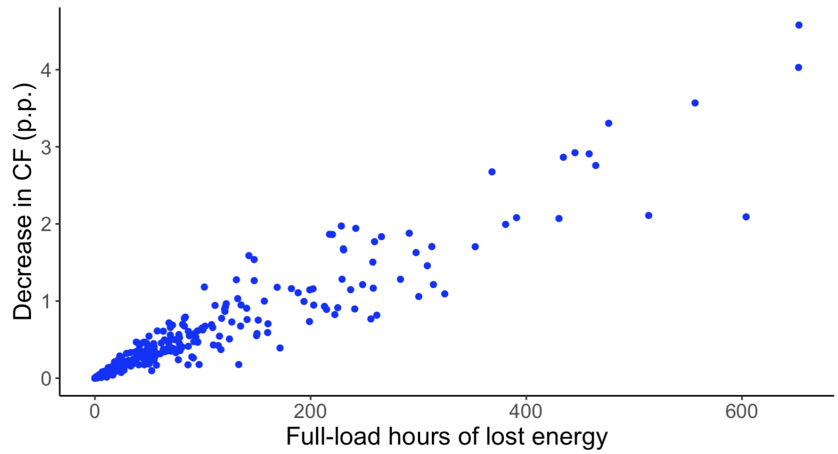

Figure A.16: Computed (mechanical) decrease in the capacity factor of nuclear units ("direct impact") as a function of the number of full-load hours of lost energy.

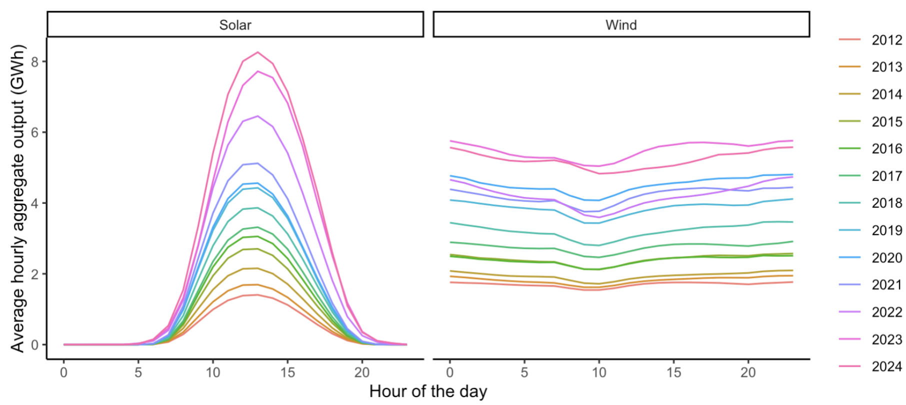

Figure A.17: Average hourly aggregate generation (GWh) from wind and solar for each year in our sample period.

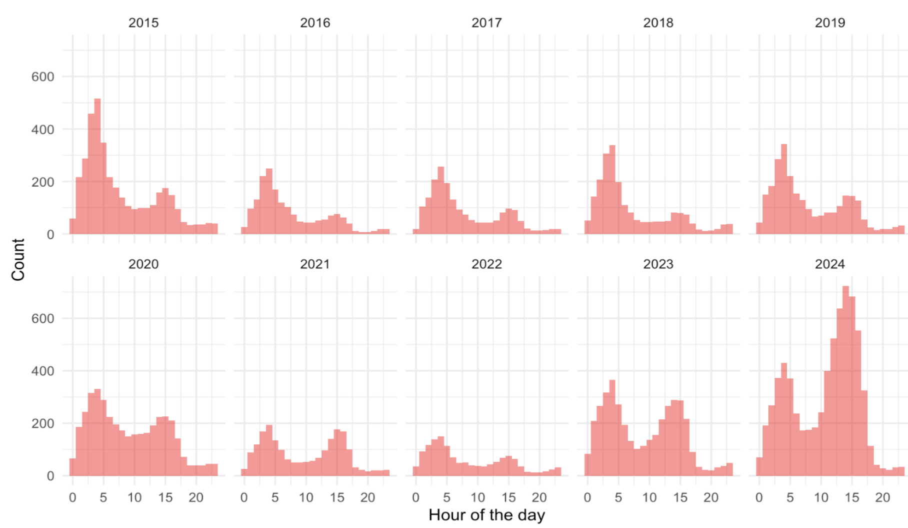

Figure A.18: Histograms (one for each year in the sample) of the hours of the day when the nominal MOC is inferred to bind.

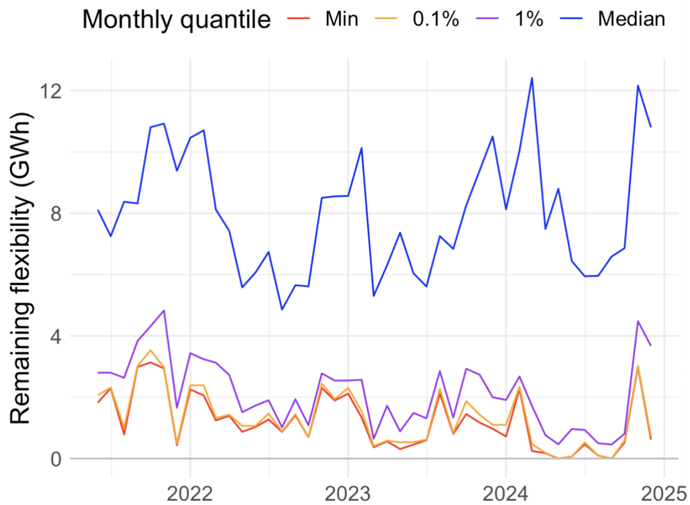

Figure A.19: Summary statistics of the monthly distributions of hourly remaining flexibility: minimum, quantile 0.001, first percentile and median.

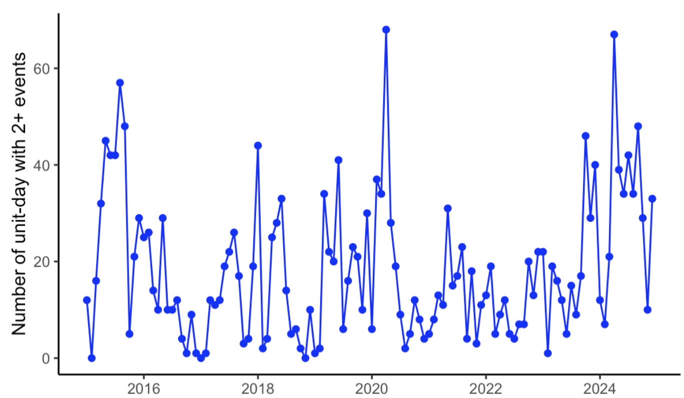

Figure A.20: Number of unit-day observations in each month with 2 or more load-following events.

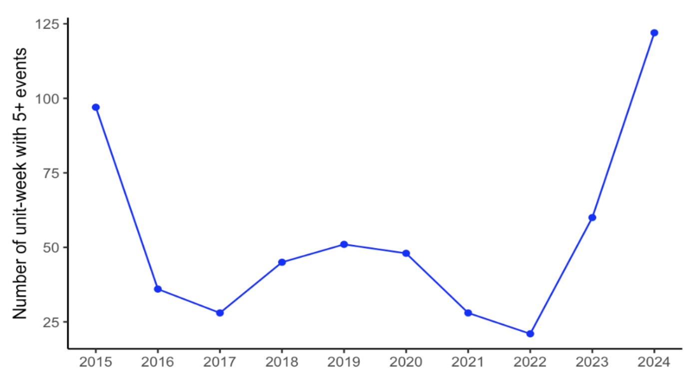

Figure A.21: Number of unit-week observations in each year with 5 or more load-following events.

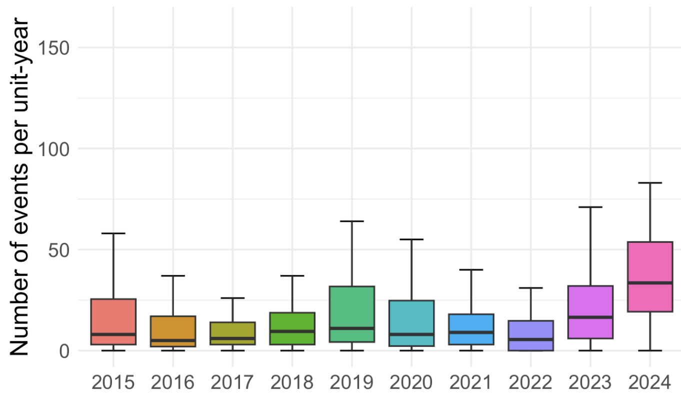

Figure A.22: Distribution (boxplots) of the number of load-following events performed by each unit in a given year.

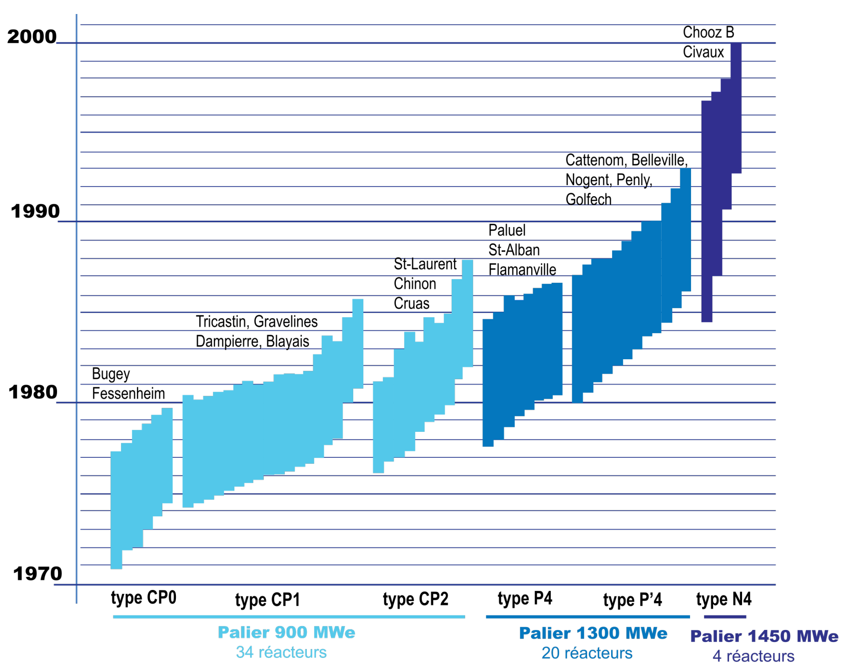

Figure A.23: Different vintages of the French nuclear fleet (source: Wikipedia).

A.2 Additional Tables

	Log(Full-load hours of output)	Log(duration (weeks) of refuel. outage)	Log(1 + # hours of planned outages)	Log(1 + # hours of forced outages)
Log(1 + Full-load hours of lost energy)	0.025	0.05	0.28	0.25
	(0.0042)	(0.029)	(0.055)	(0.050)
Unit FE	Y	Y	Y	Y
Num.Obs.	367	360	367	367
R2	0.673	0.151	0.278	0.302

Table A.5: Estimation results from Equation (4), with variables expressed in logs. The unit of observation is a fuel cycle.

	Full-load hours of output	Duration (weeks) of refuel. outage	# hours of planned outages	# hours of planned outages
Full-load hours of lost energy	0.0028	0.00	1.1	0.4
	(0.0005)	(0.0092)	(0.54)	(0.15)
Unit FE	Y	Y	Y	Y
Num.Obs.	367	360	367	367
R2	0.682	0.207	0.227	0.230

Table A.6: Estimation results from Equation (4), with variables expressed in levels. The unit of observation is a fuel cycle.

	Log(cubic meters of water intakes)	Log(cubic meters of evaporated water)	Log(mean temperature of released water)	Log(mean pH of released water)
Log(1 + lost energy in MWh)	0.000	0.02		
	(0.00407)	(0.023)		
Log(1 + lost energy in MWh per GWh produced)	0.07	0.2		
	(0.018)	(0.15)		
Log(1 + plant output in MWh)			0.049	0.003
			(0.0052)	(0.0014)
plant FE	Y	Y	Y	Y
month-of-year FE	Y	Y	Y	Y
Num.Obs.	1284	1272	764	810
R2	0.960	0.841	0.762	0.585

Table A.7: Estimated associations between load-following and cooling water. Observations are at the month-plant level, when available.

	Log(m^3 of liquid chemicals released)	Log(GBq of tritium)	Log(tritium concentration (GBq/m^3))	Log(kg of boron released)	Log(boron concentration (kg/m^3))
Log(1 + lost energy in MWh)	-0.01	0.04	0.09	-0.01	0.01
	(0.0084)	(0.0081)	(0.019)	(0.0078)	(0.0116)
Log(1 + plant output in MWh)	0.06	0.21		0.06	
	(0.011)	(0.024)		(0.027)	
plant FE	Y	Y	Y	Y	Y
month-of-year FE	Y	Y	Y	Y	Y
Num.Obs.	1,288	1,273	1,273	1,259	1,259
R2	0.560	0.518	0.351	0.400	0.188

Table A.8: Estimated associations between load-following and liquid waste. Observations are at the month-plant level, when available.

Dependent Variable:	Is nominal MOC binding? ($\gamma = 0.25$)				
	(1)	(2)	(3)	(4)	(5)
RES_h (GWh)	0.0085	0.0079			
	(6.8e-05)	(0.00015)			
RD_h (GWh)	-0.00765	-0.00762	-0.00777	-0.00777	-0.0080
	(5.8e-05)	(5.8e-05)	(5.9e-05)	(0.00006)	(6.2e-05)
$W_{RES,h,n}$ (GWh)		0.018			
		(0.0045)			
Solar$_h$ (GWh)			0.0090	0.0088	0.0088
			(7.4e-05)	(0.00017)	(0.00017)
Wind$_h$ (GWh)			0.0074	0.0074	0.0068
			(7.1e-05)	(0.00023)	(0.00023)
$W_{solar,h,n}$ (GWh)				0.006	0.014
				(0.0051)	(0.0051)
$W_{wind,h,n}$ (GWh)				0.00	0.03
				(0.0080)	(0.0083)
$X_{hydro,h,n}$					-0.7
					(0.062)
$X_{therm,h,n}$					-0.9
					(0.11)
Num.Obs.	3,188,823	3,188,823	3,188,823	3,188,823	3,188,823
R2	0.070	0.070	0.071	0.071	0.071

Table A.9: Obtained results when estimating Equation (11). Robust standard errors (HC1) clustered by unit and day-of-sample are reported in parenthesis. Exposure metrics are computed using $\gamma = 0.25$.

Dependent Variable:	Is nominal MOC binding? ($\gamma = 0.75$)				
	(1)	(2)	(3)	(4)	(5)
RES_h (GWh)	0.0085	0.0084			
	(6.8e-05)	(0.0001)			
RD_h (GWh)	-0.00765	-0.00764	-0.00777	-0.00776	-0.00791
	(5.8e-05)	(5.8e-05)	(5.9e-05)	(5.9e-05)	(0.00006)
$W_{RES,h,n}$ (GWh)		0.003			
		(0.0026)			
Solar_h (GWh)			0.0090	0.0085	0.0085
			(7.4e-05)	(0.00013)	(0.00013)
Wind_h (GWh)			0.0074	0.0076	0.0077
			(7.1e-05)	(0.00013)	(0.00013)
$W_{solar,h,n}$ (GWh)				0.017	0.022
				(0.0035)	(0.0036)
$W_{wind,h,n}$ (GWh)				-0.009	-0.007
				(0.0038)	(0.0038)
$X_{hydro,h,n}$					-0.47
					(0.022)
$X_{therm,h,n}$					-0.70
					(0.056)
Num.Obs.	3,188,823	3,188,823	3,188,823	3,188,823	3,188,823
R2	0.070	0.070	0.071	0.071	0.071

Table A.10: Obtained results when estimating Equation (11). Robust standard errors (HC1) clustered by unit and day-of-sample are reported in parenthesis. Exposure metrics are computed using $\gamma = 0.75$.

B Exposure to Other Generation Technologies

To compute our “exposure” metric, we track the evolution of municipality-level installed capacities of each generation technology other than nuclear power.

We rely on a yearly public inventory of power plants published by the French government since 2017. This inventory lists all electricity generation units as of 31 Deccember[42] and notably provides, for each unit, (i) its installed capacity, (ii) its technology (hydro, non-renewable thermal, nuclear, renewable thermal, solar, wind), (iii) the municipality where it is located,[43] and (iv) its commissioning date. In particular, knowing the installed capacity, municipality and commissioning date of each unit allows us to compute municipality-level time series of installed capacity, broken down by technology.[44]

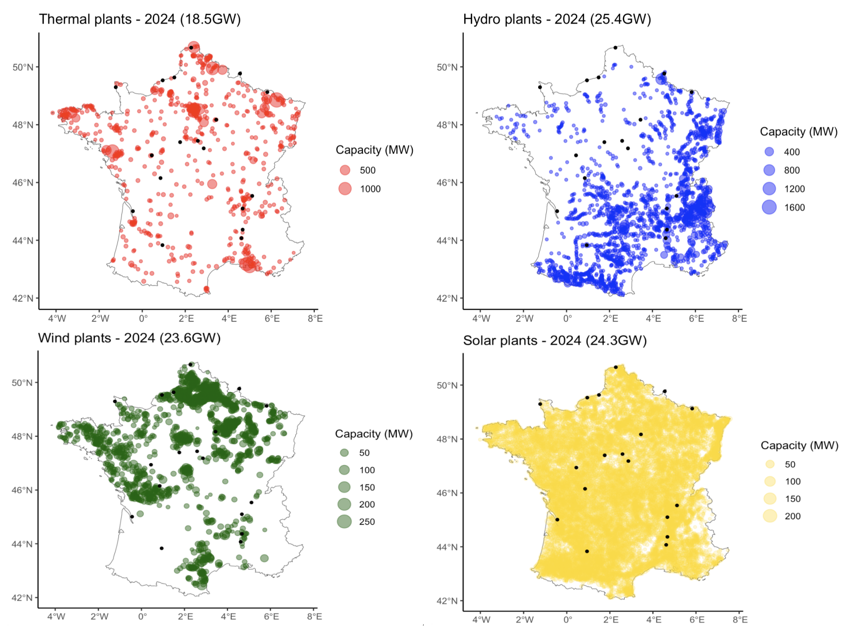

Figure B.24: Locations of non-renewable thermal, hydro, wind and photovoltaic capacities as of 31 December 2024. Nuclear plants (operational as of 31 December 2024) are represented with black circles.

We observe that the installed capacity of hydropower has been largely constant, with

[42]For example, the inventory as of 31 December 2024 is available at `https://www.data.gouv.fr/fr/datasets/registre-national-des-installations-de-production-et-de-stockage-delectricite-au-31-12-2024-1/`, last accessed on 9 April 2025.

[43]Mainland France is divided into 30,000+ such municipalities, which therefore represent a very granular spatial unit.

[44]For privacy reasons, however, small PV installations (lower than 36 kW of capacity) are aggregated at the municipality or departement level (see Astier et al. (2023) for more detail). Because they represent a small fraction of total installed solar capacity (less than 5 GW out of 24 GW as of 31 December 2024), we aggregate for simplicity small PV installations at the departement level (mainland France has 94 such departements). We interpolate these estimates linearly to proxy for the daily departement-level time series of these installations.

nonetheless a mildly increasing trend driven by the commissioning of small-scale run-of-the-river units. Fossil-fueled thermal capacity also followed an almost flat trajectory to reach about 18 GW in 2024, the closure of coal power plants being more than compensated by the opening of a few gas-fired power plants. Finally, the installed capacities of wind and solar have been growing steadily. Installed wind capacity increased from 6.7 GW as of 31 December 2011 to about 24 GW as of 31 December 2024. Solar photovoltaic capacity increased from about 2.5 GW to 24 GW over the same time window.

Figure B.24 shows where the installed capacities of the four main non-nuclear technologies (non-renewable thermal, hydro, wind and photovoltaic) are located as of 31 December 2024. Importantly, the spatial distribution of power plants differs significantly across technologies. Thermal power plants tend to be located close to the main cities, while hydro facilities locate near the main mountain areas and along rivers. Wind power plants were predominantly installed in the North. Although somewhat spread out across the whole territory, solar installations have been more intensively deployed in the South. Overall, Figure B.24 illustrates that the different nuclear power plants are "exposed" differently to other generation technologies, at least in a spatial sense.

From this data, we build a municipality-by-day-of-sample panel of the installed capacity of each technology in each municipality. Because municipalities are small geographical units in France (which has 30,000+ such municipalities), we consider units to be located at the centroid of the municipality where they sit. We denote with $d(i, n)$ the as-the-crow-flies distance between unit i and nuclear unit n. Figure B.25 illustrates these notations in a simple case where $I_{t,\tau} \equiv \{1, 2, 3\}$.

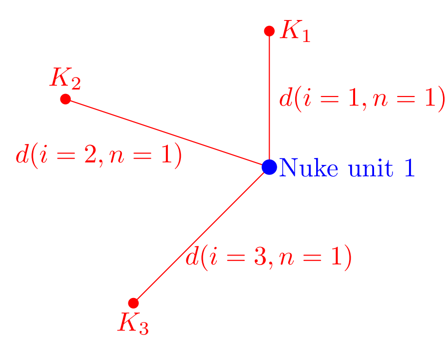

Figure B.25: Notations used in our construction of our "exposure" metric.

We then define the "exposure" $X_{n,t,\tau}$ of nuclear unit n to technology τ at date t as:

$$X_{n,t,\tau} \equiv \sum_{i \in I_{t,\tau}} (\frac{\max(d(i,n), d_0)}{d_0})^{-\gamma} K_i$$

where γ and d_0 are (positive) tuning parameters. In our main specification, we set $d_0 = 1$ km and $\gamma = 0.5$.

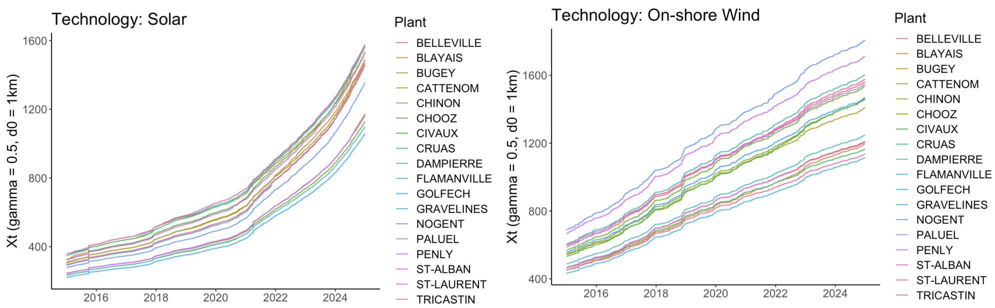

Figure B.26: Obtained exposure metric (in MW) to solar (left panel) and wind (right panel) for the 18 nuclear plants between 1 January 2015 and 31 December 2024 ($d_0 = 1$ km and $\gamma = 0.5$).

Figure B.26 shows the obtained "exposure" metric (in MW) to wind and solar for the different nuclear plants. Note that, for simplicity, we restrict the sample to the balanced panel of the 18 nuclear plants that were in operations during the whole period (i.e., we exclude from this analysis the plant of Fessenheim, which closed in 2020).