Green AI in Industry

Quasi-Experimental Evidence from the Water Treatment Sector*

Mathias Abitbol, Philippe Aghion, Céline Antonin,
Lint Barrage and Benjamin Lengereau†

March 25, 2026

Abstract

This paper studies the environmental effects of the large-scale deployment of AI-driven monitoring systems across French wastewater treatment plants operated by a global leader in water supply services. Exploiting quasi-experimental variation in both the timing of adoption and outages thanks to high frequency data, we estimate the causal impact of AI on energy use and carbon emissions. We find that AI has allowed those plants to reduce their electricity consumption and carbon emissions by 5.4% and 6% respectively, and electricity costs by 8.2%, while improving water treatment quality. The additional electricity demand generated by AI servers represents less than 1% of these savings. Beyond average effects, AI-equipped plants exhibit greater resilience to high operational stress, including extreme meteorological events and chemical pollutant peaks. They also improve load management by reallocating consumption from peak toward off-peak hours. Finally, we assess the aggregate implications of our findings for global climate and welfare using the DICE model. We find large estimated global welfare gains associated with our estimate of AI-induced CO2 reductions.

*We thank Gautier Avril and Nicolas Sourbier from Purecontrol for their critical support through the transmission and handling of the data, and the background research on AI environmental footprint. We are grateful to Morgane Bazillais, Marjolaine Demouge, Julien Lacour, Dinah Louda, Nader Moatamri, Thierry Scanff from Veolia for their help and partnership through this project, in particular for their guidance around operational aspects and the transmission of the necessary additional data. We benefitted from Patricio Velasquez's superb research assistance, from helpful comments by David Arboleda Carcamo, Francis Kramarz, and Adrien Mangold, as well as seminar participants at Collège de France - INSEAD. This project was funded by the European Union (ERC, BTI, GA Number 101198463). The views expressed herein are those of the authors only and do not represent the position of any other party.

†Abitbol: Ecole normale supérieure - PSL, Collège de France. Aghion: Collège de France, LSE, INSEAD. Antonin: Sciences-Po, OFCE. Barrage: ETH Zurich. Lengereau: Collège de France, INSEAD.

1 Introduction

The growing adoption of generative AI has reignited discussions regarding its environmental cost, particularly in terms of energy consumption and emissions (Masanet et al., 2020; Bonfiglioli et al., 2025; de Vries, 2023). However, while much attention has been paid to the potential negative environmental impact of AI, its role in optimizing resource use and facilitating the transition from polluting to green production technologies remains largely overlooked. Consequently, very few studies have attempted to calculate the overall environmental balance of AI adoption.

In this paper, we attempt to fill this gap by evaluating the net effect of AI on the carbon footprint of a company within a real-world industrial context. Specifically, we examine the impact of an AI-driven monitoring solution adopted by a Fortune 500 company and global leader in the water and environmental industry. We evaluate the extent to which AI use has enabled the company to optimize aeration systems in treatment plants, and hence to reduce their energy consumption, operational costs, and carbon emissions. Indeed, machine and deep learning models used in the industry and deployed with digital twins can help control complex processes and optimize their resource use in energy-intensive activities such as water treatment supply. Wastewater treatment is a major contributor to greenhouse gas emissions with a global share of 1.3%, comparable to aviation at 1.9% (Ritchie, 2020). In particular, aeration during the secondary treatment phase alone typically accounts for roughly half of a water treatment plant's total electricity consumption. It is representative of a wide category of energy-intensive industrial processes that operate over flexible durations, and for which predictive AI-based optimization using granular data can unlock significant efficiency gains.

Our study focuses on comparing the upstream data center energy and environmental cost of AI with the downstream energy savings AI induces when implemented for aeration monitoring. In partnership with a service provider, the company we study has rolled out a pioneer large-scale deployment of AI for industry[1] covering more than 200 plants. To our knowledge, it is the first academic study of the impact of AI deployment at scale in the industry and the first study to provide an empirical estimate documenting a net positive energy balance associated to AI. Particularly, the system we study uses high-frequency data (with 15-minute intervals) on electricity prices, weather conditions, sensor streams, and laboratory quality samples in order to schedule and regulate aeration. It aims to minimize energy cost while assuring effluent quality and meeting regulatory norms. We use both the staggered roll-out and the unplanned interruptions of the system across plants as a quasi-natural experiment to evaluate the overall energy balance of using AI to monitor

[1] As a reference, according to Eurostat (2025*b*) data, only 1.1% of European firms in utilities and 1.6% in manufacturing report using AI for decision-making in machine operations.

wastewater treatment, factoring in AI's own energy consumption. Our main finding is that AI allows plants to reduce their electricity consumption and carbon emissions by 5.4% and 6% respectively, and electricity costs by 8.2%, whereas the use of AI across plants only generates a direct upstream increase in electricity consumption of less than 1% of these gains. Even under conservative assumptions as to the required hardware installations for making AI operational in the management of aeration systems, the overall energy and carbon emission costs of AI remain well below its induced energy and emission savings in water treatment.

This paper contributes to a nascent literature on the environmental costs and benefits of AI. Studies by international institutions highlight the growing energy footprint of AI systems and their uncertain environmental impacts at the macro level (Bogmans et al., 2025; Burian and Stalla-Bourdillon, 2025; IEA, 2025). We complement this perspective by studying a distinct, micro-founded channel through which AI may affect energy use, the so-called "twin transition" (Muench et al., 2022; McKinsey, 2024): machine-learning-based predictive control and maintenance that improve the operation of industrial processes (Carvalho et al., 2019; Zonta et al., 2020). Treating AI as a single and homogeneous technology obscures critical heterogeneity between general-purpose and specialized control technologies, and between cloud-based and edge deployment. In computer science, Schwartz et al. (2020) introduced the notion of "Green AI", emphasizing performance gains that do not come at rising computational cost, and ideally reduce it. We bring this concept to the economic literature to frame circumstances under which AI can plausibly contribute to decarbonization, particularly in energy-intensive industrial settings.

While the energy costs of AI are increasingly documented, far less is known at the micro level about whether and when AI delivers energy savings in the real economy. A growing set of sectoral assessments discusses AI-driven efficiency opportunities (Yussuf and Asfour, 2024; Rolnick et al., 2022). Two studies provide micro-econometric evidence on how AI adoption affects energy use. Metcalfe et al. (2026) investigate AI's effect in households, focusing on EV charging, and find significant reallocation effects but no overall reduction in energy demand. Luo, Wang and Chen (2025) document that exposure to AI-related policy incentives is associated with improvements in corporate energy efficiency, but do not identify the operational mechanisms behind these gains. Moreover, both papers connect only indirectly to the Green AI perspective in the sense that neither provides a comprehensive accounting of AI's *net* energy footprint or climate benefits.

This paper also connects to a growing literature using laboratory, field, and quasi-experimental evidence to study the productivity impacts of artificial intelligence. AI has rapidly diffused across a wide range of domains, including services (Ashfaq et al., 2020;

Chattaraman et al., 2019; Huang and Rust, 2022; Li et al., 2021), marketing (McLean, Osei-Frimpong and Barhorst, 2021), and human resource management (Akter et al., 2021; Chatterjee et al., 2021). Recent studies document sizable productivity gains in specific tasks such as customer service (Brynjolfsson, Li and Raymond, 2025), consulting (Dell'Acqua et al., 2023), coding (Peng et al., 2023; Cui et al., 2025), and professional writing (Noy and Zhang, 2023). Yet, the implications of AI for industrial operations remain comparatively understudied, despite the fact that industrial processes are both highly energy-intensive and particularly well-suited to data-driven optimization.

The existing literature on AI and energy leaves two important gaps. First, to our knowledge, no empirical study measures the net energy impact of AI adoption, accounting simultaneously for the energy consumed by AI systems and the operational energy savings they generate. Second, most existing evidence relies on proxies or aggregate indicators, with little direct observation of energy use at the level of production processes. This paper bridges these gaps by providing a firm-level, multi-site assessment of the net energy impact of AI adoption within a major industrial group. Using detailed operational and metered data, we quantify both the electricity consumption attributable to AI systems and the energy savings and emissions reductions generated through AI-driven process optimization. This setting offers a unique opportunity to evaluate the environmental consequences of AI in a real-world industrial context, at a scale rarely observed in the literature. By combining micro-level operational data with a comprehensive assessment of the firm's total energy footprint, we deliver an integrated evaluation of AI's net energy and emissions effects within a large industrial organization. In doing so, this paper makes three key contributions: (i) it provides causal, micro-level evidence on the energy consequences of AI adoption, (ii) it offers the first empirical quantification of the net energy balance of a so-called "Green AI", and (iii) it extends AI and energy research to the industry, a critical but understudied sector for AI economics.

The remaining part of the paper is organized as follows. Section 2 describes the institutional context. Section 3 outlines our empirical strategy. Section 4 presents our data and descriptive statistics. Section 5 presents our main findings, namely the computation of the net AI energy balance. Section 6 explores two key dynamic margins of AI controls: price-sensitive intertemporal reallocation of the activity and responses to weather-induced stress. Section 7 assesses the aggregate implications of our findings for global climate and welfare using the DICE climate-economic model. Section 8 concludes.

2 Context and Institutional Setting

This section provides background on the technology and the industry setting we study. We first describe how AI-based control systems deployed as digital twins are used to optimize industrial operations with energy cost reduction objectives. We then summarize key facts about wastewater collection and treatment, with an emphasis on the energy intensity of the treatment processes. Next, we explain why the aeration process in wastewater treatment plants, which accounts for the majority of energy consumption, has significant potential for AI control. We also describe the digital-twin solution adopted by the company for its wastewater treatment plants (WWTPs). Finally, we outline the staggered roll-out that generates the variation exploited in our empirical analysis.

2.1 AI and Digital Twins for Energy Management

A growing literature argues that AI can extend the scope of traditional information and communication technologies (ICT) in the optimization of complex systems (including energy, transport, agriculture, and buildings) by improving forecasting, monitoring, and real-time control (Rolnick et al., 2022; Stern et al., 2025). In industrial settings, the main promise of AI lies in its ability to combine high-frequency sensor data with predictive models to adjust operating conditions dynamically, thereby increasing process efficiency and flexibility. This has generated substantial interest among industrial operators seeking to reduce energy costs and meet decarbonization targets (Andres, Dugoua and Dumas, 2022; Stern et al., 2025). In principle, AI-enabled tools can (i) anticipate operational shocks (including those driven by environmental conditions), (ii) optimize the timing and intensity of energy-intensive steps, and (iii) reduce the reliance on coarse, threshold-based rules (Abdessadak et al., 2025).

Digital twins are one of the most prominent AI implementations in industry. They consist in a virtual representation of a physical asset or process that is continuously updated using real-time data (typically from multiple sensors) and is used to simulate, predict, and optimize system behavior (Fuller et al., 2020; Abdessadak et al., 2025; IBM, 2024). In practice, digital twins combine engineering constraints with statistical or machine-learning models to produce operational recommendations (or automated setpoints) under uncertainty. Digital twins are increasingly viewed as an enabling technology for remote asset management and operational efficiency within the broader Internet of Things (IoT) ecosystem. Their applications span aerospace, health, smart-city infrastructure, and energy systems (Attaran and Celik, 2023). Conceptually, a digital twin can be used to optimize performance, predict failures and schedule maintenance, and evaluate counterfactual scenarios without disrupting production, as real-time feedback supports continuous model updating and learning.

Industry reports suggest substantial potential gains from digital twins, and specifically potentialities for process optimization and energy efficiency (McKinsey, 2024; Ba et al., 2025)[2]. From a climate-policy perspective, the industrial sector is frequently highlighted as a domain in which AI-enabled optimization could deliver meaningful emissions reductions, though the scalability of such gains remains uncertain (IEA, 2025). IEA (2025) estimate, for instance, that under a scenario of widespread AI adoption, the industrial sector holds the largest abatement potential by 2035, with more than 750 Mt CO_2 avoided in this sector alone. Yet these projections rest on aggregate modelling instead of micro-founded evidence of efficiency gains in actual deployments, and whether such potential translates into real-world outcomes remains an open empirical question.

2.2 Key Facts on Wastewater Treatment Plants

Process overview and the central role of aeration. A WWTP treats wastewater from households, industrial sources, and stormwater before discharge to the environment. In Appendix Section A, we present a synthetic diagram and an aerial photo of a WWTP. Treatment typically proceeds in stages. Preliminary treatment removes large debris through screens and grit removal. Primary treatment then relies on settling or flotation to remove suspended solids. Secondary treatment biodegrades organic matter using microorganisms. Like most plants in Europe, the plants in our setting operate with activated sludge systems (Rettig et al., 2018). In those processes, wastewater is aerated to provide oxygen to bacteria, which consume organic pollutants. Aeration is essential because it provides dissolved oxygen used by aerobic bacteria to sustain the biological degradation reactions that drive pollutant removal. Some plants also operate tertiary processes to meet more stringent effluent-quality standards, including filtration, nutrient removal, chemical treatment, and disinfection.

From an energy perspective, secondary treatment, and in particular aeration in activated sludge systems, is typically the most electricity-intensive component of plant operations (IEA, 2017; US DoE, 2017). Aeration must deliver sufficient oxygen to sustain biological treatment, but excessive aeration is costly. The operational problem is therefore inherently dynamic: optimal aeration depends on influent characteristics, process conditions, regulatory constraints, and short-run environmental shocks.

The water-energy nexus. Wastewater treatment is a major contributor to electricity demand. In the US and Europe, almost all WWTP of significant size engage in both

[2]Market reports project rapid growth in the digital-twin industry, from roughly $22 billion in 2025 to $150 billion by 2030 (MarketsandMarkets, 2025) and a surge in the R&D in such technologies as documented in Appendix Figure C.2.

primary and secondary treatment (Eurostat, 2025*a*; EPA, 2022). The IEA (2017) provides a systematic assessment of energy use across the water sector and estimates that the sector consumed roughly 120 million tons of oil equivalent in 2014, with electricity accounting for around 60% (about 820 TWh). In particular, wastewater treatment is widespread around the world. It currently accounts for around 0.35 - 1.1% of global electricity consumption (Magni et al. 2025), or roughly 200 TWh in 2014, comparable to the consumption of a mid-sized country like Turkey (IEA, 2017). In high-income regions where wastewater treatment is widespread, it can account for a large share of the electricity consumed by the overall water sector.

2.3 AI-Based Control in Wastewater Treatment

Why WWTPs are a natural environment for AI control. WWTP operations rely on high-frequency sensor data, nonlinear process dynamics, multiple constraints tied to effluent-quality regulation, and substantial energy costs concentrated in a small number of controllable components (notably aeration blowers). Operators have long depended on industrial automation and human oversight, typically implementing rule-based logic triggered by thresholds. While robust, threshold-based control often compresses information into a small set of indicators and may fail to exploit the richness of available data, especially when operating conditions change rapidly (e.g., during stormwater inflow events). AI-based control can draw on a larger information set to precisely adjust aeration cycles while maintaining treatment quality, by predicting the water pollution load and then accurately calculating the operation of aerators according to this load, taking into account peak and off-peak electricity prices.

The digital twin solution for WWTPs. Our setting focuses on Veolia Water, a subsidiary of the conglomerate Veolia, which is a global leader in Environmental Services and a Fortune 500 company[3]. It has adopted a digital-twin solution developed by the French startup Purecontrol[4]. The Purecontrol system builds a plant-specific digital twin by combining multiple prediction and characterization models for key process variables. In practice, the modeling strategy is pragmatic: models are selected and calibrated to deliver adequate predictive performance at reasonable computational cost, and the final system typically embeds a portfolio of methods (e.g. linear and penalized regressions, dynamic equations, neural networks, reinforcement learning algorithms, time-series components, splines), tuned to each plant .

[3]Hereafter referred to by its brand name "Veolia".

[4]Veolia is a shareholder in the start-up company, but Veolia is not the only user of this AI solution. Purecontrol has also a range of other industrial applications than WWTP, and has also partnered within the same sector with all major French water operators (Suez and Saur) as well as local authorities (e.g. the city of Rennes), reflecting broader interest in AI-based control tools for water services

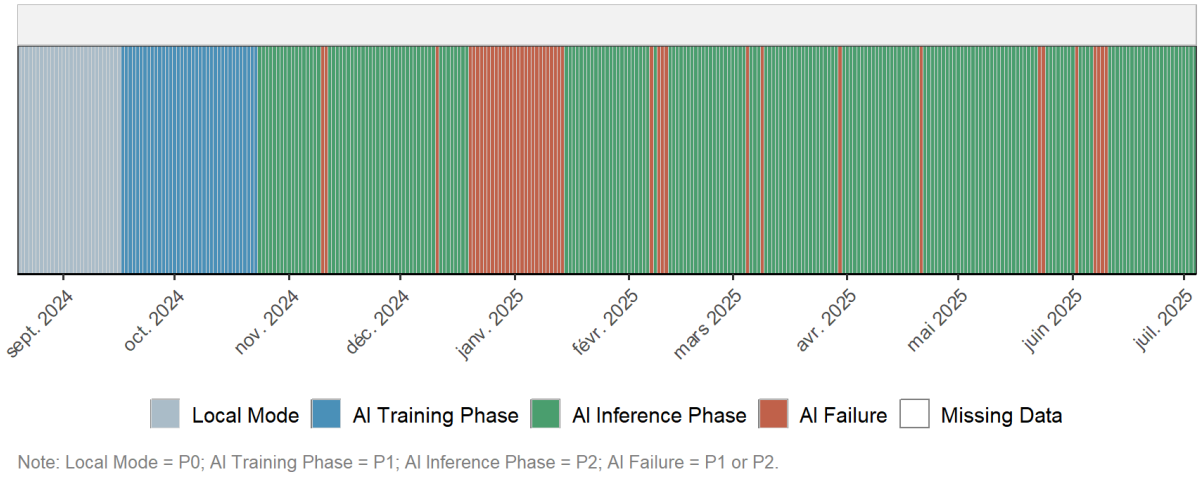
The system then optimizes aeration timing and intensity to minimize electricity costs subject to operational and regulatory constraints (effluent-quality thresholds)[5]. This set-up is aimed at reducing blower operating hours and shifting aeration cycles toward off-peak electricity periods. The solution is designed to integrate with existing industrial automation. It includes (i) a real-time monitoring interface for plant staff, (ii) a prediction layer for variables relevant to treatment dynamics, (iii) real-time tracking of energy consumption by equipment and consumption profiling, and (iv) an execution layer that transmits optimized setpoints to the plant's automation system. Inputs combine internal data (e.g. redox or dissolved oxygen measures, flow rates, manual pollutant measurements) with external information (e.g. electricity tariffs, weather forecasts, and grid carbon intensity).

2.4 AI Deployment and Rollout

Implementation stages within a plant. Deployment of the AI system in wastewater treatment plants follows a multi-stage process. Plants initially operate in "local mode" (P0), relying on on-site automation and pre-defined management thresholds. Aeration and pumping equipment is governed by Programmable Logic Controllers (PLCs) that activate or modulate actuators based on fixed setpoints (for instance, maintaining dissolved oxygen above a given concentration regardless of incoming organic load). These thresholds are set during plant commissioning and rarely updated, meaning the system cannot respond dynamically to diurnal cycles, seasonal variations, or shock loads. As a result, plants routinely aerate beyond biological needs, wasting energy while still risking compliance failures during peak stress events. Once connected to the AI platform, they transition to an intermediate stage referred to as "AI Training Phase" (P1), which typically lasts around one month. During this phase, operations remain broadly similar to local mode, but threshold-setting and calibration are performed centrally by the AI company using data streams from the plant, allowing the AI models to be progressively fine-tuned. This training phase allows to primarily collect data for model calibration, and its duration varies with plant complexity and the speed at which the digital twin converges to actual plant behavior. Although P1 might be expected to improve operational performance relative to local mode, we find no empirical evidence supporting this claim. The final stage, "AI implemented" (P2), represents full AI deployment: the system actively manages plant operations by modeling optimal behavior to minimize electricity costs while satisfying regulatory constraints. Figure 1 illustrates the typical implementation timeline for a plant.

[5]Objective functions can differ across deployments (e.g., carbon intensity minimization in some cases), but for the plants in our sample the stated objective is cost minimization, as confirmed in discussions with Purecontrol engineers.

Figure 1: Adoption Timeline of a typical plant



Notes. This figure plots the daily distribution of AI between phases for a typical facility in our sample from August 2024 to April 2025. AI Inference Phase starts in late October 2024.

Non-absorbing exposure and operational interruptions. AI operation is not an absorbing state in our context[6]. Plants can temporarily revert from "AI on" to "AI off" for several reasons. These interruptions can be deemed plausibly exogenous to local staff (e.g. connectivity outages), while others are partly endogenous (e.g. prolonged gaps in required manual sampling or deliberate managerial deactivation). Our data allow us to distinguish between (i) interruptions in which AI remains authorized to pilot but cannot operate (e.g. technical failures) and (ii) explicit deactivations in which authorization is turned off. We propose a robustness check eliminating partly endogenous deactivations from the sample, and show it does not affect the results.

Company-wide rollout. The company implemented the system in multiple waves. In a very early pilot phase, two plants were selected to test the AI system. A first deployment wave then targeted 50 plants, with recruitment beginning in March 2021 and the first AI activation occurring in October 2022. A second, larger-scale wave was launched in November 2023, with Veolia setting an objective of approximately 500 plants. As of July 4, 2025, the end of our observation period, 204 plants in this second wave had already had AI switched on, and 25 additional ones were in the training phase (P1). Across both waves, we recover data for 283 plants in total, of which 183 were AI-managed by the end of the period. Our analytical sample is restricted to 164 plants (of which 139 were AI-managed at the end of the period), as the main variables of interest are not available for all sites. Each plant might have one or multiple basins[7].

[6]In econometric terms, an "absorbing" (or staggered) treatment is one that, once adopted, cannot be exited (de Chaisemartin et al., 2022).

[7]While a plant may consist of multiple basins, AI may be installed and operating only on one. However, for simplicity, we aggregate data at the plant level for our analysis, as basin-level data is unavailable for most plants. To the extent that this aggregation assigns treatment status to only partly

Appendix Figure B.1 plots the number of plants observed in our database over time (regardless of operating mode): plants are differentiated according to AI installation status (in (a)) and AI operating mode (in (b)). While both series increase steadily, the aggregate deployment measure masks within-plant dynamics, as plants may experience intermittent interruptions that temporarily suspend AI operation.

The general eligibility conditions required plants to have a nominal capacity included between 10,000 to 50,000 population equivalents (PE) and to use activated sludge technology.[8] Plants below this threshold were excluded on the grounds that AI installation was not deemed profitable, while plants above 100'000 were managed through AI systems developed internally by Veolia, which were given preference over Purecontrol's solution. Within the eligible capacity range, Veolia operates approximately 500 plants in Metropolitan France, all of which are expected to be AI-managed by the end of the second wave. Plants volunteered for participation and were selected based on compatibility with technical requirements, including the availability of suitable automation interfaces.[9] Plants included in the initial wave were those for which AI could be installed without any modification to the existing IT infrastructure. Subsequently, the selection order among volunteer plants was quasi-random, subject only to the constraint that plants within each sub-wave were evenly distributed across French regions.

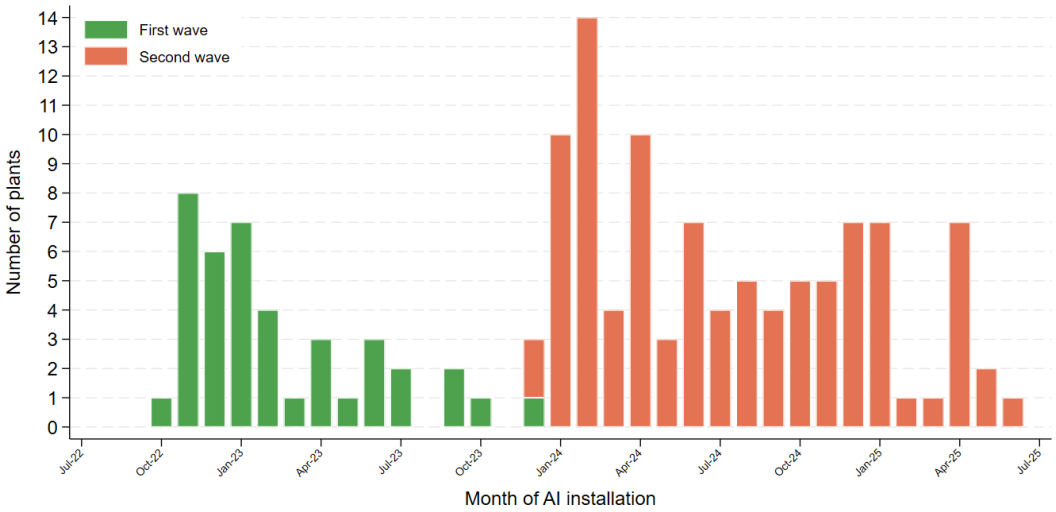
Figure 2 presents the timeline of AI adoption across plants. The initial phase targeted 50 plants, with recruitment beginning in March 2021. In the first wave, the first implementation of AI started in October 2022. In a second phase, starting in November 2023, Veolia set an objective to reach a larger-scale deployment of several hundreds plants by the end of 2025. To the date of July 4, 2025, for this second wave, 204 already had AI switched on, and 25 additional ones were in the training phase. In total, for the two periods, we recover the data for 283 plants, of which only 183 are AI-managed at the end of the period. We have exploitable data for 164 plants (including 139 AI-managed at the end of the period) because the main variables of interest are not recorded across all sites.

treated plants, it should only bias our estimates downwards.

[8] In the initial phase, a small number of plants operated under different technologies and sizes. Priority was given to activated sludge plants as this was a more standardized technology allowing for faster deployment. Toward the end of the second wave, a few plants with a nominal capacity between 50,000 and 100'000 PE were also admitted into the scheme.

[9] Some connectivity requirements exist in order for the plant to be eligible, but according to company staff, almost all plants in the relevant capacity range satisfied the connectivity condition. Process measurement requirements include real-time flow measurement, at least weekly measurements of NH_4 and NO_3, and sensors for redox potential or dissolved oxygen.

Figure 2: AI adoption across plants



Notes. This figure plots the number of plants adopting AI each month, distinguishing between the two main adoption waves.

3 Empirical Strategy

3.1 Variable Definitions

Outcome variables. Our outcomes capture the energy and environmental footprint of plant i over period t. We focus on: (i) electricity consumption, (ii) electricity expenditure, (iii) CO_2-equivalent emissions.

Let $E_{i,h}$ denote electricity consumption (in kWh) for plant i in hour h, and let $P_{i,h}$ denote the plant- and hour-specific electricity price. For any period t (e.g., day or week) consisting of hours $h \in \mathcal{H}(t)$, we define:

$$E_{i,t} = \sum_{h \in \mathcal{H}(t)} E_{i,h}, \tag{1}$$

$$\text{Expenditure}_{i,t} = \sum_{h \in \mathcal{H}(t)} P_{i,h} E_{i,h}. \tag{2}$$

Unlike physical consumption, expenditure reflects both total usage and the timing of electricity use across hours (e.g., shifting load to off-peak periods).

Finally, we compute CO_2-equivalent emissions by weighting hourly consumption by the grid emissions factor κ_h:

$$CO_2e_{i,t} = \sum_{h \in \mathcal{H}(t)} \kappa_h E_{i,h}.$$

Because κ_h is driven by the external generation mix, we interpret estimates on CO_2e as reduced-form effects that combine operational responses with exogenous variation in the

grid.

Treatment variables: AI exposure as a non-absorbing intensity. Our treatment design is non-standard in two respects. First, exposure is non-absorbing: even after installation, AI can stop operating due to connectivity outages, missing required manual inputs, or managerial deactivation. Second, the treatment intensity is continuous because AI can be used a fraction of time during the day.

We construct the three following treatment measures. For each plant i and period T, let $D_{i,T}$ denote the (continuous) intensity of AI use. We consider: (i) the intensity $D_{i,T}$ itself; (ii) an indicator for any exposure $\mathbb{1}\{D_{i,T} > 0\}$; and (iii) an "installed" indicator, equal to one once AI has operated at least once at plant i, regardless of whether it is active in period T, i.e. $\mathbb{1}\{\sum_{t=0}^{T} D_{i,T} > 0\}$. This last measure is absorbing. Hereafter, we call these three indicators respectively: *AI usage*, *Treated* and *AI installed*. Unless otherwise specified, our baseline specification relies on the continuous treatment intensity $D_{i,t}$.

3.2 DID estimators for non-binary non-absorbing treatments

The canonical TWFE estimator is subject to well-known biases under treatment-effect heterogeneity, notably contamination from already-treated units and negative weighting of cohort-specific treatment. While economists long believed that the parallel trends assumption was the only requirement for TWFE to deliver a causal estimate in staggered adoption designs, the recent difference-in-differences literature has shown that TWFE regressions yield consistent estimates only under strong homogeneity assumptions, and may be biased when treatment effects vary over time or across adoption cohorts (de Chaisemartin and D'Haultfoeuille, 2020; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Goodman-Bacon, 2021; Borusyak, Jaravel and Spiess, 2024). This concern is salient in our setting, where learning-by-doing may generate larger effects for later exposure spells or later-treated cohorts. Our core estimates rely on the robust DID estimator of de Chaisemartin and D'Haultfoeuille (2024). We additionally report TWFE results when analyzing fixed structural differences (e.g. peak versus off-peak hours), where TWFE provides a valid and transparent decomposition that the DCDH estimator is not designed to deliver.

Most staggered DID estimators and commands are designed for binary treatment variable and absorbing adoption patterns. In contrast, our setting features a non-absorbing treatment with repeated treatment status changes, as well as a potentially non-binary (discrete or continuous) treatment intensity[10]. We therefore rely on the heterogeneity-robust

[10]Treatment intensity varies not in terms of quality, but in terms of daily share of time. In our setting, the treatment variable is binary or not depending on whether we simply adopt a dummy variable or *AI usage*.

DID estimators for dynamic effects from de Chaisemartin and D'Haultfoeuille (2024)[11]. These estimators are explicitly designed to accommodate treatments that may change over time, and to remain valid under treatment-effect heterogeneity where standard two-way fixed effects can be biased.

DID estimator for continuous and non absorbing treatment. Let i index plants and define the first time the plant's exposure changes:

$$F_i = \min\{t \geq 2 : D_{i,t} \neq D_{i,t-1}\}.$$

where $D_{i,t}$ is the treatment status of unit i at time t.

For horizon $\ell \geq 1$, define the control set of "not-yet-changed" plants sharing the same baseline exposure $D_{i,1}$:

$$\mathcal{C}_{i,\ell} = \{j : D_{j,1} = D_{i,1},\ F_j > F_i - 1 + \ell\}.$$

In our setting, all plants except one start from a common baseline exposure $D_{i,1} = 0$, i.e. no plant has AI activated at the beginning of the sample period. This ensures that the control set $\mathcal{C}_{i,\ell}$ is well-defined and non-empty for all switchers, and that comparisons across units are not contaminated by heterogeneous pre-treatment exposure levels.

The group-specific event-study DID is:

$$DID_{i,\ell} = \big(Y_{i,F_i-1+\ell} - Y_{i,F_i-1}\big) - \frac{1}{\#\mathcal{C}_{i,\ell}} \sum_{j \in \mathcal{C}_{i,\ell}} \big(Y_{j,F_i-1+\ell} - Y_{j,F_i-1}\big).$$

Averaging across switchers observed at horizon ℓ yields DID_ℓ.

The estimator relies on (i) no anticipation, (ii) no interference, and (iii) a parallel-trends assumption for the relevant counterfactual "status-quo" outcomes, conditional on fixed effects and controls[12]. Under this set of assumptions, this estimator is equal to the expected difference between g's actual outcome at $F_{g-1+\ell}$ and its counterfactual "status-quo" outcome:

$$\delta_{i,\ell} = \mathbb{E}[Y_{i,F_g-1+\ell} - Y_{g,F_g-1+\ell}(D_{g,1}...D_{g,1}) \mid \mathbf{D}]$$

We report placebo (lead) estimates in an event-study plot to assess the plausibility of assumption (iii). Assumption (i) can be tested by comparing plant behavior before and

[11] Implemented in the `did_multiplegt_dyn` package.

[12] The estimator only compares switchers and non-switchers with same period-one treatment, whereas the "naively-extended" Callaway and Sant'Anna (2021) estimator compares switchers and non-switchers with different period-one treatments, thus relying on an unconditional parallel trends assumption for status-quo outcome.

after notification of future treatment. Lastly, we expect Assumption (ii) to hold given the nature of our setting. Wastewater allocation to different plants is fixed by infrastructure design that long pre-dates our analysis. Even if, e.g., managers of treated vs. untreated plants were to interact during the roll-out, the untreated plants' managers cannot utilize the AI technology without installation, which would be observable in our setting.

We estimate the effects of AI by aggregating the data at the weekly frequency for the use of the de Chaisemartin and D'Haultfoeuille (2024) estimator. In this framework, the choice of time aggregation is consequential because the estimator constructs counterfactual outcomes over the same horizon as the unit of observation. While we report results for daily, weekly, and monthly time aggregation, our preferred specification considers weekly data. Monthly aggregation is too coarse as it smooths substantial within-month variation in AI activation into average treatment intensities and reduces the number of identifiable treatment episodes as well as the pool of units that can serve as valid controls. Daily aggregation, by contrast, is subject to more noise as electricity consumption exhibits pronounced intra-week volatility and may obfuscate the benefits of shifting electricity consumption within days of the week. Weekly aggregation preserves meaningful variation in treatment while allowing effects to accumulate over several days. Appendix Section E.3 compares results across alternative time aggregations.

4 Data and descriptive statistics

4.1 Data

Our analysis relies on high-frequency monitoring data from Purecontrol on Veolia wastewater treatment plants under its management, complemented by Veolia administrative data and weather data from Meteo-France. The main variables of interests derived from the Purecontrol database forming our main sample are :

(1) Share of time of each plant is in control mode, measuring the fraction of time during which the plant operates under threshold-based control either by the plant management (P0) or by Purecontrol (P1), and the optimized AI control mode (P2). Control modes are reported as a continuous variable between 0 and 1 at the daily level[13];

(2) Electricity active power, observed mostly at a 10-minute resolution and expressed in kilowatts (kW). In the analysis, it is aggregated at the hourly level and converted into electricity consumption, expressed in kilowatt-hours (kWh);

(3) Carbon footprint, retrieved from RTE (2026)'s open source data. This variable measures the CO2-equivalent emissions associated with the plant's electricity consumption[14],

[13]The data is only available at the daily level.

[14]The emissions coefficient is national in scope as is the French electricity market. One caveat is that

expressed in kg and observed at the hourly level;

(4) Hourly electricity prices. As the data typically distinguish only two tariff periods at the day-level, we use this variable to get a peak / off-peak hours indicator, with off-peak hours typically occurring between 10 pm and 5 am (Figure D.7);

(5) Electricity costs (in current €) are reconstructed by computing, at the hourly level, the product of electricity consumption and the corresponding electricity price.

(6) Plant management authorization status. This dummy variable captures whether plant management authorizes Purecontrol to control the aeration process using either AI-optimized mode (P2) or manual mode (P1).

To enrich our primary dataset, we include additional variables which are less commonly available than the three score indicators mentioned above, including:

(7) Nominal capacity, available for most plants and measured in Population Equivalent (PE), which corresponds to the maximum load of organic pollution that the plant is built to handle under standard operating conditions;

(8) Inflow rate, which represents the volume of water treated by the plant, measured in cubic meters (m^3);

(9) Redox, measured in millivolts (mV), which reflects the oxidation-reduction potential (ORP) of treated water and is central to aeration control;

(10) Daily average concentrations of ammonium, nitrate, and phosphate, measured at the end of the water treatment process, as well as the plant and time-specific quality thresholds for each variable. All AI-managed plants were required to conduct at least every ten days a test on ammonium and nitrate concentration[15];

(11) Rainfall data from Meteo-France, extracted at the hourly level and matched with plant-level data thanks to plants' geolocation data.

4.2 Summary Statistics

Table 1 presents the summary statistics of the sample, divided into four groups: all plants (All); plants where AI has not been deployed during our sample period yet (Never treated); observations of plants before the first deployment of AI (Treated, pre); and observations after the first deployment of AI (Treated, post). The sample comprises 2.2 million hourly observations on electricity consumption, carbon footprint, electricity prices and the share of time during which AI is used, across 164 plants. Within the sample, approximately one quarter of the plants are never equipped with AI. The sample data spans from January 2022 to July 2025 (Appendix Figure D.1). The first plant began operating under AI con-

the coefficient only reflects French electricity production and thus misses emissions from imports.

[15]Within the week, the day at which the sampling was conducted was not random, but there is no reason to suspect that the pattern of sampling was different before and after AI adoption.

trol in October 2022. The start date of electricity power measurement varies across plants but typically precedes the adoption of AI-managed (P2) mode by several months[16]. AI is deployed on selected plants through two main deployment waves. Appendix table D.1 provide the summary statistics by adoption wave. There are no major differences between the stations from the two adoption waves, apart from the fact that the stations from the first wave are slightly smaller. In Appendix table D.2 we report summary statistics for the subsample restricted to observations for which all control variables - inflow, rainfall and redox - are available.

Figure 3 displays the geographic distribution and size, measured by nominal capacity in Population Equivalent (PE), of the plants in our sample. Facilities are spread relatively evenly across metropolitan France. As shown in Appendix Section D.2, the sample is heavily skewed toward medium and large plants: while roughly 80% of the 21,474 French WWTPs recorded in 2016 fall below 2,000 PE, only two such plants are included. By contrast, the sample comprises 68 medium-sized plants (2,000 - 10,000 PE) and 93 large plants (10,000 - 100,000 PE), representing 3% and 8% of their respective national categories. This reflects a technical constraint: the AI device is poorly suited to smaller installations. That being said, small plants collectively account for only 9% of total installed capacity nationwide (CGDD, 2019), so their absence from the sample has limited consequences for the representativeness of our estimates in terms of overall wastewater sector footprint.

In Figure 4, we plot the raw distribution of the main variables of interest, presented in the previous section: the outcome variables (Electricity consumption, Carbon footprint, Electricity cost) of the never-treated group (AI never installed), as well as the treated group when AI is operating (With AI), the treated group before AI installation and when AI is not operating (Without AI), and the explanatory variable (AI control) for treated units. We observe that plants operating with AI consume less electricity and emit less carbon emissions relative to both plants never-treated and plants from the treated group operating without AI. However, AI does not seem to have an impact on electricity costs, but it can be explained by the average increase in electricity prices between pre-treatment and post-treatment periods (Figure D.5) especially during winter (Figure D.6).

Table 2 presents summary statistics on AI interruptions. Three distinct sources may cause AI to cease operating. Connectivity outages prevent the transmission of plant sensor data to Purecontrol, generating failures that can be considered exogenous to both

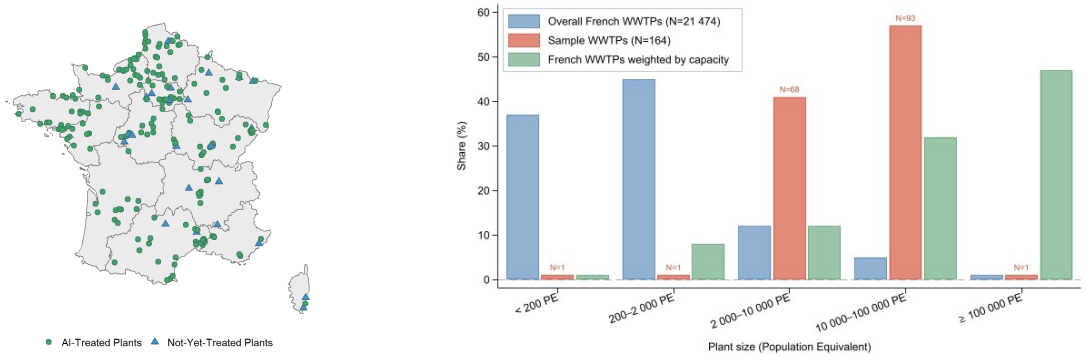
[16]For a plant to get enrolled in the AI-managed system, it needs at least a AI training phase during which Purecontrol teams take control over the plant that is therefore no longer managed locally. It is only when the training phase is completed that the plant gets AI-managed. The duration of this step might vary across plants but typically lasts 2 to 3 months.

plant operations and AI performance. In other cases, the nature of the interruption is less clear-cut, as the underlying cause cannot be straightforwardly classified as exogenous, arising either from plant managers' decisions to disable AI management or from the Pure-control team reverting to manual control (P1). Leveraging data on manager authorization status and the intraday distribution of control modes at the time of failure, we classify interruptions as follows: *unauthorized* when AI management has not been enabled by the plant; *AI-team decision* when local plant control (P0) accounts for less than 10% of the day, indicating that the AI team has assumed manual control for most of that day; and *connection failure* otherwise. As shown in Table 2, the majority of interruptions are endogenous: 46% of failure-hours correspond to unauthorized AI management, and 24% reflect deliberate reversions to manual control by the AI team, leaving 30% attributable to exogenous connectivity outages. These endogenous interruptions introduce a potential downward bias in our estimates of the AI treatment effect, insofar as AI is more likely to be disabled precisely in circumstances where plant outcomes would have been unfavorable regardless. We address this concern in the robustness checks of Section 5.1 by restricting the sample to weeks free of potentially endogenous failures.

Table 1: Summary Statistics

Variable	**All**	**Never treated**	**Treated, pre**	**Treated, post**
Observations	2,242,511	148,211	585,712	1,508,588
Plants	164	25	138	139
Nominal Capacity (PE)	19,804	17,913	20,283	20,144
AI usage (Share of time)	0.39	0	0	0.58
Power (kWh)	66.9	80.4	65.4	66.1
Carbon em. ($kgCO_2e/h$)	1.8	1.9	2.25	1.62
Electricity price (€/kWh)	0.11	0.12	0.09	0.11
Electricity cost (€/h)	7.23	10.4	6.16	7.32

Notes: This table shows summary statitics of our dataset, presenting the number of observations and plants in each group and then the average of the main outcome variables. The first column consists of all plants in our sample, the second column includes plants who never receive coverage by AI. The third column presents statistics for treated plants before they receive AI installation, and the fourth column includes treated plants after the first AI deployment.

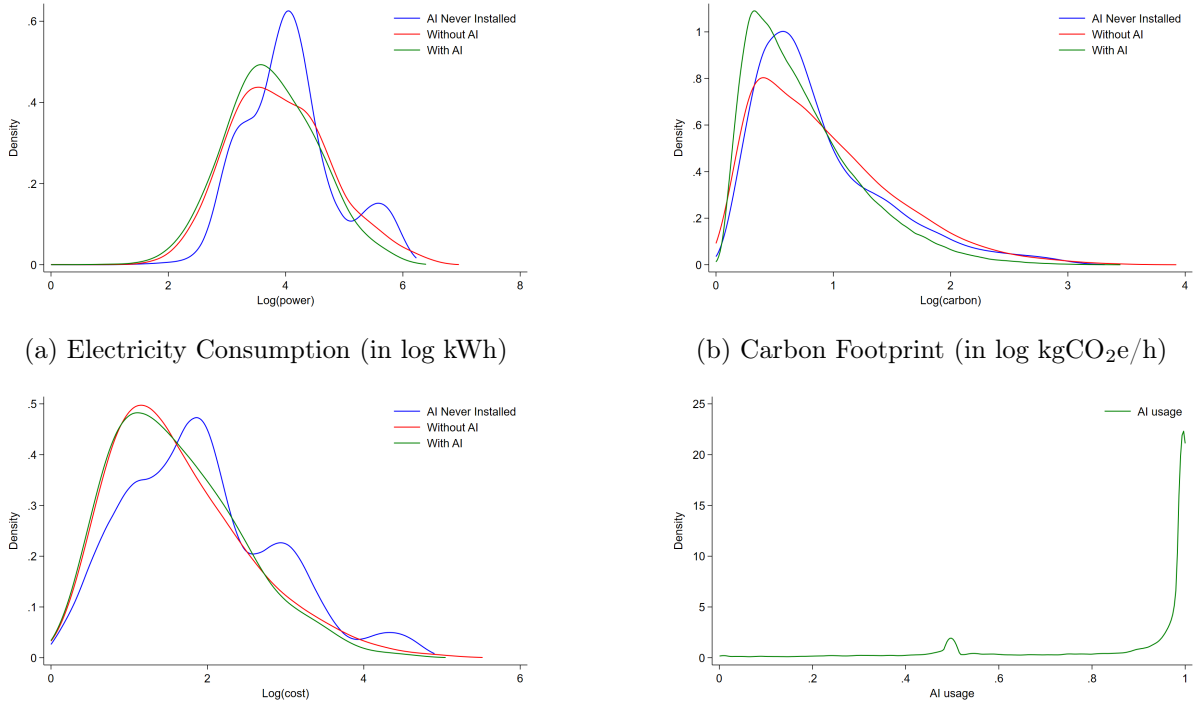


(a) Plant geographical location

(b) Plant nominal capacity

Figure 3: Plant Location and Nominal Capacity

Note: This figure shows the geographical and capacity distribution of plants within our sample. The plant nominal capacity are compared with the French national distribution in 2016 (CGDD, 2019).



(a) Electricity Consumption (in log kWh)

(b) Carbon Footprint (in log $kgCO_2e/h$)

(c) Electricity cost (in log €/h)

(d) AI control share of time

Figure 4: Distributions of the variables of interest

Notes. Observations for electricity consumption, carbon footprint and electricity costs are at the hour-level, and for AI control at the day-level. The first three variables are expressed in log. "AI Never Installed" refers to never treated plants; "Without AI" refers to observations from plants that have AI installed but where AI control is disabled (either before the first treatment date or after the first treatment date but where AI is not operating); "With AI" refers to observations from plants that will have AI installed when AI is operating.

Table 2: AI interruptions summary statistics

Failure Type	**Frequence**	**Percent**
Unauthorized	206,307	46.11
AI-team decision	107,419	24.01
Connection failures	133,677	29.88
Total	447,403	100.00

Notes. Observations are at the hourly level, restricted to periods of AI failures, after the deployment of AI. The first type of failures correspond to observations where the plant withdrew AI authorization. The second type of failures correspond to hours with any presence of P1 mode (AI team manually overriding the AI). Connection failures capture the residual exogeneous failures, attributed to connection outages.

5 Main Results: the net balance of AI use

5.1 Gross impact

Conceptual Framework. In Appendix Section F, we develop a stylized model to extend the scope of our empirical findings and clarify the mechanisms through which AI affects energy consumption in an industrial context. We consider an environment in which plants combine two factors: a generic input and an energy-intensive input. Their objective function consists in meeting a service target under operational uncertainty and with minimal costs. The required level of the energy input is stochastic, and plants face an asymmetric loss function: undershooting (an insufficient service level) incurs steep penalties due to non-compliance risk, while overshooting (excessive service) generates smaller linear losses as a consequence of wasted energy. AI adoption is modeled as an information improvement: it provides a more precise signal about operational requirements, formally characterized as a mean-preserving contraction of the predictive distribution of the required input. From this framework, we derive five testable implications, each consistent with our empirical results. First, AI generates energy savings on average by allowing operators to reduce the safety margin they hold against costly undershoots. Second, savings are larger when inflow rates are low, precisely when baseline operations are least efficient. Third, AI enables price responsiveness: optimal input usage decreases in the electricity price, and this sensitivity is amplified post-adoption as operators can fine-tune consumption more precisely. Fourth, AI reduces over-provisioning: operators require smaller safety margins to maintain compliance, compressing the right tail of the aeration distribution. This is confirmed in the data: examining the oxido-reduction potential, we document a significant reallocation of hours away from over-aeration toward more optimal levels, with no corresponding increase in under-aerated hours. Fifth, the model accommodates dynamic learning: as AI models improve over time, energy savings accumulate gradually with diminishing marginal gains.

Table 3: Effects on weekly electricity consumption, carbon footprint and costs

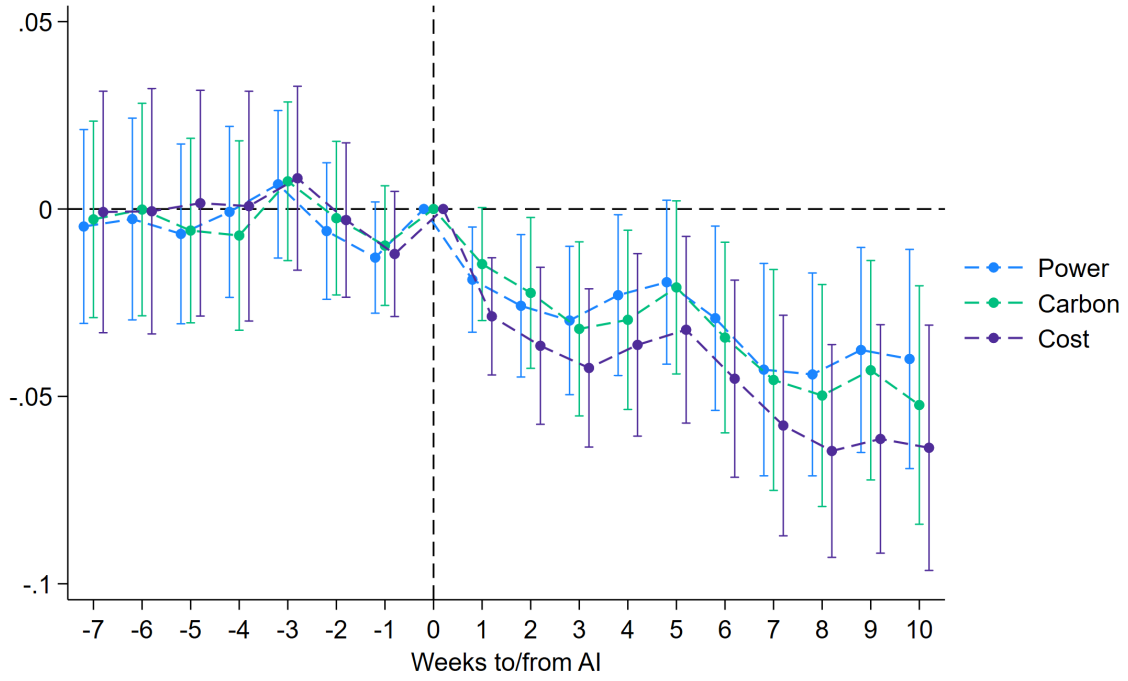
	AI usage	Treated	AI installed
Electricity Consumption			
	-0.054***	-0.035***	-0.030***
	(0.017)	(0.013)	(0.01)
Carbon footprint			
	-0.060***	-0.040***	-0.033***
	(0.019)	(0.012)	(0.012)
Electricity cost			
	-0.082***	-0.055***	-0.046***
	(0.017)	(0.012)	(0.011)
Observations	13,708	13,708	13,708
Number of plants	164	164	164

Notes: We report the average cumulative total effect measured over 10 weeks using de Chaisemartin and D'Haultfoeuille (2024). Electricity consumption, carbon footprint, and electricity cost are expressed in log values. The definition of the measure of AI use differs across columns. Effect on electricity cost is obtained controlling for the mean individual weekly electricity price. Standard errors are clustered at the plant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Overall Impacts. Table 3 examines the impact of AI implementation, with several measures of its implementation, on energy consumption, carbon footprint and electricity costs, using de Chaisemartin and D'Haultfoeuille (2024). The columns differ in the definition of the measure of AI use. In the column *AI usage*, the treatment variable is the share of the day during which the AI piloting mode is switched on. AI generates energy savings of 5.4 %, reduces the carbon footprint by 6 % and electricity costs by 8.2% on electricity cost assuming AI is switched on 100% of time. In the second column *Treated*, we present the observed effect when AI is switched on at least part of the day. On average, on days where it has been active at least once during the day, AI has generated electricity consumption savings of 3.5%, carbon footprint savings of of 4%, and electricity cost savings of 5.5%. Finally, in the third column, the dummy *AI installed* captures the effect of AI installation, irrespective of whether AI is actually switched on. AI installation has generated energy savings of 3%, a carbon footprint reduction of 3.3%, and a reduction in electricity costs of 4.6%. The weaker effect of AI installation is explained by the fact AI is switched off 42% of the time after its installation because of connectivity outages, or unauthorized AI control by the plant managers. Importantly, electricity consumption, electricity costs, and carbon footprint are measured at the plant level rather than at the aeration process level. As a result, when focusing solely on the aeration process, the efficiency gains are approximately twice as large in magnitude, since the aeration represents approximately 50% of the overall plant electricity consumption.

In Figure 5, the event-study graph suggests the absence of detectable pre-trends, followed by a gradual, statistically significant effect in the post-treatment period. The event-study graph shows, for each week relative to treatment, the average observed effect without weighting by the plant's actual intensity of AI use. The effect fully materializes after approximately 8 weeks (roughly two months). While Figure 5 uses log values, Appendix Figure E.6 also shows results with the absolute values of the effects, that we will use to assess the net energy balance of AI.

Figure 5: Event study, AI effect



Notes. This figure plots the coefficients and 95% confidence intervals from the event-study regression using de Chaisemartin and D'Haultfoeuille (2024). Dependent variables are in log terms. Effect on electricity cost is obtained controlling for the mean individual weekly electricity price.

Effect on Effluent Quality. We verify that AI monitoring and the reduction of electricity consumption do not come at the expense of effluent quality, especially in the removal of nitrogen compounds in treated water. We investigate how AI impacts the nitrification process, whose goal is to remove ammonium (NH_4^+) and convert it into nitrates (NO_3^-) through aeration of the effluent, and the denitrification process, whose goal is to remove nitrates and convert them into gaseous nitrogen (N_2), which then naturally escapes from the water to the atmosphere[17]. We also examine the effect on phosphate (PO_4^{3-}) concen-

[17]We provide a further description of the nitrogen removal process in Appendix Section J.1 and J.2.

trations. Unlike nitrogen removal, phosphate elimination is less dependent on aeration[18]. We rely on manual sampling of treated water for these pollutants. Appendix Table J.1 reports TWFE estimates using plant and date fixed effects. The results show that AI significantly reduces nitrate concentrations by 0.51 mg/L and ammonium concentrations by 0.29 mg/L, relative to baseline levels of 1.97 mg/L and 1.73 mg/L respectively. By contrast, we find no significant effect on phosphate concentrations, which suggests that AI-driven aeration management do not interfere with the phosphate elimination process.

Wastewater treatment plants are subject to pollutant discharge limits that vary across plants and over time due to local regulations. Using data on compliance with ammonium and phosphate thresholds, we estimate the effect of AI adoption on the likelihood of meeting these standards. Appendix Table J.2 reports the results from conditional logistic regression models including plant fixed effects. AI usage has a positive and statistically significant effect on ammonium compliance. The estimated coefficient implies that, within a given plant, AI usage increases the odds of meeting ammonium discharge limits by about 79%. We find no statistically significant effect of AI usage on phosphate compliance, consistent with our previous findings on phosphate concentrations.

We then focus on the heterogeneity of these results depending on the ammonium concentration. The marginal benefits of AI are concentrated in periods of high ammonium load, precisely when the biological process is most stressed and conventional rule-based aeration performs poorly. When influent ammonium is elevated, oxygen demand rises sharply; standard control logic typically responds by increasing aeration, which both (i) drives up electricity consumption and (ii) erodes the anoxic (i.e. oxygen-free) conditions required for denitrification[19]. By contrast, AI delivers a "triple dividend" in these high-load states, as we illustrate in Figure 6, by dynamically maintaining a good oxygenation balance[20] as a function of basin conditions and by anticipating load spikes. First, AI avoids over-aeration, shifting redox conditions toward denitrification: we observe a decline in oxidation-reduction potential at high ammonium deciles, which coincides with progressively larger reductions in effluent nitrate concentrations, while effects are negligible in the lowest deciles. Second, the reduced reliance on aeration translates into sizable electricity savings, concentrated in the upper six ammonium deciles. Third, despite lower aeration intensity on average, AI improves nitrification performance where it matters

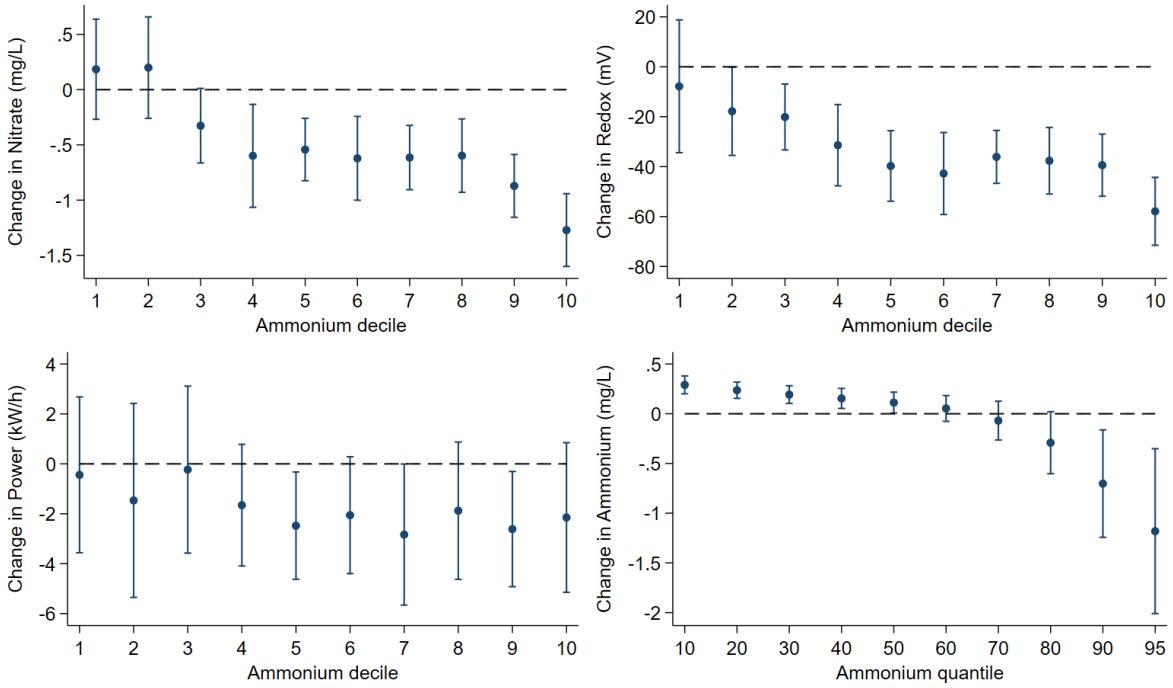
[18]It is mainly achieved through phosphorus uptake by polyphosphate-accumulating organisms (PAOs) and through chemical precipitation. But poor aeration management can nonetheless impair its removal (Deronzier and Choubert, 2004)

[19]In activated-sludge systems, excessive aeration raises dissolved oxygen in the basin; since heterotrophs preferentially use oxygen as an electron acceptor, denitrification is inhibited, leading to higher effluent nitrate.

[20]Also known as a anoxic/aerobic balance. We use "aerobic" to denote oxygenated conditions and "anoxic" to denote low-oxygen conditions in which nitrate (rather than oxygen) can serve as the terminal electron acceptor.

most: quantile regressions reported in Appendix Table J.3 show that AI adoption sharply reduces ammonium concentrations in the top decile, consistent with AI primarily mitigating over-limit episodes, while effects on lower quantiles are small and reflect a leftward shift of the overall distribution.

Figure 6: Effect of AI depending on ammonium concentration heterogeneity



Notes. The first three panels present TWFE estimates of the marginal effect of AI usage on nitrate (NO_3^-) concentrations (mg/L), redox potential (mV), and electricity consumption (kWh) across deciles of the NH_4^+ distribution, controlling for plants and date fixed effect, with standard errors clustered at the plant level. Results tables are in appendix J.4. The bottom-right figure shows quantile regression with plant fixed effects estimates of the effect of AI usage on NH_4^+ concentrations across the conditional distribution from the 10th to the 95th quantile. The result table for the quantile regression is in Appendix J.3. Confidence intervals are bootstrapped with standard errors clustered at the plant level.

Robustness. The causal interpretation of our results rests on a set of identifying assumptions that we probe through a battery of robustness checks, varying the estimation strategy, the treatment definition, and the composition of the sample.

1. Alternative estimators. We show similar results using alternative estimators robust to treatment effect heterogeneity. Appendix Figure G.1 presents event-study estimates from the estimator de Chaisemartin and D'Haultfoeuille (2024), alongside a standard TWFE and the estimator of Borusyak, Jaravel and Spiess (2024). We estimate here the effect of the AI installation which constitutes an absorbing treatment. All three specifications present a reduction in electricity consumption of approximately 2 to 4% in the weeks

following installation, with no evidence of pre-trends in the six weeks preceding treatment.

2. Restriction to exogenous failures. To address the potential downward bias in our estimates of the AI effect caused by endogenous failures, we restrict our sample to weeks unaffected by endogenous interruptions. Specifically, we drop all post-installation weeks where AI management was authorized less than 80% of the time, and all weeks where the AI team switched to manual control (P1 mode) more than 20% of the time. This last filter allows us to flag deliberate decisions by the AI team to disable AI rather than a simple connectivity outage. We don't apply this filter in the first two weeks post AI installation as manual interventions in this period are part of the installation process and are not deliberate decisions to disable the AI. In total, we remove 3,383 weeks, representing 36% of post-installation observations and 24% of the overall sample. The results of the AI effect, reported in Appendix Table G.1, are almost not affected by this restriction: we estimate reductions of 4.4%, 4.9%, and 6.9% in energy consumption, carbon footprint, and electricity cost, respectively, when AI is activated 100% of the time. The effect of AI installation remains statistically significant and broadly consistent with our baseline estimates.

3. Rebound effect. Since the reduction in electricity demand is limited to periods when AI is switched on, we might wonder whether AI induced energy efficiency gains are compensated by a rebound effect when AI is switched off. Two causes might drive this rebound out of AI-use time: first, a behavioral attitude of the plant management (for example in the situation he is less careful, provide less guidelines because the plant is AI managed); second, a reallocation effect between hours. We test for the rebound effect hypothesis by regressing electricity consumption, carbon footprint and electricity costs in a TWFE specification with both a dummy for AI being installed and piloting time variables. Should the rebound hypothesis be verified, we would observe an increase in consumption when the plant is AI-installed once controlled for energy gains induced by AI piloting time. In Appendix Table G.3, we find only limited evidence in defense of a rebound effect. If any rebound effect (out of AI operations), it is: (i) limited to electricity consumption and carbon footprint, and does not affect electricity costs and (ii) extremely limited in comparison to the gains (3 to 4 times smaller in magnitude of the coefficient).

4. Effect across AI adoption wave. We compute the separate effectiveness of AI usage across the two adoption waves. Appendix Table G.2 reports the estimates of AI usage on electricity consumption for each wave separately, using never-treated plants as controls in both subsamples. The average total effect for the second wave is -4.7%, consistent with our main results. For the first wave, the estimated effect is -6.0% but does not reach the significance levels. In both cases, the joint placebo test fails to reject the null hypothesis of

parallel trends. The consistency of the estimates across waves suggests that the baseline results are not driven by a particular adoption cohort.

5. Additional robustness checks. We check that our results are not sensitive to several other changes in specifications. In Appendix Table G.4, we run a Diff-in-diff specification with the addition of an inflow control; the result is unaffected in significance and magnitude. On the contrary, when we add redox as a control, the result remains significant, but the magnitude is smaller; this is in itself consistent with the important endogenous changes in redox values due to AI implementation. Table present the fraction of failures of AI, showing that a large share of failure may be due to unauthorized pilotage. Otherwise said, comparing the impact of AI on two same redox values is meaningless given the redox value is affected in itself by AI: as we documented it, part of the effect is driven by an increase of the low-ORP hours which are less energy-consuming. In complement, since we observe that the never-treated units had different characteristics in Section 4, in Appendix Table G.5, we check that the restriction of the sample to non-never treated units does not affect our results.

Heterogeneity effects. We examine the heterogeneity of AI effects according to plant size, water inflow levels, and water oxidation levels using hourly data[21].

1. Plant size and flow amplitude. We examine size heterogeneity along two dimensions of scale: the size of the plant, measured by nominal treatment capacity, and the effective operating scale, measured by received water inflows. We find no evidence that the effectiveness of AI varies systematically with the nominal capacity or the average inflow levels. As shown in Appendix Figure H.1, treatment effects are broadly flat across the capacity distribution, suggesting that AI does not deliver larger gains (in %) in larger installations. By contrast, we observe substantial heterogeneity with respect to absolute inflow levels. When we interact AI use with inflow (measured in logarithm) in specifications analogous to our baseline TWFE model, we find that the energy-saving effect of AI reaches 12.8 and 13.6 percent for electricity consumption and environmental footprint respectively at very low inflow levels (Appendix Table H.1). We also find that energy intensity (measured by the electricity consumption per unit of inflow) decreases sharply as inflow levels rise (Appendix Table H.2), indicating that the water treatment process exhibits strong economies of scale and operates more energy-efficiently at higher flow rates. Appendix Figure H.2 also shows that the energy-saving impact of AI diminishes at higher inflow levels, and that AI efficiency gains are most pronounced at low inflow. AI reaches allows to improve energy efficiency up to 19% at very low levels of inflows, and this effect becomes less strong at higher inflow levels. Taken together, these results suggest that AI effectively

[21] The piloting time is only available at the day-level, so we assume a fixed level of piloting time across hours.

mitigates the energy inefficiencies that occur at low-operating regimes.

2. The Oxidation - Reduction Potential (ORP). A key concern is whether AI saves energy by genuinely improving efficiency or simply by lowering the intensity of the treatment task. To rule out this latest possibility, we examine ORP (redox), a standard indicator of oxygenation conditions and treatment quality. Checking ORP permits to rule out the possibility that AI reduces electricity simply by relaxing treatment (e.g. operating under more reducing conditions that might compromise nitrification or effluent standards). Following the consensus in the technical literature on water treatment, we classify redox values into three standard operating ranges: very low (below -200 mV), moderately anoxic (-200 to 0 mV), and oxidizing (above 0 mV)[22] When we compare the distribution of ORP values by AI treatment status (Appendix Figure D.9), plants with AI spend much more time within this optimal range. This shows that AI does not reduce electricity by relaxing treatment standards. Instead, AI improves process control by keeping the system more consistently in the biologically efficient operating window, achieving both better compliance and lower energy use. We interact the bins with AI use in a TWFE specification (Appendix Figure I.2). The results show strong heterogeneity across redox regimes. AI delivers the largest gains when the biological process is close to binding or unstable conditions. At very low redox levels (below -200 mV), where small timing errors in aeration can generate large energy costs or quality risks, AI substantially improves operational efficiency, yielding energy savings of approximately 20 p.p. In moderately anoxic conditions (-200 mV to 0 mV), where the process is already relatively well controlled, AI delivers more modest but still significant gains (-7p.p). By contrast, when redox values are positive, indicating an already oxidizing environment, there is little room left for optimization, and we find no significant effect.

5.2 Net environmental balance

What is Green AI? Assessing AI's environmental balance, that is to say jointly evaluating its operational benefits and its own footprint, is both central and technically challenging, and lies at the core of the notion of Green AI (Rojahn and Grum, 2025). Following the canonical definition proposed by Schwartz et al. (2020) and widely adopted in the computer science literature (Xu et al., 2021), we define Green AI as the ability "to obtain novel results without increasing computational cost, or ideally reducing it." While this perspective is widely viewed as promising, empirical evidence remains scarce on whether AI delivers net environmental gains once the energy used to deploy and op-

[22]The water treatment literature identifies a target ORP range between -200 and 0 mV (Luccarini et al., 2017; Cheng, Kumar and Lin, 2012). In particular, values below -200 mV risk poor nitrification (Charpentier et al., 1998; Myers, Myers and Okey, 2006). Conversely, ORP values above 0 mV suggest insufficient anoxic conditions for denitrification. Following discussions with the Purecontrol engineers we exclude data below -450 mV and above 450 mV that are mismeasurements.

erate AI systems is taken into account. In this section, we take a first step toward such a net accounting by (i) bounding the direct electricity consumption of the AI infrastructure used to operate the control system at scale, and (ii) combining these measurements with conservative life-cycle estimates for the broader AI system.

Direct electricity consumption from AI use: metered upper bounds. A key advantage of our setting is that Purecontrol operates its computing infrastructure internally rather than relying exclusively on third-party cloud providers. This allows us to directly measure the electricity use of the AI system using internal server monitoring data. All AI workloads (production, development, and testing) run on a set of virtual machines hosted on a limited number of physical computer servers. Because these services can migrate across servers over time, we adopt a deliberately conservative strategy: we construct an upper bound at the level of physical servers. Using internal monitoring data (Grafana), we (i) map virtual AI services to the physical servers that host them, and (ii) sum conservative server level caps based on one month of observed power consumption. Since the mapping changes over time, this bound is intentionally coarse but safely conservative.

Using computer server monitoring data over a one-month window, we observe server-level peak power draws ranging from roughly 0.3-0.4 kW for the smaller production nodes up to about 1.5 kW for the largest node. Electricity consumption is stable over time. To obtain an upper bound robust to short-lived peaks and measurement noise, we apply a conservative uplift (peak +10%) and an additional safety margin at the node level, and then sum across all nodes used to host AI-related services across environments. This yields a total upper bound of approximately

$$P_{\max}^{\mathrm{AI}} \approx 5.6 \text{ kW}$$

for the entire compute equipment supporting all its customers in the water treatment sector.

Because peak load is an extreme scenario, we also compute a more representative bound using the observed node-level average power draw over the same monitoring window, again applying a conservative margin. This yields an upper bound of about

$$P_{\mathrm{avg}}^{\mathrm{AI}} \approx 4.0 \text{ kW}.$$

Assuming continuous operation, we convert these power bounds into annual electricity use by multiplying by the 8,760 hours in a year. Applying this to the conservative average and peak power estimates yields:

$$E^{\mathrm{AI}} \in [P_{\mathrm{avg}}^{\mathrm{AI}}, P_{\max}^{\mathrm{AI}}] \times 8760 \quad \Rightarrow \quad E^{\mathrm{AI}} \in [35, 49] \text{ MWh/year}.$$

Crucially, this computing infrastructure infrastructure serves roughly 500 plants, of which about half are Veolia sites. When we distribute the entire electricity use of this computing assembly across plants, the implied order of magnitude is on the order of 70–100 kWh per plant per year. Even if we conservatively attribute the entire electricity use to the Veolia fleet, the implied figure remains below twice this amount. When benchmarked against the operational electricity reductions documented in the previous subsection, the direct electricity used to run the AI infrastructure is well below 1% of the electricity savings generated by AI-driven optimization. We estimate the energy saving of AI in absolute values to be around 2600 kWh/month on average across plants, which corresponds to 31,2 MWh/year per plant on average (see Appendix Table E.6.)

Lifecycle emissions of AI: a conservative accounting framework. Direct electricity is only one component of AI's footprint. A broader accounting must also include upstream emissions associated with deployment and maintenance (e.g., additional on-site hardware, software development, and deployments). Producing a full life-cycle assessment (LCA) that matches standardized protocols is beyond the scope of this paper; however, we can build a transparent and conservative estimate based on a detailed internal assessment of the AI product life cycle.

The internal LCA framework defines the functional unit as the continuous optimization of a site over one year and decomposes impacts into (i) incremental hardware on the client site when required (connectivity, cybersecurity, metering), (ii) software development and platform infrastructure, (iii) deployment (logistics and potential on-site interventions), and (iv) operational use (data collection, storage, model execution, and monitoring). It does not quantify end-of-life impacts and explicitly excludes emissions from office buildings and employee commuting, which are second-order for the present application.

On the hardware side, the incremental equipment required for deployment is highly heterogeneous: in the best case, existing infrastructure can be reused and incremental embodied emissions are close to zero; in the most equipment-intensive cases, the sum of embodied emissions associated with routers, firewalls, gateways, local servers, and additional probes can reach several hundred $kgCO_2e$ (with an upper range around 750 $kgCO_2e$ for the most equipment-heavy configurations, before annualization). To remain conservative, we consider scenarios in which additional on-site hardware and specialized interventions are required.

On the "use" side, the internal assessment provides an order-of-magnitude decomposition of annual digital energy use per site: data collection and storage (tens of kWh/year), daily AI characterization workloads (roughly 100–150 kWh/year), and a real-time control loop (roughly 270–300 kWh/year), for a total of about 400–500 kWh per site-year. Using a French electricity emission factor of 0.08–0.1 $kgCO_2e$/kWh, this corresponds to about 30–40 $kgCO_2e$ per site-year for the operational (electricity) component.

Putting these components together, and excluding raw material extraction and end-of-life stages (while allowing for conservative assumptions on deployment requirements), the resulting estimate of life-cycle emissions in the conservative case scenario when we attribute hardware changes to the AI system is approximately

$$\text{LCA}^{\text{AI}} \approx 250\text{–}400 \text{ kgCO}_2\text{e per plant-year}$$

in the conservative scenario we consider[23]. Based on the results from the previous section, we estimate average annual carbon savings of approximately 550 kg CO_2e per plant, indicating that the net environmental impact of AI adoption remains positive. It is important to emphasize that these energy savings are observed in plants located in France, where the electricity mix is among the least carbon-intensive globally. Consequently, the associated reductions in carbon emissions are likely to be substantially larger in countries with more carbon-intensive electricity systems.

Implications for the net environmental balance. Taken together, the monitoring-based upper bounds on direct electricity use and the conservative LCA scenario suggest that the environmental footprint of the AI system is small relative to the operational savings it generates in our setting. In particular, our metered upper bound implies that the electricity required to operate the AI compute stack represents less than 1% of the electricity saved through process optimization. Moreover, even under conservative life-cycle assumptions, back-of-the-envelope comparisons indicate that net emissions savings remain positive once AI-related emissions are accounted for.

6 Additional Dynamic Margins of AI control

6.1 AI cost-effectiveness and intertemporal reallocation effects

Effects between peak and off-peak hours. To further examine how AI impacts electricity costs, we explore the evolution of the intertemporal elasticity of electricity consumption. We follow a methodology similar to Jessoe and Rapson (2014) and Fabra et al. (2021), by measuring the response of consumption to prices. Unlike their approach, which relied on instrumental variables (IV), our setting does not require such a technique: given the extremely micro industrial context, individual plants have no influence on electricity prices, making endogeneity concerns minimal. We examine how the intertemporal elasticity of electricity changes with AI adoption by analyzing the interaction between price and treatment. However, relative to our concerns, the price signal is extremely noisy, because

[23]We thank the Purecontrol teams for sharing their internal assessment used to construct this conservative scenario. Discussions with Veolia and Purecontrol indicate that equipment-heavy deployments and repeated on-site interventions were seldom necessary in practice, and that plants were selected *ex ante* to minimize such requirements.

it is affected by weekly and seasonal shifts, while the intertemporal effects, if any, only occurs at the infra-day level[24]; otherwise said, the structure of electricity pricing exhibits limited intra-day variations but important heterogeneity across days: prices are generally lower on Sudays, and there are substantial seasonal patterns. To focus on intra-day price variation with daily data, we create a binary variable indicating whether the price is high relative to its within-day distribution (above the median), and zero otherwise, as most prices are heavily polarized across days. This variable is consistent with our empirical observation (most plants have only two different prices per day). An additional concern is that AI piloting data is only available at the daily level; we therefore focus on highly treated plants.

We find that reallocation is a critical mechanism driving the pattern of electricity consumption. Activating AI does not lead to a uniform reduction in electricity consumption across hours. Demand remains relatively stable during off-peak hours, whereas demand sharply decreases during peak hours. Figure 7 illustrates the effect of AI piloting share, estimated separately with de Chaisemartin and D'Haultfoeuille (2024) for peak and off-peak hours[25]. We find the overall reduction in electricity consumption to be entirely driven by a substantial decrease during peak hours. Average weekly electricity consumption during peak hours declines by 7.4%, whereas consumption during off-peak hours remains essentially unchanged (-1.7%, n.s.), resulting in an average reduction of 5.4% when considering all hours combined. In complement, in a TWFE specification, in Appendix Figure K.3, we find that the effect is substantial enough for the price elasticity to reverse: while electricity consumption was 7 percentage points higher during peak hours before AI adoption, it becomes 3 percentage points lower during off-peak hours after AI implementation. These findings show that the AI in industry permit indeed to reallocate electricity demand from peak toward off-peak hours, consistently with previous findings on households energy demand (Metcalfe et al., 2026).

To further validate our findings and explore the mechanisms driving these shifts, we compare the relative electricity consumption patterns across different hours of the day. In Figure 8, we compare hourly consumption profiles when controlled and not controlled by AI[26]. While this analysis is descriptive in nature, it provides valuable insights into how electricity consumption evolves throughout the day. Our results clearly show that the differentiated electricity consumption patterns during peak and off-peak hours are driven by a reallocation mechanism. Notably, the reduction in consumption across these periods

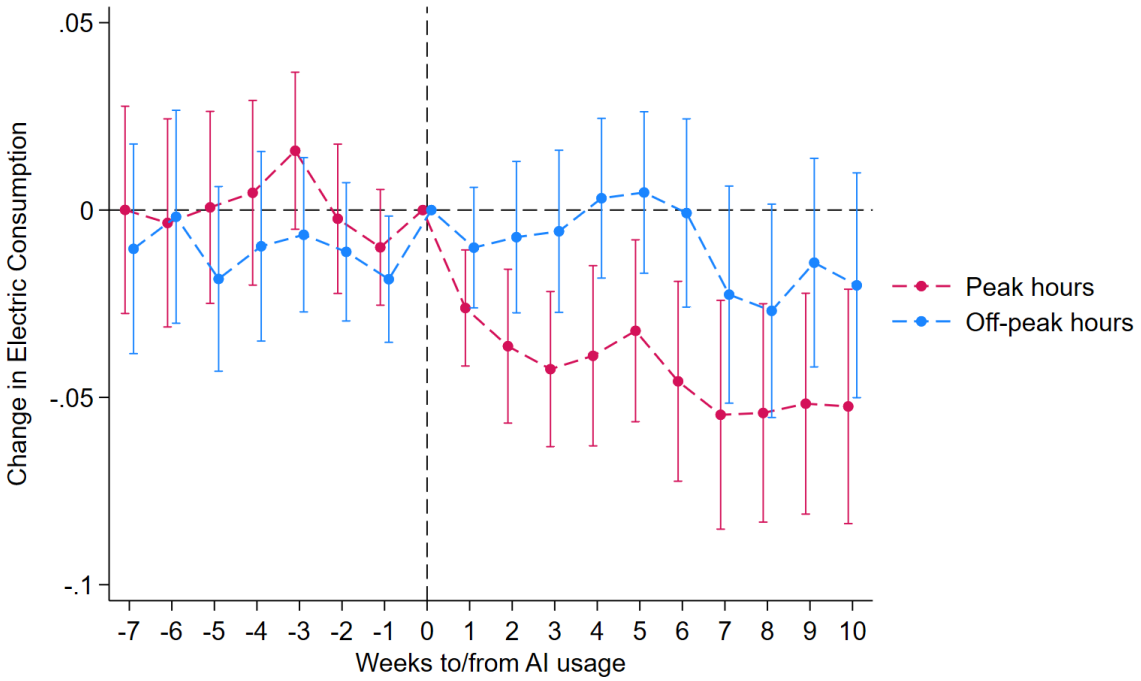
[24]Maintaining safety margins and operational constraints typically restrict the ability of plants to shift consumption across days. The reallocation might only occur across a couple of hours

[25]The division between peak and off-peak hours is generally consistent across days and seasons for a given plant, which justifies the relevance of this specification.

[26]We compute a weighted average using weights equal to the inverse of the number of observations per id and hour, so that each id contributes equally to the hourly mean.

Figure 7: Power savings between peak and off-peak hours

(a) Event study



(b) Average cumulative total effect per plant

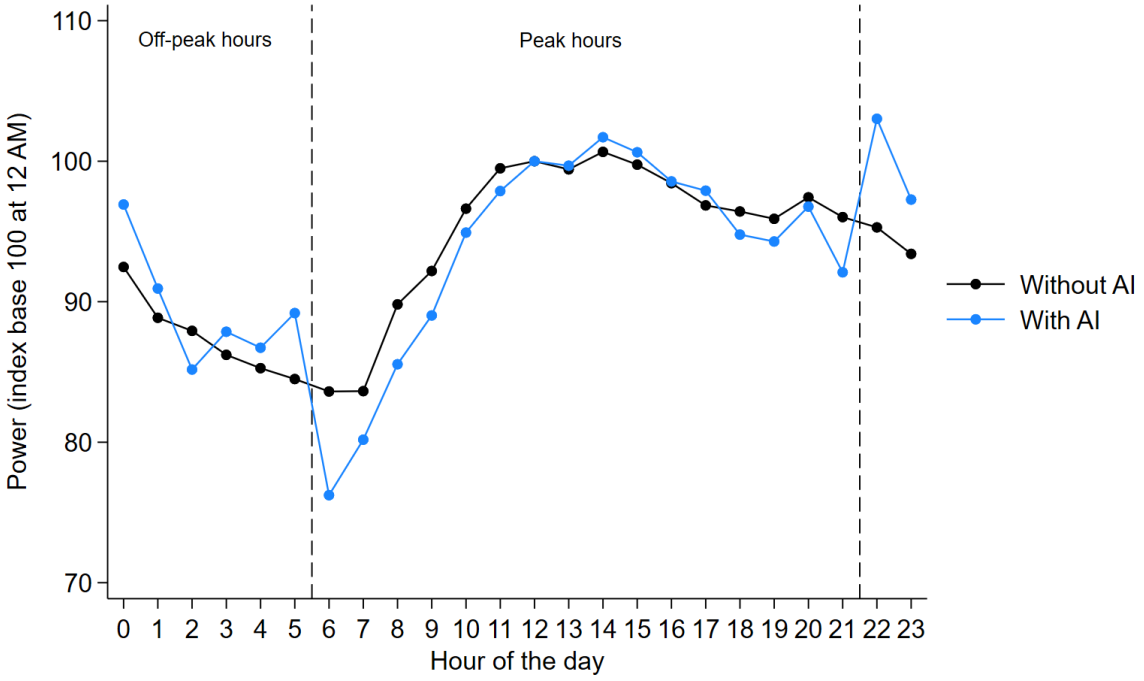
	Estimate	Std. Error	95% CI Lower	95% CI Upper
All hours	-0.054***	0.017	-0.087	-0.021
Off-peak hours	-0.017	0.017	-0.051	0.016
Peak hours	-0.074***	0.017	-0.108	-0.04

Notes. This figure plots the coefficients and 95% confidence intervals from the event-study regression using de Chaisemartin and D'Haultfoeuille (2024). Subfigure (a) separates the sample between off-peak and peak hours and aggregates outcomes at the weekly level. Subfigure (b) reports the average cumulative total effect measured over 10 weeks. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

is not uniform: we observe two clear kink points at the beginning and end of the peak hours. At the onset of the peak period, there is a sharp 20% reduction in electricity consumption, followed by a 10% increase towards the end of the peak period. Furthermore, we find that electricity consumption during the beginning and end of the peak period is lower in AI-managed plants compared to the counterfactual without AI, while the reverse holds true for the off-peak period. These pronounced discontinuities at the start and end of the peak period suggest that the price signal effectively drives behavior, allowing the system to anticipate and shift consumption ahead of the peak and defer it to the end within a few hours of the threshold. On the contrary, during the middle of the peak period (from 12 p.m. to 5 p.m.), consumption remains largely unaffected by the reallocation,

indicating that the mechanism operates primarily around the edges of the peak window.

Figure 8: Daily relative consumption profiles with and without AI



Notes. This figure compares the weighted average of the electricity consumption profiles where each plant receive an equal weight per hour, between treated days ("With AI") and non-treated days ("Without AI"). We focus on relative consumption levels: for each group, consumption is normalized to 100 at 12 a.m. The figure also highlights peak and off-peak hours, which are common across plants and do not vary over the year.

In theory this pattern is likely to generate both reduction in electricity costs and in environmental footprint, because the carbon intensity of electricity demand in generally stronger in peak periods. However, as we show in Appendix Figure K.1, the peak / off-peak shift is not discriminative enough: in our sample, the average carbon intensity during peak hours is only about 5% higher than during off-peak hours. It explains why we cannot, counterintuitively, detect a stronger effect of AI use on carbon footprint than on electricity consumption in Figure 3, despite the reallocation effect. This result advocates for a reinforcement of the price signal sensitivity through a more comprehensive consideration of supply-side heterogeneity and a shift towards more dynamic, real-time pricing schemes for industrial electricity consumers. Indeed, a more dynamic pricing system, reflecting intra-day variability in carbon intensity (that we see in Appendix Figure K.2) and electricity demand, could amplify AI's benefits and encourage industries to further optimize their energy usage patterns in alignment with both environmental goals and economic incentives. Hence, although price reallocation is not a direct lever for reducing carbon footprints in this context, it could become a significant factor in a grid for which the

carbon intensity of electricity supply varies across hours, coupled with a pricing scheme that reflects this variability. It highlights the potential of AI not only as an operational tool for energy efficiency but also as a mechanism for driving policy changes in electricity pricing.

6.2 AI in Periods of Weather Stress

Extreme precipitation is a first-order operational stressor for wastewater treatment plants (WWTPs) in our setting: rainfall episodes are strongly correlated with short-run movements in electricity use, consistent with the idea that inflow surges and hydraulic constraints mechanically translate into higher pumping and process demand. This motivates the hypothesis that the value of AI-based control is concentrated in extreme-weather episodes, when conventional operations face binding constraints and electricity use tends to spike.

Empirical specification and precipitation definitions. To study whether AI attenuates the operational and energy consequences of extreme precipitation, we estimate:

$$\log(Y_{i,d}) \;=\; \beta D_{i,d} + \gamma R_{i,d} + \delta\left(D_{i,d} \times R_{i,d}\right) + \mu_i + \lambda_d + \varepsilon_{i,d}, \tag{3}$$

where $Y_{i,d}$ denotes electricity consumption of plant i on day d, $D_{i,d}$ captures AI operation (we consider alternative measures, including a binary treatment indicator, a higher-intensity treatment definition, and a continuous "pilotage share" measure), and μ_i and λ_d are plant and day fixed effects. Because the dependent variable is in logs, γ and δ can be interpreted as semi-elasticities.

We operationalize "extreme precipitation" using two complementary indicators. The first one captures hourly precipitation volume:

$$R_{i,h}^{\text{hourly}} = \mathbb{1}\{\text{Rain}_{i,h} > \tau\},$$

where $\text{Rain}_{i,h}$ is hourly rainfall. The second one captures the overall effect of rainfall episode at the day-level:

$$R_{i,d}^{\max} = \mathbb{1}\left\{\max_h(\text{Rain}_{i,h}) > \tau\right\}.$$

We sweep the threshold τ over a range of high cutoffs in the right tail of the empirical distribution, plotted in the Appendix Section D.8.

The role of precipitation intensity and concentration. Using the hourly precipitation measure $R_{i,h}^{\text{hourly}}$, Panel (a) Figure 9 logically shows that exceedances of high rainfall thresholds are associated with sizeable contemporaneous increases in WWTP electricity

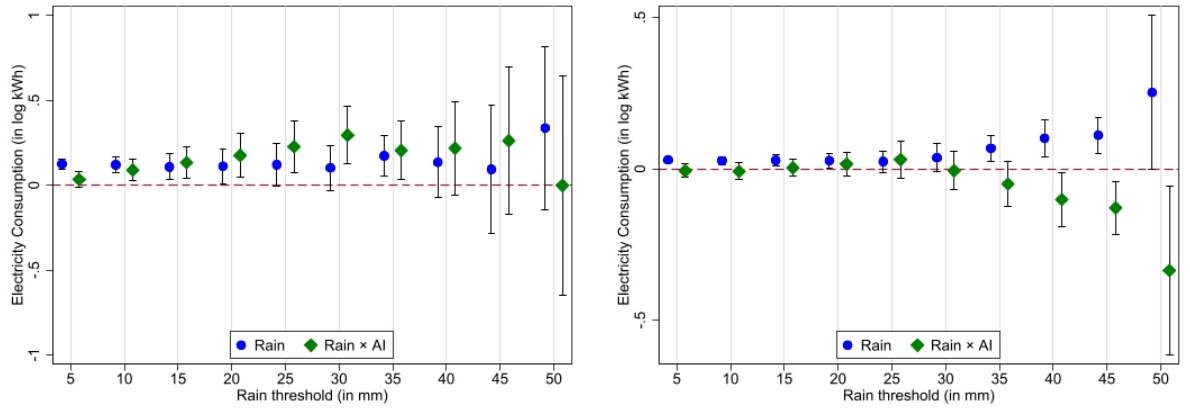
consumption (a positive γ). More importantly to our topic, we estimate a positive interaction with AI operation (a positive δ), and the magnitude of this interaction varies systematically with shock intensity: the incremental response under AI is negligible at low thresholds, becomes statistically significant at the 10 mm threshold, and rises gradually up to 30 mm. This pattern suggests that AI leads plants to react more sharply and more proportionally at the time the shock arrives. In other words, during the hours in which extreme rainfall occurs, AI does not reduce electricity use; instead, it reallocates effort toward real-time adjustment, increasing consumption in a targeted manner when inflows spike.

Panel (b) of Figure 9 examines the same episodes at the daily level using $R_{i,d}^{\max}$. Up to 30 mm, we detect little effect on daily consumption. Starting around 35 mm, however, rain shocks generate markedly larger increases in daily electricity demand under conventional (non-AI) operation, consistent with acute inflow surges and binding capacity constraints. For these high-intensity events, the rain-induced increase is of similar order of magnitude at the hourly and daily horizons, indicating limited scope for intertemporal smoothing under rule-based automation.[27] Crucially, these are precisely the episodes for which AI delivers its largest mitigation effects. The interaction term δ is of comparable magnitude and opposite sign to γ, implying that the incremental daily electricity consumption associated with extreme rainfall is largely offset under AI operation. Across thresholds, the evidence is consistent with near-complete compensation: while extreme precipitation mechanically raises electricity demand under conventional control, AI brings daily consumption back toward the counterfactual level observed on non-extreme days. Appendix Figure L.1

The attenuation therefore does not appear to operate by suppressing consumption at the peak hour of the incident. Instead, the mitigation shows up when aggregating over the day. This pattern is consistent with AI acting as a supervisory control that reallocates effort intertemporally: it may respond aggressively (and energy-intensively) at the time of the shock to maintain process stability or compliance, but then prevents a prolonged elevation in electricity use over subsequent hours, delivering almost-normalized daily totals. Appendix Figure L.1 corroborates this interpretation: while both curves surge sharply at rain onset (t = 0), the AI-equipped plants reach far more rapidly a higher instantaneous peak, yet their index also falls back below baseline faster thereafter. Non-AI plants, by contrast, sustain a persistently elevated plateau throughout the post-event window. In other words, the main benefit of AI in extreme precipitation episodes appears to be improved within-day management of the operational shock, smoothing the daily electricity profile rather than reducing instantaneous consumption at the moment of maximum stress.

[27] This finding is consistent with our assumption that, absent AI, plants adjust sluggishly to sudden inflow surges.

Figure 9: Intense rainy episodes and their attenuation through AI use.



Thr.	5	10	15	20	25	30	35	40	45	50
IPW	10^2	10^3	10^3	10^4	10^4	10^4	10^4	10^5	10^5	10^5

(a) At the hour-level

Thr.	5	10	15	20	25	30	35	40	45	50
IPW	10^1	10^2	10^2	10^2	10^3	10^3	10^3	10^3	10^3	10^4

(b) At the day-level

Notes. Panel (a) reports estimated effects using the hourly precipitation measure $R_{i,h}^{\text{hourly}}$. The first series corresponds to the effect of exceeding the rainfall threshold τ, while the second series reports the interaction between threshold exceedance and AI use (scaled by the AI-use intensity). Error bars display 5% confidence intervals. Panel (b) reports the corresponding estimates at the daily level using the maximum daily precipitation measure $R_{i,d}^{\max}$. Markers and confidence intervals are defined as in panel (a). The table below the panels reports the inverse probability weights (IPW), i.e., the inverse of the frequency at which an observation in the sample exceeds each rainfall threshold.

Appendix Figure L.2 repeats the analysis using total daily precipitation, a daily analogue of $R_{i,h}^{\text{hourly}}$. The qualitative pattern is similar but estimates are less precise, with statistically significant effects only for very large daily totals (above 80 mm/day). Taken together, these results point to concentration as a key state variable: AI appears particularly effective at mitigating highly concentrated rainfall episodes, precisely the situations in which simple threshold-based automation is least suited to regulate the process.

Other meteorological conditions. We further examine whether AI adoption differentially moderates the effect of other meteorological conditions on electricity consumption. Using the same empirical strategy, we test for extreme low minimum temperatures, high maximum temperatures, extreme humidity, and water pressure. Across all specifications, we find no statistically significant interaction effects. Although these meteorological variables are correlated with electricity consumption at the plant level, the magnitude of these relationships is substantially smaller than for intense rainfall episodes. As a result, any moderating effect of AI, if present, is likely too limited for the interaction terms to be estimated with sufficient statistical power.

7 Welfare and Climate Significance

In order to gauge the potential importance of our findings for global climate and welfare, we integrate AI-induced energy efficiency improvements in the DICE climate-economic model (Barrage and Nordhaus (2024)). We consider three scenarios which respectively apply our estimated treatment effects to: (1) global wastewater treatment electricity consumption ("WWTP only"), (2) wastewater treatment plus energy demand from industrial drying and steel galvanization - two other processes for which Purecontrol has already developed similar solutions - ("WWTP+"), and (3) all industrial energy and CO_2 emissions ("All"). Extrapolating beyond wastewater treatment is motivated by the fact that there are many industrial *endpoint control* processes where an unobservable target must be met and where digital twins could help improve energy efficiency over fixed-time application or rules of thumb (e.g., industrial drying, cooling, fermentation, certain metallurgical processes like galvanization, etc.).

We formally consider a permanent level shift in the baseline industrial CO_2 intensity of production in the DICE-2023 model (base scenario) of -6% applied to the relevant emissions shares, which are (1) 0.4% of emissions for "WWTP only",[28] (2) 5% of emissions for "WWTP+",[29] and (3) 100% of energy-related emissions for "All."

Table 4 summarizes the results. The estimated global welfare gains associated with our estimate of AI-induced CO_2 reductions are (1) $75 billion ($2025 PPP) in the "WWTP only" scenario, (2) $940 billion in the "WWTP+" scenario, and (3) $18.8 trillion in the "All" scenario. We note that these estimates are lower bounds as they reflect only the value of avoided CO_2 emissions and not of the associated energy savings. Despite these conservative assumptions, the modeling results show large global stakes from AI-induced energy efficiency improvements.

Table 4: Global Welfare Gains from AI-Induced Energy Efficiency Improvements

Processes Covered	% Emissions	Welfare Gain (2025 PPP)
WWTP	0.4%	$75 billion
WWTP + drying + galvanization	5%	$940 billion
All	100%	$18.8 trillion

[28] We combine IEA (2016)'s estimate that WWTP account for 1% of global electricity consumption with the global average CO_2 intensity of electricity (0.0005 GtCO2/TWh, calculated from International Energy Agency (2025) and International Energy Agency (2016)) and divide by total industrial CO_2 emissions (International Energy Agency (2016)) to arrive at 0.4%.

[29] Drying is frequently estimated to account for 10-25% of industrial energy demand (e.g., da Silva et al. (2025); Bayati et al. (2025); Mujumdar (2007)). Combining a central value of 15% with a 30% estimated share of total energy demand from industry (IEA (2024)) yields an energy emissions share of 4.5% for drying. For galvanization, we combine a CO_2 intensity of 0.2 kg CO_2 per kg of steel from Álvarez-Álvarez et al. (2025) with a lower bound on global galvanized steel production from WSA (2025) to arrive at an estimated global emissions share of 0.1%.

8 Conclusion

This paper studies whether industrial AI can deliver net environmental benefits once its own footprint is accounted for. Using high-frequency operational data from a large fleet of wastewater treatment plants and quasi-experimental variation from staggered adoption and non-absorbing exposure to AI control, we estimate the causal impact of an AI-driven digital twin on electricity consumption and CO_2 emissions. We find economically meaningful reductions in both outcomes: 5.4% and 6% in the first ten weeks of AI operation for electricity consumption and carbon footprint respectively, and 8.2% for electricity costs. Overall, our findings provide micro-level evidence consistent with the view that AI-enabled process optimization can generate substantial emissions reductions in industry (IEA, 2025).

AI enhances process efficiency in aeration, one of the most energy-intensive stages of wastewater treatment, by leveraging a richer real-time information set than rule-based control and by dynamically adjusting operating setpoints. The evidence points to several distinct channels through which AI improves energy efficiency. First, a large share of the gains stems from responding to short-run heterogeneity in inflows: AI reduces hours of over-aeration and scales back activity when inflow levels are low. Second, AI limits the electricity surges triggered by extreme and concentrated rainfall events, adjusting more flexibly to abrupt inflow spikes than conventional automation rules allow. Third, AI improves the management of chemical pollution peaks. When influent ammonium load is elevated, conventional rule-based aeration tends to over-aerate, simultaneously driving up electricity consumption and eroding the anoxic conditions required for denitrification. AI avoids this trade-off by dynamically maintaining the aerobic/anoxic balance and anticipating load spikes, delivering simultaneous gains in energy efficiency and effluent quality precisely when the biological process is most stressed. Taken together, these three channels suggest that AI-based control is particularly valuable under operational stress, when the limitations of fixed automation rules are the most binding. In addition, AI reshapes the intraday profile of electricity use. Reductions are concentrated during peak hours, which helps explain why electricity costs fall by more than consumption and carbon emissions. Under efficient grid pricing, this load-shifting capacity could constitute an additional lever for carbon emissions reduction.

Beyond the mechanisms at play, a central contribution of the paper is to move beyond indirect proxies and provide a net accounting of Green AI's environmental implications in an industrial setting. Using provider-side metering data and conservative life-cycle assumptions, we bound the electricity use and associated emissions attributable to operating the AI system, and compare them to the operational savings achieved at plants. The AI system's own electricity consumption is small relative to the downstream reductions

it induces, implying positive net emissions savings even under conservative assumptions. While our estimates pertain to a particular AI tool deployed in a single industry, the mechanism – machine-learning-enabled control that improves the efficiency of energy-intensive operations – is likely to be relevant in other process industries. Consistent with this interpretation, the technology provider reports deploying closely related digital-twin systems beyond wastewater aeration, including steel galvanization in metallurgy, anaerobic digestion for biogas production, water-network operations, and agri-food processes. Given the breadth of processes where similar endpoint-control logic could apply, we complement our micro-level estimates with a macroeconomic assessment. Integrating our treatment effects into the DICE climate-economic model, we find welfare gains of $940 billion from avoided CO_2 emissions when extrapolating to processes where comparable solutions are already proved to be technically feasible, rising to $18.8 trillion under full industrial extrapolation. These figures illustrate the potential aggregate significance of AI-enabled efficiency gains, though the broader scenarios require extrapolation and should be interpreted accordingly. Extending the analysis to additional sectors, operating environments, and electricity systems would help assess external validity and quantify heterogeneity in impacts.

References

Abdessadak, Abdelali, Hicham Ghennioui, Nadège Thirion-Moreau, Brahim Elbhiri, Mounir Abraim, and Safae Merzouk. 2025. "Digital twin technology and artificial intelligence in energy transition: A comprehensive systematic review of applications." *Energy Reports*, 13: 5196–5218.

Akter, Shahriar, Samuel Fosso Wamba, Marcello Mariani, and Umme Hani. 2021. "How to Build an AI Climate-Driven Service Analytics Capability for Innovation and Performance in Industrial Markets?" *Industrial Marketing Management*, 97: 258–273.

Álvarez-Álvarez, Tania, A Barbón, L Bayón, and CA Silva. 2025. "Energy, exergy, sustainability, environmental emission, and fuel cost analysis of a hot-dip galvanised steel wire process." *Energy*, 319: 134897.

Andres, Pia, Eugenie Dugoua, and Marion Dumas. 2022. "Directed Technological Change and General Purpose Technologies: Can AI Accelerate Clean Energy Innovation?" Centre for Climate Change Economics and Policy, Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science Working Paper CCCEP Working Paper No. 403 / Grantham Working Paper No. 378.

Ashfaq, Muhammad, Jiang Yun, Shubin Yu, and Sandra Maria Correia Loureiro. 2020. "I, Chatbot: Modeling the determinants of usersâ satisfaction and continuance intention of AI-powered service agents." *Telematics and Informatics*, 54: 101473.

Attaran, Mohsen, and Bilge Gokhan Celik. 2023. "Digital Twin: Benefits, use cases, challenges, and opportunities." *Decision Analytics Journal*, 6: 100165.

Ba, Labouda, Fatma Tangour, Ikram El Abbassi, and Rafik Absi. 2025. "Analysis of Digital Twin Applications in Energy Efficiency: A Systematic Review." *Sustainability*, 17(8): 3560.

Barrage, Lint, and William Nordhaus. 2024. "Policies, projections, and the social cost of carbon: Results from the DICE-2023 model." *Proceedings of the National Academy of Sciences*, 121(13): e2312030121.

Bayati, Alisina, Amber Srivastava, Vedant Mundada, Srinivasa Salapaka, Hao Feng, and Amir Malvandi. 2025. "Energy efficiency in industrial drying: A hybrid ultrasonic system with a novel dynamic optimization framework." *Energy Conversion and Management*, 345: 120392.

Bogmans, Mr. Christian, Patricia Gomez-Gonzalez, Ganchimeg Ganpurev, Mr. Giovanni Melina, Mr. Andrea Pescatori, and Sneha D Thube. 2025. "Power Hungry: How AI Will Drive Energy Demand." International Monetary Fund IMF Working Papers 2025/081.

Bonfiglioli, Alessandra, Rosario Crinó, Mattia Filomena, and Gino Gancia. 2025. "Data, Power and Emissions: The Environmental Cost of AI." CESifo CESifo Working Paper Series 12158.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. "Revisiting Event-Study Designs: Robust and Efficient Estimation." *The Review of Economic Studies*, rdae007.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. "Generative AI at Work*." *The Quarterly Journal of Economics*, 140(2): 889–942.

Burian, Vlad, and Arthur Stalla-Bourdillon. 2025. "The increasing energy demand of artificial intelligence and its impact on commodity prices." *Economic Bulletin Boxes*, 2(None): None.

Callaway, Brantly, and Pedro H.C. Sant'Anna. 2021. "Difference-in-Differences with multiple time periods." *Journal of Econometrics*, 225(2): 200–230. Themed Issue: Treatment Effect 1.

Carvalho, Thyago P., Fabrízzio A. A. M. N. Soares, Roberto Vita, Roberto da P. Francisco, Joao P. Basto, and Symone G. S. Alcalá. 2019. "A systematic literature review of machine learning methods applied to predictive maintenance." *Computers Industrial Engineering*, 137: 106024.

CGDD. 2019. "L'assainissement collectif des eaux usées." *notre-environnement.gouv.fr*, Accessed: March 2026.

Charpentier, José, Guy Martin, Hervé Wacheux, and Pierre Gilles. 1998. "ORP regulation and activated sludge: 15 years of experience." *Water Science and Technology*, 38: 197–208.

Chattaraman, Veena, Wi-Suk Kwon, Juan E. Gilbert, and Kassandra Ross. 2019. "Should AI-Based, conversational digital assistants employ social- or task-oriented interaction style? A task-competency and reciprocity perspective for older adults." *Computers in Human Behavior*, 90: 315â330.

Chatterjee, Sheshadri, Nripendra P. Rana, Yogesh K. Dwivedi, and Abdullah M. Baabdullah. 2021. "Understanding AI adoption in manufacturing and production firms using an integrated TAM-TOE model." *Technological Forecasting and Social Change*, 170: 120880.

Cheng, Hong-Bang, Mathava Kumar, and Jih-Gaw Lin. 2012. "Interpretation of redox potential variation during biological denitrification using linear non-equilibrium thermodynamic model." *International Biodeterioration Biodegradation*, 67: 28–39.

Cui, Zheyuan, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz. 2025. "The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers."

da Silva, Ramon PP, Ikenna J Okeke, Kiran Thirumaran, Manish Mosalpuri, Forooza Samadi, Sachin U Nimbalkar, and Joe Cresko. 2025. "Current state and future projections of drying processes in the US food and pulp and paper sectors: Energy, economic, and environmental assessment." *Energy Conversion and Management: X*, 101181.

de Chaisemartin, Clément, and Xavier D'Haultfoeuille. 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review*, 110(9): 2964â96.

de Chaisemartin, Clément, and Xavier D'Haultfoeuille. 2024. "Difference-in-Differences Estimators of Intertemporal Treatment Effects." *The Review of Economics and Statistics*, 1–45.

de Chaisemartin, Clément, Xavier d'Haultfoeuille, Félix Pasquier, and Gonzalo Vazquez-Bare. 2022. "Difference-in-Differences Estimators for Treatments Continuously Distributed at Every Period." *SSRN Electronic Journal*.

Dell'Acqua, Fabrizio, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani. 2023. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." Harvard Business School Technology Operations Mgt. Unit The Wharton School Research Paper.

Deronzier, Ga¨elle, and Jean-Marc Choubert. 2004. "Phosphorus removal in small activated sludge wastewater treatment plants - Comparison of technical and economical aspects for biological and physico-chemical phosphorus removal." *HAL CCSD, Cemagref Editions*, 50.

de Vries, Alex. 2023. "The growing energy footprint of artificial intelligence." *Joule*, 7(10): 2191–2194.

EPA. 2022. "Clean Watersheds Needs Survey 2022." *https://sdwis.epa.gov/ords/sfdw_pub/r/sfdw/cwns_pub/wastewater-dashboard?session=15965764314405*, Accessed: 2025-11-23.

Eurostat. 2025*a*. "Population connected to at least secondary wastewater treatment." *https://ec.europa.eu/eurostat/databrowser/view/sdg_06_20/default/table?lang=en*, Accessed: 2025-11-11.

Eurostat. 2025*b*. "Use of artificial intelligence in enterprises (ISOC_EB_AI)." *Eurostat online data table*, Accessed 2026-01-23.

Fabra, Natalia, David Rapson, Mar Reguant, and Jingyuan Wang. 2021. "Estimating the Elasticity to Real-Time Pricing: Evidence from the Spanish Electricity Market." *AEA Papers and Proceedings*, 111: 425â29.

Fuller, Aidan, Zhong Fan, Charles Day, and Chris Barlow. 2020. "Digital Twin: Enabling Technologies, Challenges and Open Research." *IEEE Access*, 8: 108952â108971.

Goodman-Bacon, Andrew. 2021. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics*, 225(2): 254–277. Themed Issue: Treatment Effect 1.

Huang, Ming-Hui, and Roland T. Rust. 2022. "A Framework for Collaborative Artificial Intelligence in Marketing." *Journal of Retailing*, 98(2): 209â223.

IBM. 2024. "What is a digital twin?" *https://www.ibm.com/think/topics/what-is-a-digital-twin*, Accessed: 2025-09-25.

IEA. 2016. "World Energy Outlook 2016." International Energy Agency (IEA), Paris.

IEA. 2017. "Water-Energy Nexus." International Energy Agency Technical Report.

IEA. 2024. "World Energy Mix." *https://www.iea.org/world/energy-mix*, Accessed: March 2026.

IEA. 2025. "Energy and AI." International Energy Agency Technical Report.

International Energy Agency. 2016. "Decoupling of Global Emissions and Economic Growth Confirmed." *https://www.iea.org/news/decoupling-of-global-emissions-and-economic-growth-confirmed*, Accessed: March 2026.

International Energy Agency. 2025. "Global CO2 Emissions from Electricity Generation, 2015–2027." *https://www.iea.org/data-and-statistics/charts/global-co2-emissions-from-electricity-generation-2015-2027*, Accessed: March 2026.

Jessoe, Katrina, and David Rapson. 2014. "Knowledge Is (Less) Power: Experimental Evidence from Residential Energy Use." *American Economic Review*, 104(4): 1417â38.

Li, Minglong, Dexiang Yin, Hailian Qiu, and Billy Bai. 2021. "A systematic review of AI technology-based service encounters: Implications for hospitality and tourism operations." *International Journal of Hospitality Management*, 95: 102930.

Luccarini, Luca, Dalila Pulcini, Davide Sottara, Roberto Di Cosmo, and Roberto Canziani. 2017. "Monitoring denitrification by means of pH and ORP in continuous-flow conventional activated sludge processes." *Desalination and Water Treatment*, 61: 319–325.

Luo, Kaikai, Fen Wang, and Xuezhen Chen. 2025. "Impact of artificial intelligence on energy efficiency in Chinese enterprises." *International Review of Economics & Finance*, 103(C): None.

Machado, JosÃ© A.F., and J.M.C. Santos Silva. 2019. "Quantiles via moments." *Journal of Econometrics*, 213: 145–173.

Magni, Michele, Edward R Jones, Marc FP Bierkens, and Michelle TH van Vliet. 2025. "Global energy consumption of water treatment technologies." *Water Research*, 277: 123245.

MarketsandMarkets. 2025. "Digital Twin Market - Size, Share, Trend and Forecast (2025 - 2030)."

Masanet, Eric, Arman Shehabi, Nuoa Lei, Sarah Smith, and Jonathan Koomey. 2020. "Recalibrating global data center energy-use estimates." *Science*, 367(6481): 984–986.

McKinsey. 2024. "Digital twins: The next frontier of factory optimization." *https://www.mckinsey.com/capabilities/operations/our-insights/digital-twins-the-next-frontier-of-factory-optimization*, Accessed: 2025-09-22.

McLean, Graeme, Kofi Osei-Frimpong, and Jennifer Barhorst. 2021. "Alexa, do voice assistants influence consumer brand engagement? â Examining the role of AI powered voice assistants in influencing consumer brand engagement." *Journal of Business Research*, 124: 312â328.

Metcalfe, R.D., R.S. Schein, C.R. Schein, and Y. Sun. 2026. "AI in Charge: Large-Scale Experimental Evidence on Electric Vehicle Charging Demand." NBER Working Paper 34709.

Muench, S., E. Stoermer, K. Jensen, T. Asikainen, M. Salvi, and F. Scapolo. 2022. "Towards a green and digital future." EU JRC, Publications Office of the European Union, Luxembourg.

Mujumdar, Arun S. 2007. "An overview of innovation in industrial drying: current status and R&D needs." *Transport in porous media*, 66(1): 3–18.

Myers, Matthew, Leland Myers, and Robert Okey. 2006. "The Use of Oxidation-Reduction Potential as a Means of Controlling Effluent Ammonia Concentration in an Extended Aeration Activated Sludge System."

Noy, Shakked, and Whitney Zhang. 2023. "Experimental evidence on the productivity effects of generative artificial intelligence." *Science*, 381(6654): 187–192.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. 2023. "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot."

Rettig, Stefan, Karin Schulz, Matthias Barjenbruch, and Iyad Al-Zreiqat. 2018. "Key figure data for energy efficiency. Benchmarking the Baltic Sea Region in the project IWAMA - Interactive Water Management." IWAMA Project.

Ritchie, Hannah. 2020. "Sector by sector: where do global greenhouse gas emissions come from?" *Our World in Data*. https://archive.ourworldindata.org/20251125-173858/ghg-emissions-by-sector.html.

Rojahn, Marcel, and Marcus Grum. 2025. "Green AI: A systematic review and meta-analysis of its definitions, lifecycle models, hardware and measurement attempts."

Rolnick, David, Priya L. Donti, Lynn H. Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, Alexandra Sasha Luccioni, Tegan Maharaj, Evan D. Sherwin, S. Karthik Mukkavilli, Konrad P. Kording, Carla P. Gomes, Andrew Y. Ng, Demis Hassabis, John C. Platt, Felix Creutzig, Jennifer Chayes, and Yoshua Bengio. 2022. "Tackling Climate Change with Machine Learning." *ACM Comput. Surv.*, 55(2).

RTE. 2026. "éCO$_2$mix - Les émissions de CO2 par kWh produit en France." *https://www.rte-france.com/donnees-publications/eco2mix-donnees-temps-reel/emissions-co2-par-kwh-produit-france*, Accessed: 2026-01-04.

Schwartz, Roy, Jesse Dodge, Noah A. Smith, and Oren Etzioni. 2020. "Green AI." *Communications of the ACM*, 63(12): 54–63.

SDEA. 2017. "Schémas des cycles de l'eau." *sdea.fr*, Syndicat des Eaux et de l'Assainissement Alsace-Moselle.

Stern, Nicholas, Mattia Romani, Roberta Pierfederici, Manuel Braun, Daniel Barraclough, Shajeeshan Lingeswaran, Elizabeth Weirich-Benet, and Niklas Niemann. 2025. "Green and intelligent: the role of AI in the climate transition." *npj Climate Action*, 4.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics*, 225(2): 175–199. Themed Issue: Treatment Effect 1.

US DoE. 2017. "Energy Data Management Manual for the Wastewater Treatment Sector." *https://www.energy.gov/sites/prod/files/2018/01/f46/WastewaterTreatmentDataGuide_Final_0118.pdf*, Accessed: 2025-11-23.

WSA. 2025. "World Steel in Figures 2025." *https://worldsteel.org/data/world-steel-in-figures/world-steel-in-figures-2025/*, Accessed: March 2026.

Xu, Jingjing, Wangchunshu Zhou, Zhiyi Fu, Hao Zhou, and Lei Li. 2021. "A Survey on Green Deep Learning."

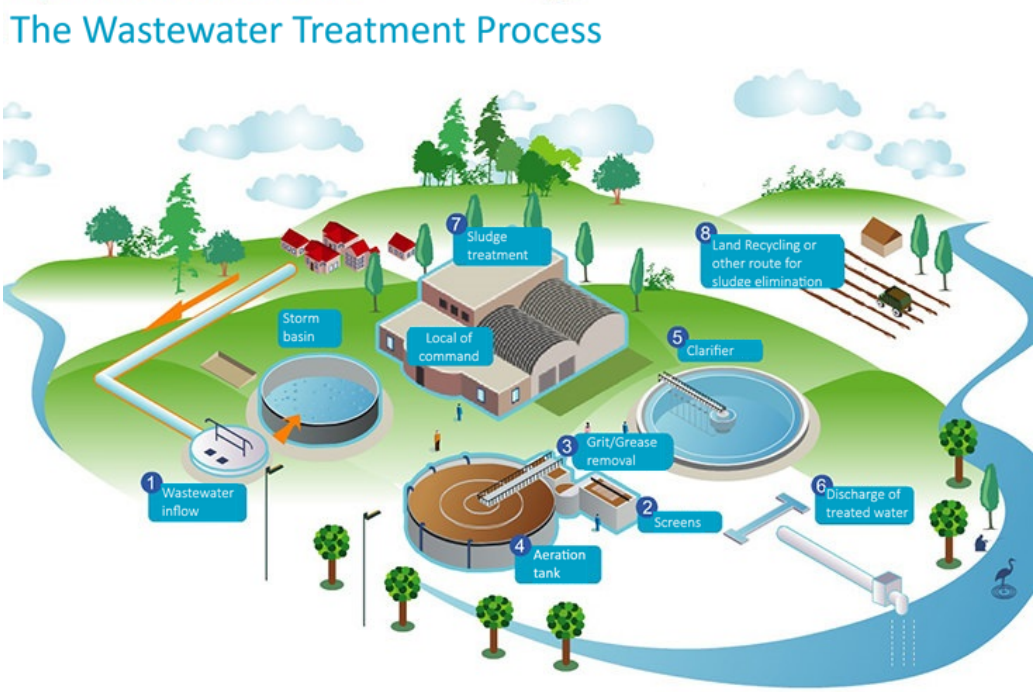
Yussuf, Raheemat O., and Omar S. Asfour. 2024. "Applications of artificial intelligence for energy efficiency throughout the building lifecycle: An overview." *Energy and Buildings*, 305: 113903.

Zonta, Tiago, Cristiano André da Costa, Rodrigo da Rosa Righi, Miromar Joé de Lima, Eduardo Silveira da Trindade, and Guann Pyng Li. 2020. "Predictive maintenance in the Industry 4.0: A systematic literature review." *Computers Industrial Engineering*, 150: 106889.

Appendix

A Wastewater Treatment Plant Diagram and Photo

Figure A.1: Diagram of a typical activated sludge WWTP



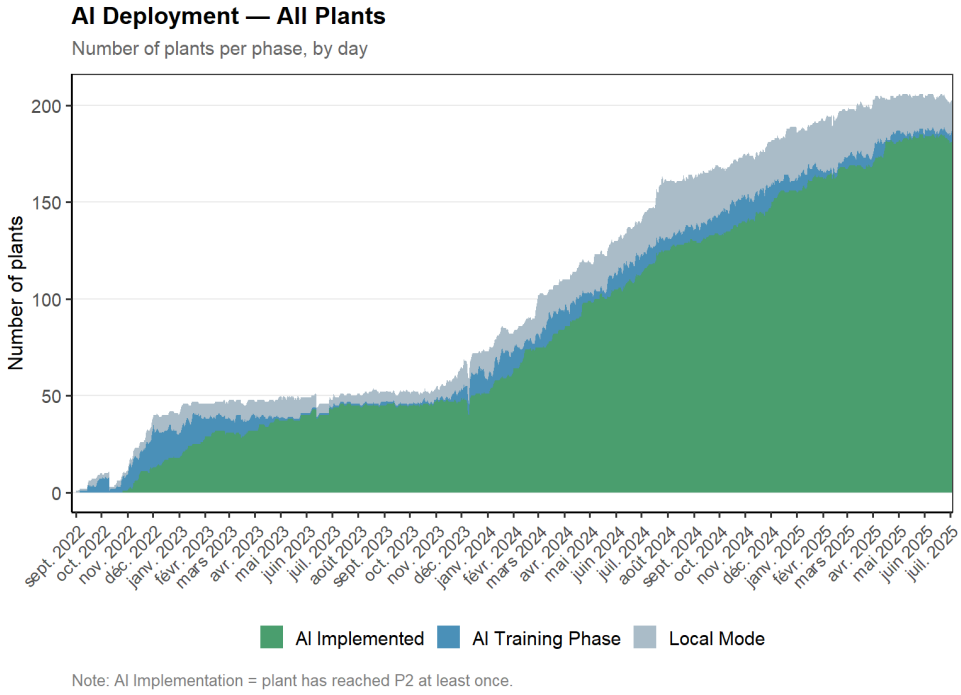
Notes. This diagram is adapted from a presentation of SDEA (2017).

Figure A.2: Aerial Photo of the WWTP of Bonneville (activated sludge)

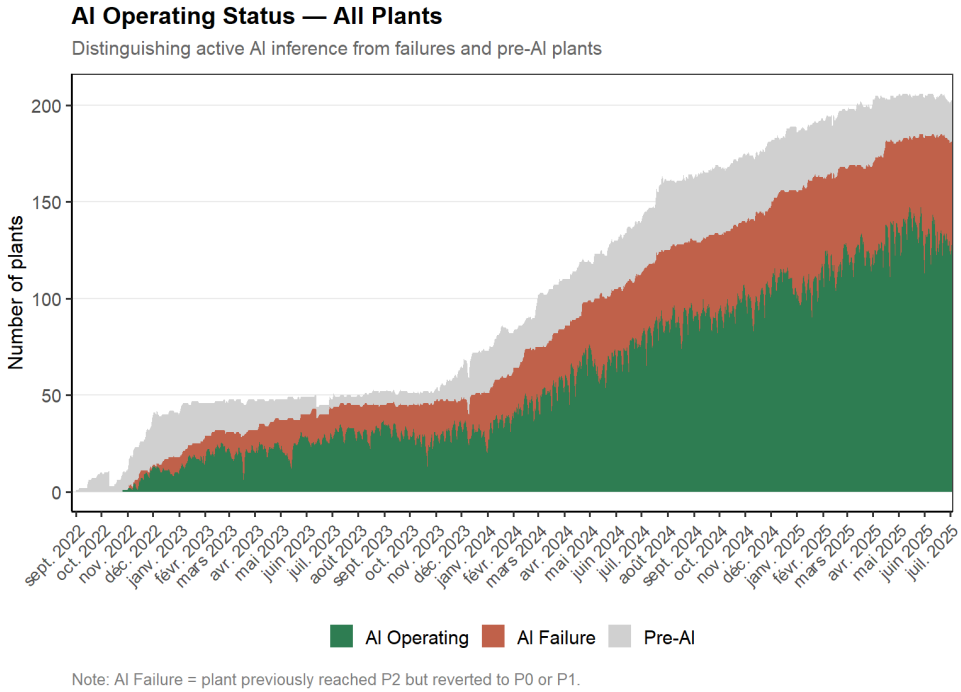


Notes. This photo was provided by Veolia.

B AI deployment



(a) Plant with AI installed



(b) Plant with operating AI

Figure B.1: Roll-out of AI adoption and use.

Note: In subfigure (a), we are interested to the phase of AI implementation (not implemented, in training or installed). In subfigure (b), we are interested in AI operating status (not yet installed, operative or failing).

C Digital Twins

Figure C.1: Generic Diagram of a Digital Twin

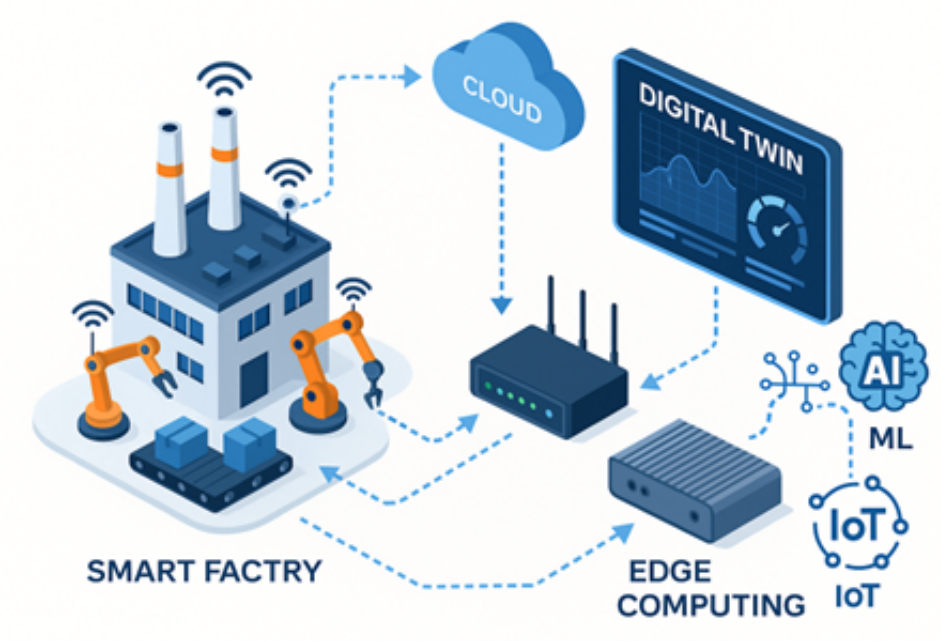
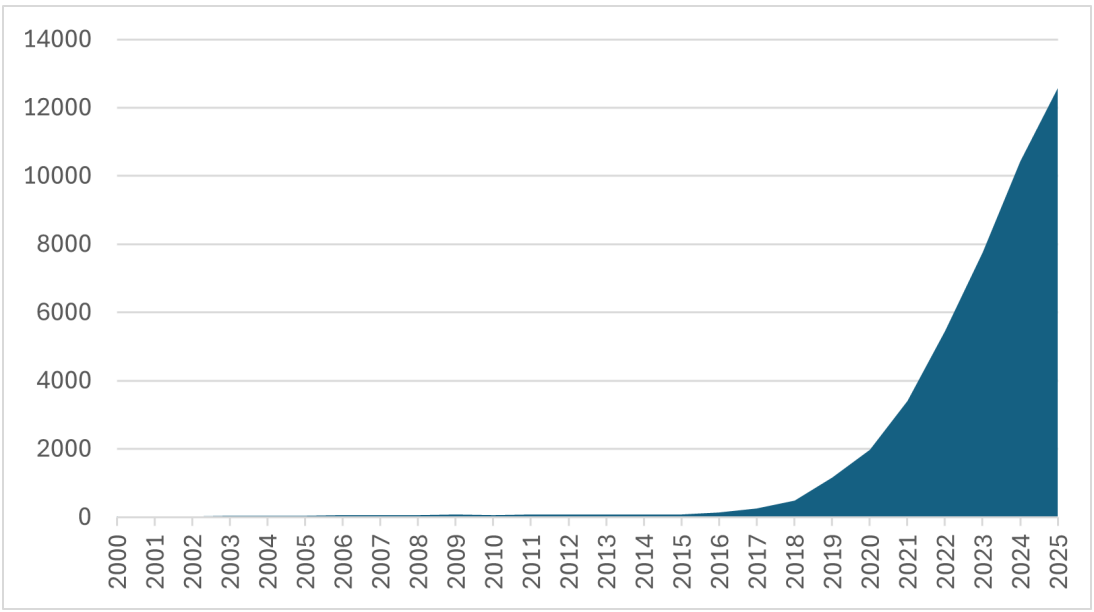


Figure C.2: Occurrence of "Digital Twins" in Research

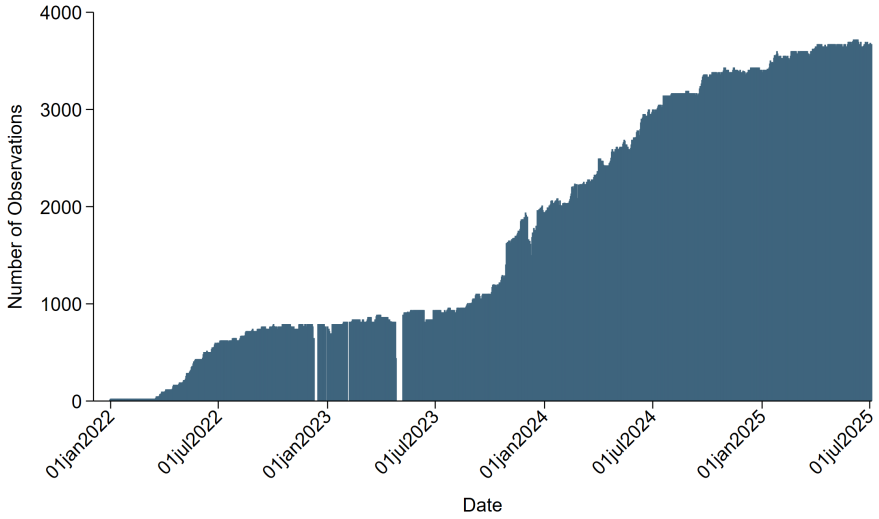


Notes. This figure plots the occurrence of "Digital Twins" (article title, abstract, keywords) by year in the Scopus database.

D Additional Descriptive Statistics

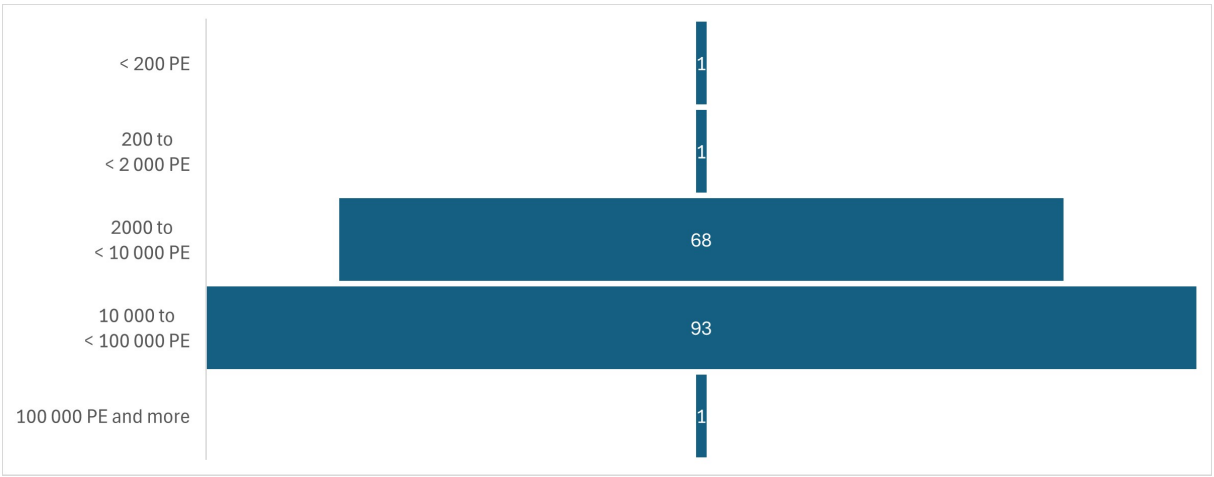
D.1 Number of observations by date

Figure D.1: Distribution of observations over the sample



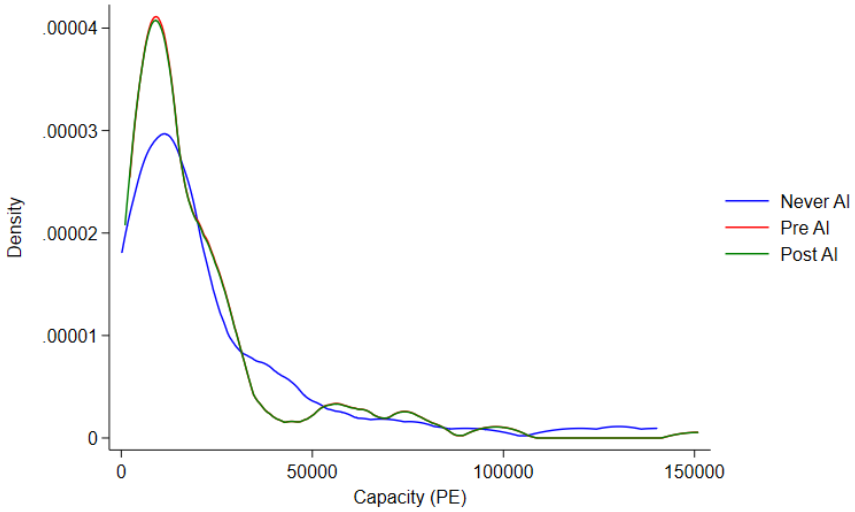
D.2 Comparison of the sample to the French national distribution of WWTP

Figure D.2: Distribution of the number of plants per size in our sample



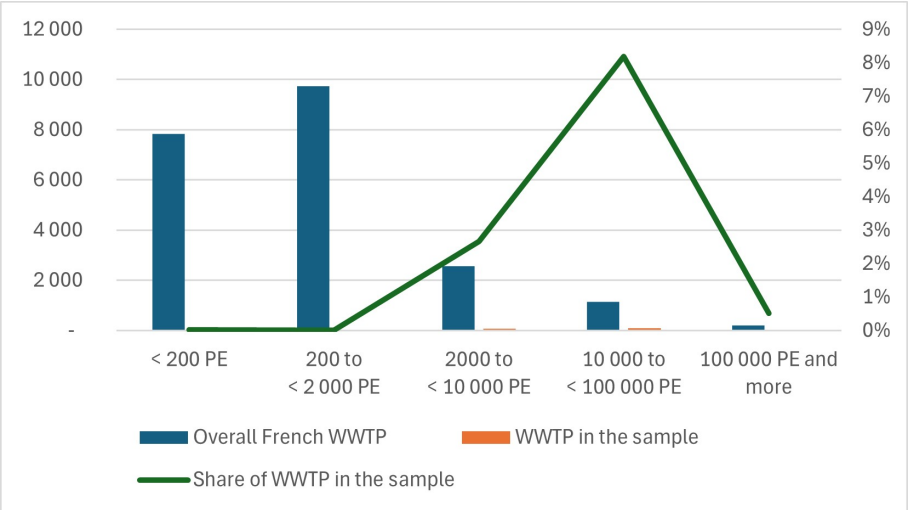
Notes. This graph shows the number of plant per size category in our sample.

Figure D.3: Distribution of plant nominal capacity



).

Figure D.4: Distribution of observations over the sample



Notes. This graph shows the absolute number of plants in our sample vs. the French national distribution (Source: CGDD (2019)).

D.3 Statistics by AI adoption wave

Table D.1: Summary statistics by AI adoption waves

Variable	**First wave**		**Second wave**	
	(Pre AI)	(Post AI)	(Pre AI)	(Post AI)
Plants	40	-	99	-
Mean adoption date	24 Feb. 2023	-	25 Jul. 2024	-
Nominal capacity (PE)	14,809	-	22,301	-
Observations	188,292	719,060	395,979	785,188
AI usage (Share of time)	0	0.585	0	0.579
Power (kW)	59	56	68	75
Carbon em. ($kgCO_2e/h$)	3.2	1.5	1.8	1.7
Electricity cost (€/h)	3.5	5.4	7.4	9.1

Notes: This table presents the summary statistics of the sample by AI adoption wave, before and after the AI implementation.

D.4 Sample with additional variables

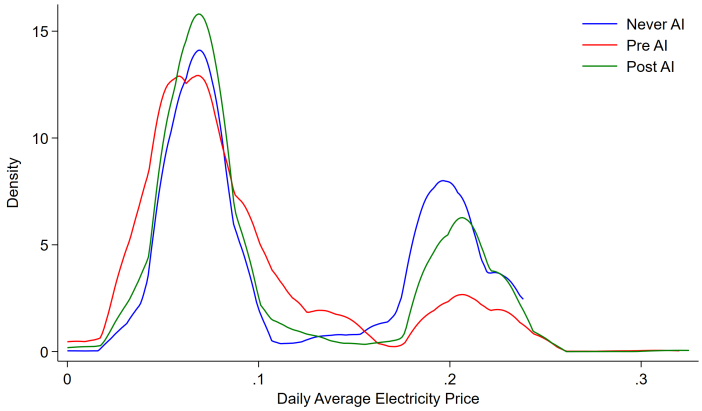
Table D.2: Summary statistics - sample with additional variables

Variable	**All**	**Never AI**	**Pre-AI**	**Post-AI**
Observations	1,635,615	108,329	264,388	1,262,898
Plants	150	21	90	129
Nominal capacity (PE)	19,199	17,992	21,316	19,396
AI usage (Share of time)	0.45	0	0	0.6
Power (kW)	66.6	85.8	68.6	64.1
Carbon em. ($kgCO_2e/h$)	1.58	2.1	1.7	1.51
Electricity price (€/kWh)	0.12	0.13	0.11	0.12
Electricity cost (€/h)	7.7	11.6	7.7	7.27
Inflow (m3)	115	164	129	109
Rain (mm)	0.13	0.11	0.15	0.13
Redox (mV)	22.17	16.5	30	21

Notes: This table presents the summary statistics of the subsample restricted to observations for which the control variables (inflow, rainfall and redox) are available. The first column consists of all plants in our sample, the second column includes plants who never receive coverage by AI. The third column presents statistics for treated plants before they receive AI installation, and the fourth column includes treated plants after the AI deployment.

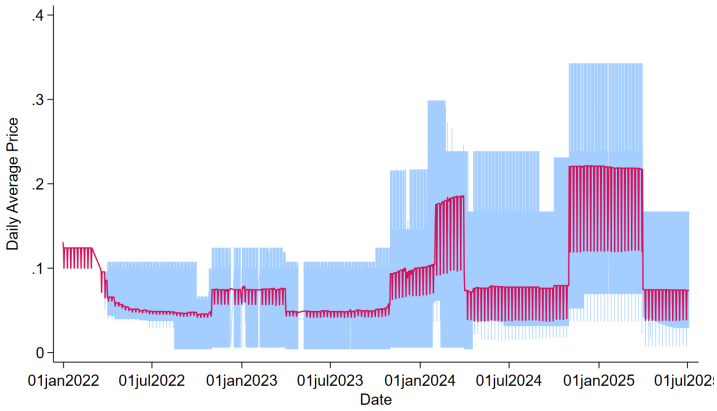
D.5 Electricity price

Figure D.5: Daily electricity price distribution



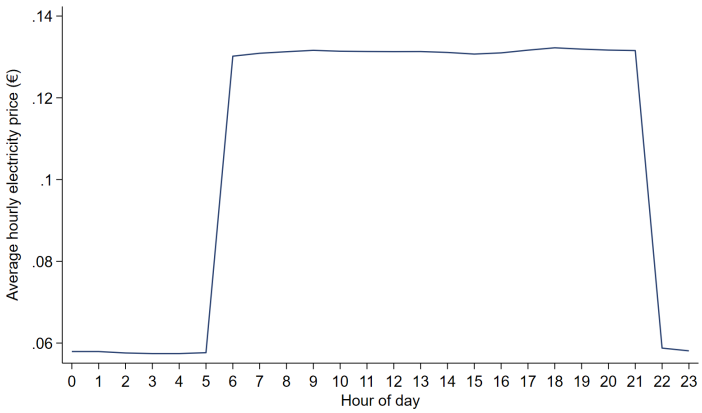
Notes. This graph shows the distribution of the average daily electricity price.

Figure D.6: Daily electricity price trajectories across plants



Notes. The average daily electricity price across plants is plotted in red, while the daily electricity prices of each plants appear in blue.

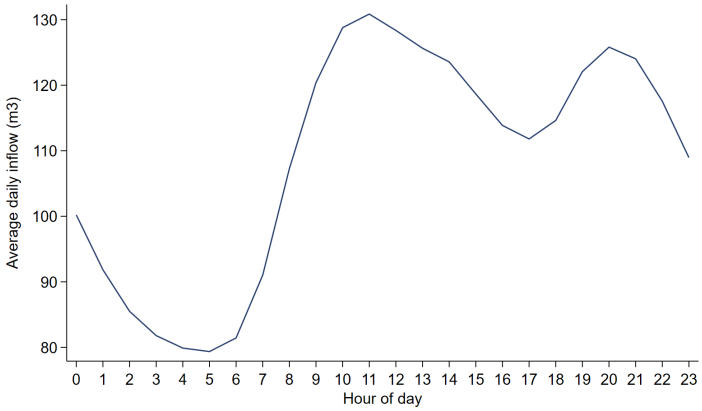
Figure D.7: Daily average electricity price



Notes.

D.6 Inflow

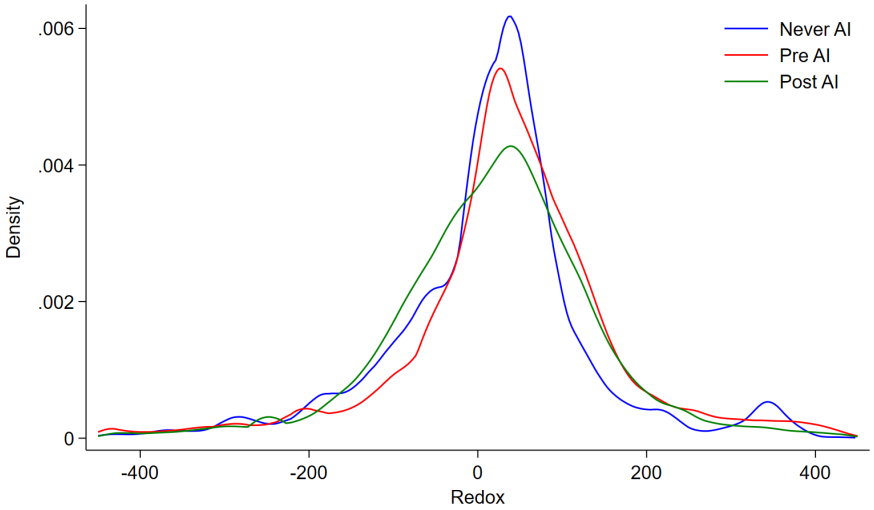
Figure D.8: Inflow during the day



Notes.

D.7 ORP Distribution

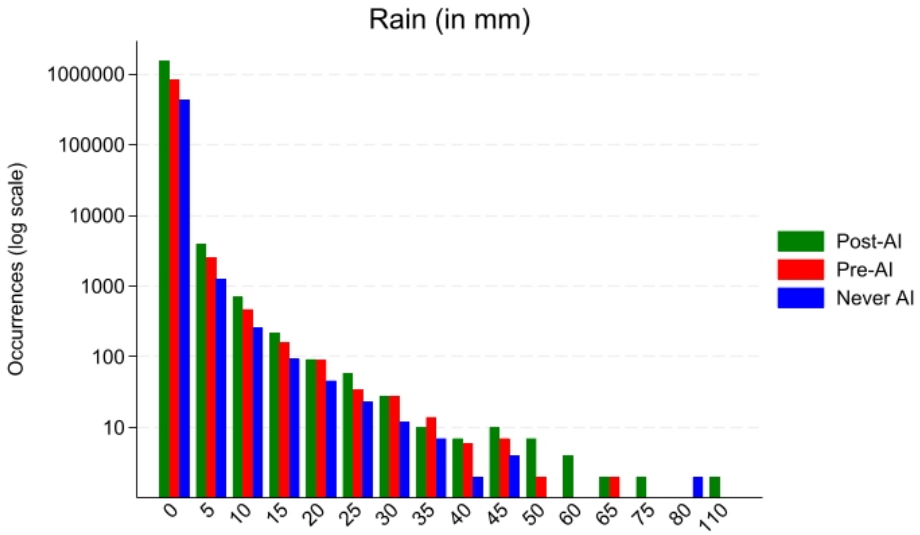
Figure D.9: Distribution of Oxido-Reduction Potential (ORP) values



Notes. Observations for redox values are at the hour-level.

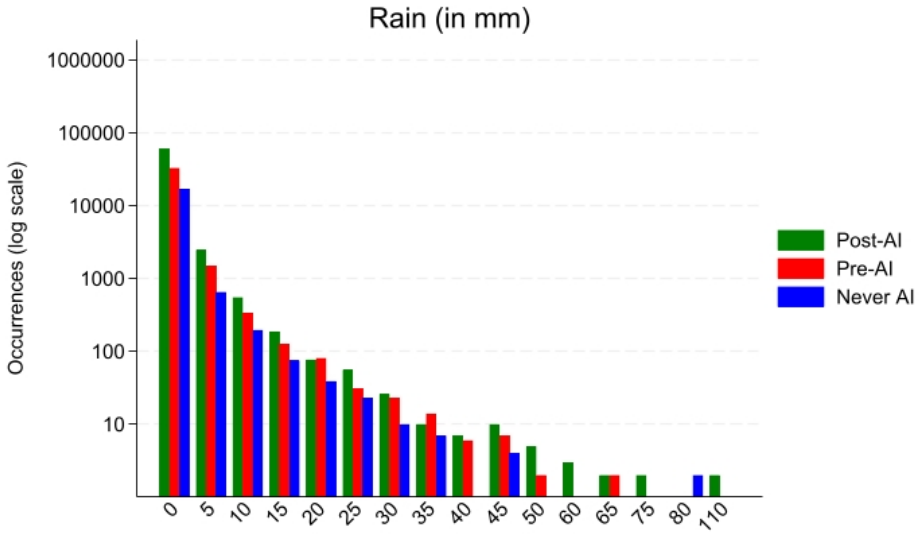
D.8 Rainfall Distribution

Figure D.10: Distribution of rainfall by hour



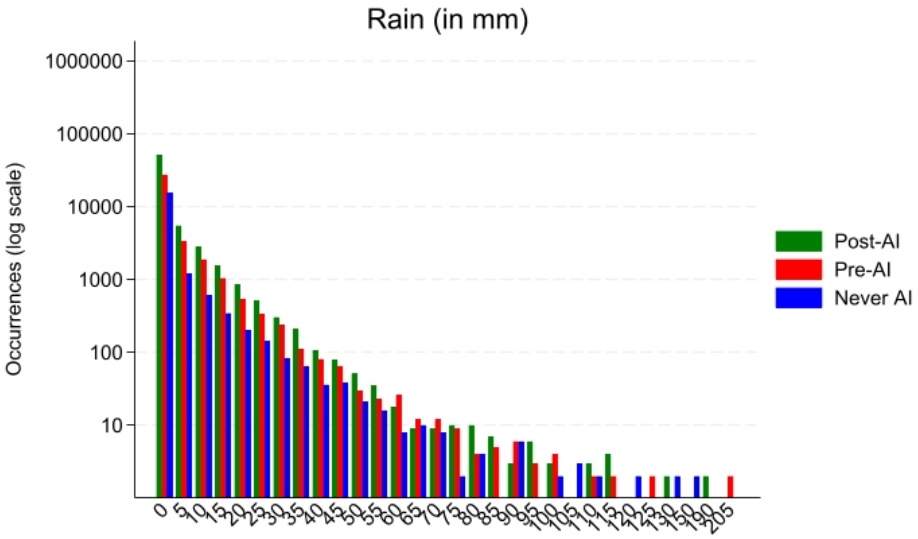
Notes. Observations for rainfall (in mm) are at the hour-level.

Figure D.11: Distribution of daily peak hourly rainfall



Notes. Observations for rainfall (in mm) are at the hour-level but plotted by day.

Figure D.12: Distribution of rainfall by day

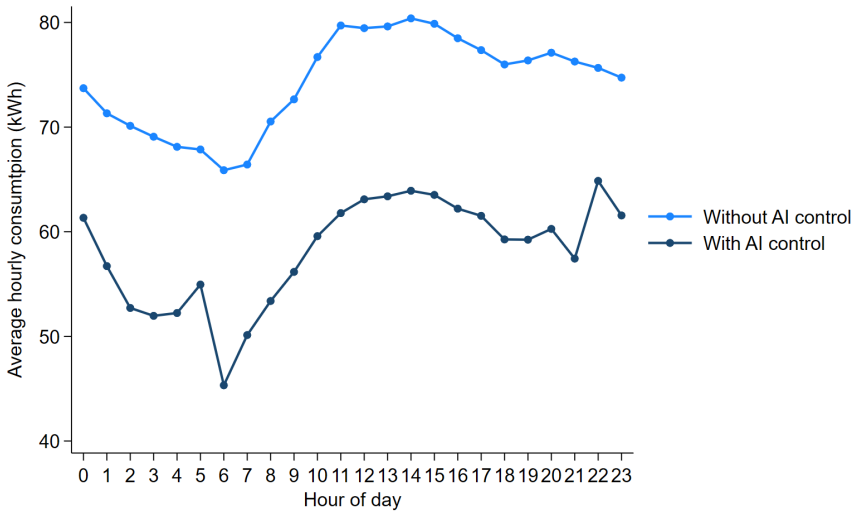


Notes. Observations for rainfall (in mm) are at the day-level.

E Effect of AI on electricity consumption and carbon footprint

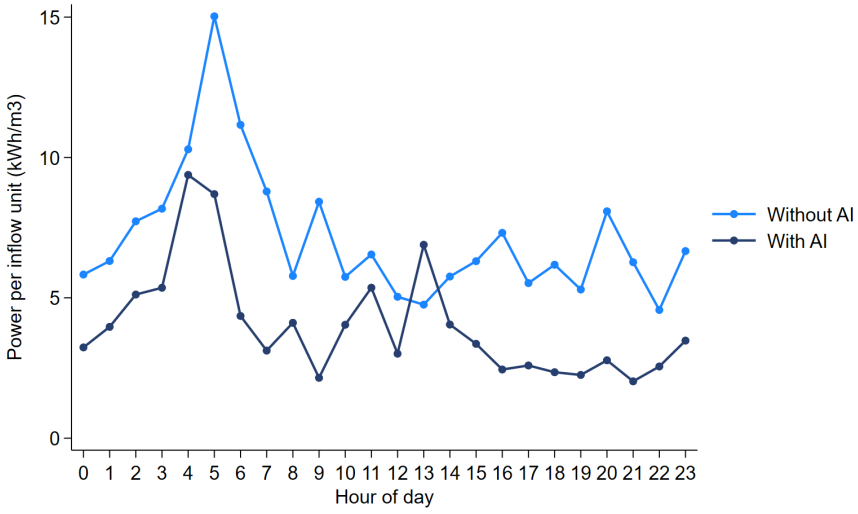
E.1 Descriptive analysis on electricity consumption and efficiency

Figure E.1: Hourly electricity consumption with and without AI control



Notes.

Figure E.2: Hourly electricity efficiency with and without AI control



Notes.

E.2 Additional TWFE

E.2.1 Hourly level

Table E.1: Result at the hourly level with a TWFE-OLS

	AI usage	Highly Treated	Treated	Ever Treated
Electricity Consumption				
	-0.043^{***}	-0.039^{***}	-0.03^{**}	-0.002
	(0.009)	(0.007)	(0.008)	(0.011)
Carbon footprint				
	-0.044^{***}	-0.04^{***}	-0.031^{***}	-0.001
	(0.008)	(0.007)	(0.008)	(0.012)
Electricity cost				
	-0.032^{***}	-0.029^{***}	-0.021^{***}	-0.024^{**}
	(0.008)	(0.007)	(0.008)	(0.012)
Observations	2,236,729			
Number of plants	164			

Notes: Electricity consumption, carbon footprint, and electricity cost are expressed in log values. The definition of the measure of AI use differs across columns. Effect on electricity price is obtained controlling for the mean individual daily electricity price. All specifications include plant and date fixed effects. Standard errors are clustered at the plant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. R^2 exceeds 0.95 in all specifications.

E.2.2 Daily level

Table E.2: Result at the daily level with a TWFE-OLS

	AI usage	Highly Treated	Treated	Ever Treated
Electricity Consumption				
	-0.043^{***}	-0.032^{***}	-0.029^{***}	-0.003
	(0.009)	(0.007)	(0.008)	(0.011)
Carbon footprint				
	-0.044^{***}	-0.03^{***}	-0.029^{***}	-0.001
	(0.008)	(0.007)	(0.008)	(0.012)
Electricity cost				
	-0.045^{***}	-0.033^{***}	-0.036^{***}	-0.035^{***}
	(0.009)	(0.009)	(0.008)	(0.014)
Observations	93,713			
Number of plants	164			

Notes: Electricity consumption, carbon footprint, and electricity cost are expressed in log values. The definition of the measure of AI use differs across columns. All specifications include plant and date fixed effects. Effect on electricity price is obtained controlling for the mean individual daily electricity price. Standard errors are clustered at the plant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. R^2 exceeds 0.95 in all specifications.

E.3 Estimator comparison

E.3.1 Daily level

Table E.3: Daily effects comparing DiD estimators

	TWFE-OLS (Std. Error)	De Chaisemartin-D'Haultfoeuille (Std. Error)
Electricity consumption	-0.044***	-0.015*
	(0.007)	(0.015)
Carbon footprint	-0.044***	-0.027*
	(0.007)	(0.016)
Electricity Cost	-0.05***	-0.044**
	(0.009)	(0.022)

Notes. This table reports the estimated impact of AI usage on daily electricity consumption expressed in log values at the plant level. We compare a standard two-way fixed effects (TWFE) estimator, with the average cumulative total effect per plant estimated over 15 days with de Chaisemartin and D'Haultfoeuille (2024). Effect on cost is obtained controlling for the mean individual weekly price. Standard errors are clustered at the plant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

E.3.2 Weekly level

Table E.4: Weekly effects comparing DiD estimators

	TWFE-OLS (Std. Error)	De Chaisemartin-D'Haultfoeuille (Std. Error)
Electricity consumption	-0.044***	-0.054***
	(0.009)	(0.017)
Carbon footprint	-0.041***	-0.06***
	(0.009)	(0.018)
Electricity Cost	-0.052***	-0.082***
	(0.012)	(0.017)

Notes. This table reports the estimated impact of AI usage on weekly electricity consumption expressed in log values at the plant level. We compare a standard two-way fixed effects (TWFE) estimator, with the average cumulative total effect per plant estimated over 10 weeks with de Chaisemartin and D'Haultfoeuille (2024). Effect on cost is obtained controlling for the mean individual weekly price. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

E.3.3 Monthly level

Table E.5: Monthly effects comparing DiD estimators

	TWFE-OLS (Std. Error)	De Chaisemartin-D'Haultfoeuille (Std. Error)
Electricity consumption	-0.05***	-0.03
	(0.011)	(-0.022)
Carbon footprint	-0.046***	-0.044*
	(0.011)	(0.025)
Electricity Cost	-0.063***	-0.073***
	(0.01)	(0.027)

Notes. This table reports the estimated impact of AI usage on monthly electricity consumption expressed in log values at the plant level. We compare a standard two-way fixed effects (TWFE) estimator, with the average cumulative total effect per plant estimated over 3 months with de Chaisemartin and D'Haultfoeuille (2024). Effect on cost is obtained controlling for the mean individual weekly price. Standard errors are clustered at the plant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

E.4 Absolute effects

Table E.6: Absolute weekly effects

	TWFE-OLS (Std. Error)	De Chaisemartin-D'Haultfoeuille (Std. Error)
Electricity consumption	-314.9***	-606.8***
	(97.8)	(233.2)
Carbon footprint	11.3	-22.8
	(10.5)	(18.9)
Electricity Cost	-73.7*	-199**
	(38.5)	(87.9)

Notes. This table reports the estimated impact of AI usage on monthly electricity consumption at the plant level. We compare a standard two-way fixed effects (TWFE) estimator, following Equation (**??**), with the average cumulative total effect per plant estimated over 10 weeks with de Chaisemartin and D'Haultfoeuille (2024). Absolute effects are expressed in kilowatt-hours (kWh) for electricity consumption, $kgCO_2e$ for the carbon footprint, and euros for the electricity cost. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

F Conceptual Framework

F.1 Environment and Technology

Consider a wastewater treatment plant (or industrial facility) indexed by $i = 1, \ldots, N$ operating over periods $t = 1, \ldots, T$. Each plant must deliver a target service level Q_{it} (in this context required aeration of water) using two inputs: a "generic factor" L_i (for simplicity we say it is flow rate) and an energy-intensive input A_{it} (aeration of water).

Production technology. Service is produced via a constant elasticity of substitution (CES) technology:

$$Q_{it} = [\alpha L_i^{\rho} + (1-\alpha) A_{it}^{\rho}]^{1/\rho}, \quad \alpha \in (0,1),\ \rho \in (0,1).$$

The elasticity of substitution between L and A is $\sigma = 1/(1-\rho) > 1$. For a given target Q_{it} and generic factor L_i, the minimal aeration (absent shocks) is:

$$A(Q_{it}, L_i) = \left[\frac{Q_{it}^{\rho} - \alpha L_i^{\rho}}{1-\alpha}\right]^{1/\rho}.$$

Operational uncertainty. Actual conditions (influent load, temperature, weather) introduce a stochastic "difficulty shifter" χ_{it}. The required aeration to achieve Q_{it} is:

$$R_{it} = A(Q_{it}, L_i) \cdot g(\chi_{it}) + \varepsilon_{it}, \tag{4}$$

where $g(\cdot)$ is increasing, $\mathbb{E}[g(\chi_{it})] = 1$, and ε_{it} is idiosyncratic noise with $\mathbb{E}[\varepsilon_{it}|\chi_{it}] = 0$. The plant operator observes an information set $\mathcal{I}_{it}$ and must choose a setpoint a_{it} before realizing R_{it}.

F.2 Decision Problem: Asymmetric Loss and Price

Cost structure. The operator faces: (i) an energy cost $p_{it}a_{it}$, where p_{it} is the electricity price, and (ii) a non-monetary loss from deviations between chosen aeration a_{it} and required aeration R_{it}. We model this loss as an asymmetric piecewise-linear function:

$$\ell(a, R) = \phi \cdot (a - R)_+ + \psi \cdot (R - a)_+, \quad \psi > \phi > 0, \tag{5}$$

where

$$(x)_+ = \max\{x, 0\}$$

. The interpretation:

- **Overshoot** ($a > R$): excess aeration wastes energy (cost ϕ per unit).
- **Undershoot** ($a < R$): insufficient aeration risks non-compliance, effluent violations, fines (cost ψ per unit, $\psi > \phi$).

Objective. The operator solves:

$$\min_{a \geq 0} \mathbb{E}\left[p_{it} a + \ell(a, R_{it}) \mid \mathcal{I}_{it}\right]. \tag{6}$$

Optimal decision rule. Let $F_{R|\mathcal{I}}$ denote the CDF of R_{it} conditional on $\mathcal{I}_{it}$. The first-order condition is:

$$p_{it} + \phi \cdot \Pr(R_{it} < a^* \mid \mathcal{I}_{it}) - \psi \cdot \Pr(R_{it} > a^* \mid \mathcal{I}_{it}) = 0. \tag{7}$$

Rearranging:

$$F_{R|\mathcal{I}}(a^*) = \frac{\psi - p_{it}}{\phi + \psi}. \tag{8}$$

Define the critical probability:

$$\tau(p_{it}) := \frac{\psi - p_{it}}{\phi + \psi}. \tag{9}$$

where $\tau(p_{it}) \in (0, 1)$ requires $p_{it} < \psi$, i.e. the electricity price must be below the unit cost of non-compliance, which we assume holds throughout.

Then the optimal setpoint is:

$$a^*_{it} = F^{-1}_{R|\mathcal{I}_{it}}\big(\tau(p_{it})\big).$$

Remark. When $p_{it} = 0$, $\tau = \psi/(\phi+\psi) > 1/2$, so the setpoint exceeds the median: the operator "over-provisions" to guard against costly undershoots. Higher prices p_{it} reduce τ, lowering the margin.

F.3 Artificial Intelligence as Information Refinement

Pre-AI information. Before AI adoption, the operator relies on a baseline information set $\mathcal{I}^{\text{pre}}_{it}$ (simple heuristics, historical averages, limited instrumentation). The predictive distribution of R_{it} has mean μ_{it} and dispersion θ^{pre}_{it}.

Post-AI information. AI systems (sensor fusion, machine learning, real-time forecasting) provide a refined information set $\mathcal{I}^{\text{AI}}_{it}$. We model this as a mean-preserving contraction (or Blackwell-more-informative experiment (Blackwell, 1953)):

$$\mathbb{E}[R_{it} \mid \mathcal{I}^{\text{AI}}_{it}] = \mathbb{E}[R_{it} \mid \mathcal{I}^{\text{pre}}_{it}] = \mu_{it}, \quad \text{Var}[R_{it} \mid \mathcal{I}^{\text{AI}}_{it}] < \text{Var}[R_{it} \mid \mathcal{I}^{\text{pre}}_{it}].$$

Denote $\theta^{\text{AI}}_{it} < \theta^{\text{pre}}_{it}$.

Location-scale family (working assumption). For tractability, assume:

$$R_{it} = \mu_{it} + \theta_{it} \cdot Z, \quad Z \sim F_Z, \ \mathbb{E}[Z] = 0, \ \text{Var}(Z) = 1^{30}. \tag{10}$$

[30]The variance of Z is supposed to be equal to 1 without loss of generality.

Then:

$$a_{it}^{*} = \mu_{it} + \theta_{it} \cdot q_{\tau(p_{it})}(Z), \tag{11}$$

where $q_\tau(Z) = F_Z^{-1}(\tau)$ is the τ-quantile of Z. AI reduces θ_{it}, directly lowering the "safety margin" $\theta_{it} \cdot q_\tau$.

F.4 Testable Implications

We derive five empirical predictions from the model, indexed as Implications 1–5.

Implication 1: Average Energy Savings

Prediction. AI adoption reduces expected aeration and energy consumption.

Proof. From (11):

$$\mathbb{E}[a_{it}^{\text{AI}}] = \mathbb{E}[\mu_{it}] + \mathbb{E}[\theta_{it}^{\text{AI}}] \cdot q_\tau, \quad \mathbb{E}[a_{it}^{\text{pre}}] = \mathbb{E}[\mu_{it}] + \mathbb{E}[\theta_{it}^{\text{pre}}] \cdot q_\tau.$$

Since $\theta_{it}^{\text{AI}} < \theta_{it}^{\text{pre}}$ and $q_\tau > 0$ (for $\tau > 1/2$):

$$\mathbb{E}[a_{it}^{\text{AI}}] < \mathbb{E}[a_{it}^{\text{pre}}].$$

Energy consumption $E_{it} = \kappa \cdot a_{it}$ (with $\kappa > 0$), so $\mathbb{E}[E_{it}^{\text{AI}}] < \mathbb{E}[E_{it}^{\text{pre}}]$. □

Implication 2: Heterogeneity by Inflow Level

Prediction. Energy savings are larger (in absolute or percentage terms) for low flow rate levels (low L_i).

Mechanism. Suppose baseline uncertainty θ_i^{pre} is decreasing in L_i (larger plants have more robust processes, better pre-AI instrumentation). Define the energy saving:

$$S_i := \mathbb{E}[E_i^{\text{pre}}] - \mathbb{E}[E_i^{\text{AI}}] = \kappa \cdot (\theta_i^{\text{pre}} - \theta_i^{\text{AI}}) \cdot q_\tau.$$

If AI reduces dispersion proportionally ($\theta_i^{\text{AI}} = k \cdot \theta_i^{\text{pre}}$, $k < 1$), then:

$$S_i = \kappa \cdot (1 - k) \cdot \theta_i^{\text{pre}} \cdot q_\tau.$$

Thus $\partial S_i / \partial L_i < 0$ if $\partial \theta_i^{\text{pre}} / \partial L_i < 0$. □

Implication 3: AI Creates and Amplifies Price Responsiveness

Prediction. AI adoption generates price responsiveness where pre-AI operations exhibit none, enabling temporal reallocation of energy use toward low-price periods.

Setup. Partition time into high-price periods (p^H) and low-price periods ($p^L < p^H$), with average price $\bar{p} = \lambda p^H + (1-\lambda)p^L$, $\lambda \in (0,1)$.

Pre-AI behavior. Without real-time optimization capability, the pre-AI operator sets a fixed setpoint based on average price conditions, unresponsive to realized intraday price variation:

$$a^{\text{pre}} = \mu + \theta^{\text{pre}} \cdot q_{\tau(\bar{p})}, \qquad \frac{\partial a^{\text{pre}}}{\partial p_{it}} = 0. \tag{12}$$

Post-AI behavior. The AI system continuously adjusts the setpoint to the prevailing real-time electricity price:

$$a_{it}^{\text{AI}} = \mu + \theta^{\text{AI}} \cdot q_{\tau(p_{it})}, \qquad \frac{\partial a^{\text{AI}}}{\partial p_{it}} = -\frac{\theta^{\text{AI}}}{(\phi+\psi)\, f_Z(q_\tau)} < 0. \tag{13}$$

Part (i) â Existence of price elasticity. From (9), $\partial\tau/\partial p = -1/(\phi+\psi) < 0$. Since $q_\tau = F_Z^{-1}(\tau)$ is strictly increasing, $\partial a^{\text{AI}}/\partial p < 0$, while $\partial a^{\text{pre}}/\partial p = 0$. AI strictly creates price elasticity. □

Part (ii) - Temporal reallocation. Define the reallocation gain as the energy saving in high-price periods relative to the flat pre-AI baseline:

$$\Delta E^H := \kappa\big(a^{\text{pre}} - a^{\text{AI}}(p^H)\big) = \kappa\Big[\theta^{\text{pre}}\, q_{\tau(\bar{p})} - \theta^{\text{AI}}\, q_{\tau(p^H)}\Big]. \tag{14}$$

Since $p^H > \bar{p}$ implies $\tau(p^H) < \tau(\bar{p})$, hence $q_{\tau(p^H)} < q_{\tau(\bar{p})}$, we decompose:

$$\Delta E^H = \underbrace{\kappa\,(\theta^{\text{pre}} - \theta^{\text{AI}})\, q_{\tau(\bar{p})}}_{\text{(a) uncertainty reduction}} + \underbrace{\kappa\,\theta^{\text{AI}}\,\big(q_{\tau(\bar{p})} - q_{\tau(p^H)}\big)}_{\text{(b) price response}} > 0. \tag{15}$$

Both terms are strictly positive: term (a) reflects the average energy saving from reduced forecast uncertainty (Implication 1); term (b) is the additional saving from active setpoint adjustment to real-time prices, which is zero pre-AI by construction and strictly positive post-AI. Hence:

$$\Delta E^H\big|_{\text{AI}} > \Delta E^H\big|_{\text{pre-AI}} = 0.$$

□

Remark. The price-response component (b) is independent of the level-effect channel, providing a clean empirical separation between the two mechanisms.

Implication 4: Dynamic Learning and Decreasing Returns

Prediction. If the AI model improves over time (θ_{it} decreasing in t), energy consumption declines gradually, exhibiting diminishing marginal gains.

Proof. Suppose $\theta_{it} = \theta_0 \cdot h(t)$, with $h'(t) < 0$ and $h''(t) > 0$ (concave learning). Then:

$$a^*_{it} = \mu_{it} + \theta_0 \cdot h(t) \cdot q_\tau \quad \Rightarrow \quad \frac{\partial a^*_{it}}{\partial t} = \theta_0 \cdot h'(t) \cdot q_\tau < 0,$$

and:

$$\frac{\partial^2 a^*_{it}}{\partial t^2} = \theta_0 \cdot h''(t) \cdot q_\tau > 0.$$

Thus, aeration (and energy) decrease over time at a decelerating rate. □

Implication 5: Reduction in Overshoot Magnitude

Prediction. AI adoption reduces the frequency and magnitude of overshoot events ($a_{it} > R_{it}$).

Proof. Define an overshoot event as $\mathbb{1}[a^*_{it} > R_{it}]$. From (8):

$$\Pr(a^*_{it} > R_{it} \mid \mathcal{I}_{it}) = \tau(p_{it}).$$

Ex ante, the frequency of overshoots is $\tau \approx \psi/(\phi + \psi)$, independent of θ (determined by loss asymmetry). However, ex post overshoot frequency depends on realized deviations.

Under location-scale (10):

$$a^*_{it} - R_{it} = \theta_{it} \cdot (q_\tau - Z).$$

The ex post probability that $a^*_{it} > R_{it}$ is: $\Pr(q_\tau > Z) = \tau$, unchanged by θ_{it}.

But the magnitude of overshoot, conditional on $a^* > R$, is:

$$\mathbb{E}[a^*_{it} - R_{it} \mid a^*_{it} > R_{it}] = \theta_{it} \cdot \mathbb{E}[q_\tau - Z \mid Z < q_\tau].$$

Since $\theta^{\text{AI}}_{it} < \theta^{\text{pre}}_{it}$, the conditional overshoot magnitude is smaller post-AI.

Interpretation: define overshoot relative to a realized threshold (e.g., Oxydo-Reduction Potential target). If monitoring is imperfect pre-AI, operators set higher margins; AI reduces this margin. Let $\bar{R}$ be a compliance threshold. Pre-AI, the operator may target $a^{\text{pre}} = \mu + \theta^{\text{pre}} q_\tau$; post-AI, $a^{\text{AI}} = \mu + \theta^{\text{AI}} q_\tau$. The deterministic overshoot is:

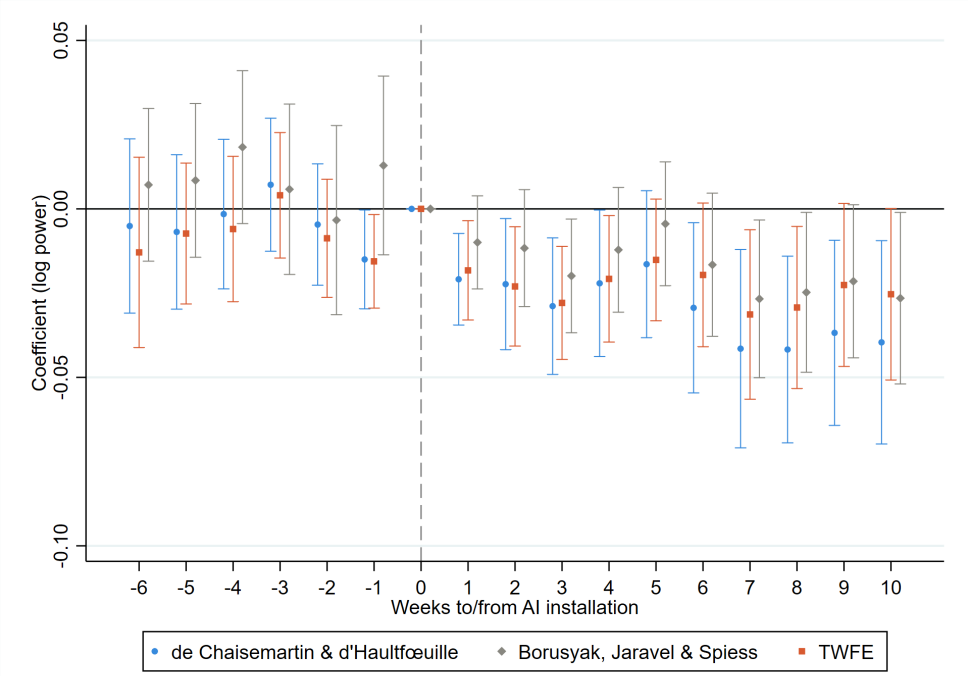
$$\Delta a := a^{\text{pre}} - a^{\text{AI}} = (\theta^{\text{pre}} - \theta^{\text{AI}}) \cdot q_\tau > 0.$$

Hence, AI reduces systematic over-provisioning, which is extremely consistent with our findings on the sharp reduction in ORP > 0 mV. □

G Robustness checks

G.1 Event studies using different DiD estimators

Figure G.1: Event studies, Electricity consumption



Notes. This figure presents the effect of AI deployment on the electricity consumption of WWTP using the robust dynamic difference-in-differences estimators introduced in Borusyak, Jaravel and Spiess (2024) and de Chaisemartin and D'Haultfoeuille (2024), and a standard two-way fixed effects regression model. Data are are the weekly level and cover a period up to 10 weeks after AI installation.

G.2 Keeping only exogenous failures of AI

Table G.1: Effects on weekly electricity consumption, carbon footprint and costs

	AI usage	Treated	Ever Treated
Electricity Consumption			
	-0.044***	-0.033***	-0.03***
	(0.014)	(0.011)	(0.01)
Carbon footprint			
	-0.049***	-0.037***	-0.033***
	(0.016)	(0.012)	(0.011)
Electricity cost			
	-0.069***	-0.052***	-0.048***
	(0.018)	(0.012)	(0.012)
Observations	10,544	10,544	10,544
Number of plants	164	164	164

Notes: We report the average cumulative total effect measured over 10 weeks using de Chaisemartin and D'Haultfoeuille (2024). Electricity consumption, carbon footprint, and electricity cost are expressed in log values. The definition of the measure of AI use differs across columns. Effect on electricity cost is obtained controlling for the mean individual weekly electricity price. Standard errors are clustered at the plant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

G.3 Effects of AI by adoption wave

Table G.2: Effects of AI adoption by wave

	First wave		Second wave	
	dCdH	TWFE	dCdH	TWFE
Average total effect	-0.060		-0.047**	
	(0.042)		(0.018)	
TWFE coefficient		-0.046**		-0.049***
		(0.015)		(0.012)
Placebo test (p-value)	0.347		0.618	
Joint effects test (p-value)	0.444		0.065	

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at the plant level in parentheses. dCdH refers to the de Chaisemartin and D'Haultfoeuille (2024) estimator. The dependent variable is log power consumption, and the independent variable is the AI usage share of time. Each wave subsample includes never-treated plants as controls.

G.4 Rebound Effect

Table G.3: Rebound Effect

	AI Pilotage Effect (Std. Error)	AI Implementation Effect (Std. Error)
Electricity consumption	-0.056***	0.016
	(0.010)	(0.011)
Carbon footprint	-0.055***	0.019*
	(0.008)	(0.011)
Electricity Cost	-0.047***	-0.009
	(0.010)	(0.015)

Notes. This table reports the estimated impact of AI usage and AI implementation on daily electricity consumption, carbon footprint and electricity cost expressed in log values at the plant level estimated with a TWFE specification with both Piloting and Implementation variables. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

G.5 Controlling for inflow

Table G.4: Average cumulative total effect per plant

	Estimate	Std. Error
Electricity consumption	-0.049***	0.017
Carbon footprint	-0.056***	0.019
Electric costs	-0.079***	0.019

Notes. This table reports the average cumulative total effect measured over 10 weeks estimated with de Chaisemartin and D'Haultfoeuille (2024) controlling for the inflow level. Effect on electricity cost is obtained controlling for the mean individual weekly electricity price. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

G.6 Without Never Treated from the counterfactual

Table G.5: Average cumulative total effect per plant

	Estimate	Std. Error
Electricity consumption	-0.045***	0.017
Carbon footprint	-0.05***	0.019
Electricity cost	-0.075***	0.019

Notes. This table reports the average cumulative total effect measured over 10 weeks estimated with de Chaisemartin and D'Haultfoeuille (2024) by removing plants that were never treated from the counterfactual group. Effect on electricity cost is obtained controlling for the mean individual weekly electricity price. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

G.7 Keeping only optimized control mode in the counterfactual

Table G.6: Average cumulative total effect per plant

	Estimate	Std. Error
Electricity consumption	-0.052***	0.017
Carbon footprint	-0.056***	0.025
Electricity cost	-0.078***	0.021

Notes. This table reports the average cumulative total effect measured over 4 weeks estimated with de Chaisemartin and D'Haultfoeuille (2024) by removing observations daily observations where the most important control mode was the P0 "local mode". Effect can only be observed through 4 weeks due to the lack counterfactual data. Effect on electricity cost is obtained controlling for the mean individual weekly electricity price. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table G.7: TWFE

	Estimate	Std. Error
Electricity consumption	-0.048***	0.009
Carbon footprint	-0.044***	0.009
Electricity cost	-0.061***	0.012

Notes. This table reports the effect measured with the standard TWFE estimator, by removing observations daily observations where the most important control mode was the P0 "local mode". Effect on electricity cost is obtained controlling for the mean individual weekly electricity price. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

H Heterogeneous AI impact

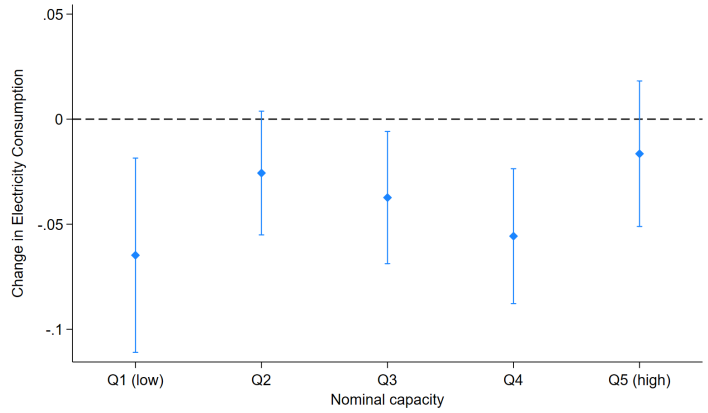
H.1 By plant

To explore cross-plant heterogeneity, we classify plants into quintiles based on their nominal capacity and average inflow size over the sample period. We the estimate the following TWFE regression intereacting AI use with these quintile indicators:

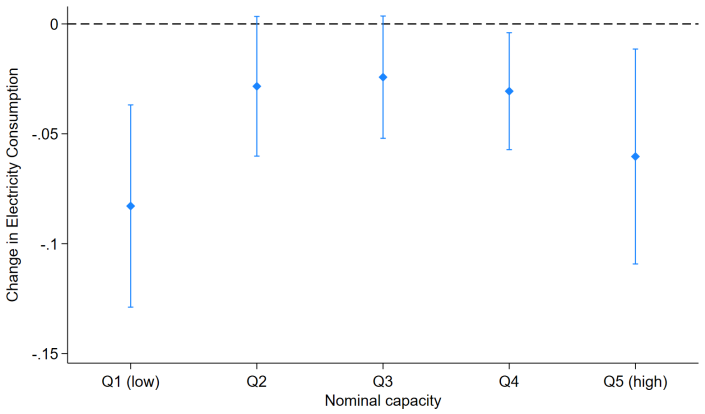
$$\ln(\text{Power}_{it}) = \sum_{q=1}^{5} \beta_q \left(\text{AI usage}_{it} \times \mathbb{1}\{i \in Q_q\}\right) + \alpha_i + \delta_t + \varepsilon_{it},$$

Figure H.1: AI impact by plant size - in relative value

(a) by Nominal Capacity



(b) by Average Inflow



Notes. This figure plots the coefficients and 95% confidence intervals of the interaction term between AI usage intensity and the quintile of nominal capacity, measured in People Equivalent (subfigure a), and average inflow (subfigure b) of each plant using a standard TWFE regression on the logarithm of electricity consumption.

H.2 By inflow

Table H.1: Heterogeneity by inflow: TWFE

	AI Usage	Highly Treated	Treated	Ever Treated
Electricity Consumption				
AI use	-0.112***	-0.100***	-0.098**	-0.087
	(0.040)	(0.038)	(0.039)	(0.053)
Inflow	0.059***	0.059***	0.059***	0.057***
	(0.012)	(0.012)	(0.012)	(0.014)
Inflow * AI use	0.016*	0.016*	0.016*	0.016
	(0.12)	(0.12)	(0.12)	(0.14)
Carbon footprint				
AI use	-0.114***	-0.100***	-0.098**	-0.085
	(0.039)	(0.037)	(0.038)	(0.052)
Inflow	0.059***	0.059***	0.058***	0.057***
	(0.012)	(0.012)	(0.012)	(0.014)
Inflow * AI use	0.016*	0.016*	0.016*	0.016
	(0.012)	(0.012)	(0.012)	(0.014)
Plant FE	170	170	170	170
Date-Hour FE	31'449	31'449	31'449	31'449
Observations	2'042'110	2'042'110	2'042'110	2'042'110

Note: Electricity Consumption, Carbon footprint and Inflows are expressed in log values. The definition of the measure of AI use differs indifferent columns. Fixed effects at the day and id levels and SE clustering at the id level are included. *p<0.1; **p<0.05; ***p<0.01. R^2 are at 79.7% for electricity consumption and at 86.6% for carbon emissions.

Table H.2: Intensity heterogeneity by inflow - TWFE

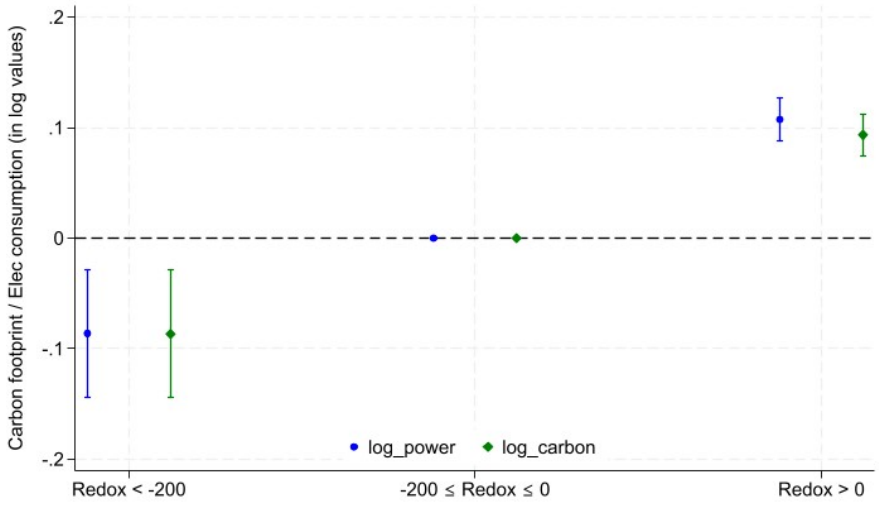
	Log(Electricity intensity)
AI usage	-0.19^{**}
	(0.079)
Log(inflow)	-1.06^{***}
	(0.09)
Log(inflow) * AI usage	0.034^{*}
	(0.018)
Observations	2,021,413
Number of plants	155

Notes: Electricity intensity refers to the electricity consumption per unit of inflow. The specification include plant and date fixed effects. Standard errors are clustered at the plant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

I Oxydo-Reduction Potential and AI

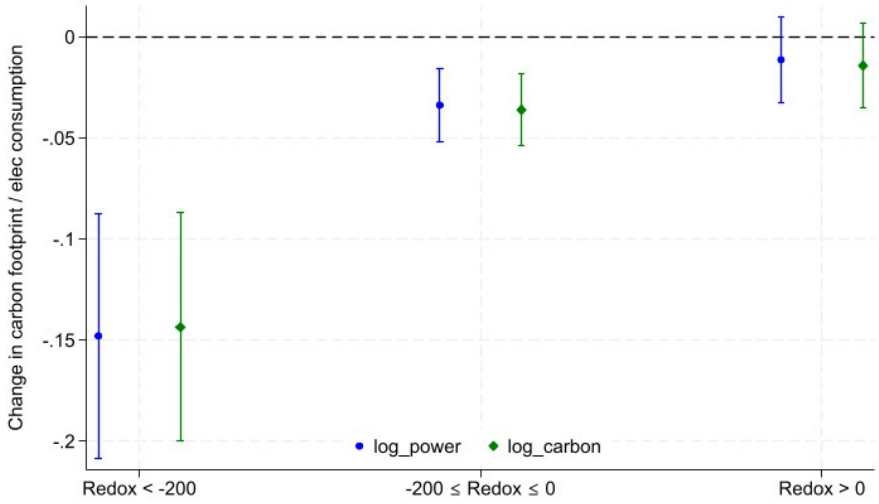
I.1 Empirical Results

Figure I.1: Heterogeneous baseline consumption depending on the ORP level



Notes. Electricity Consumption and Carbon footprint are expressed in log values. Redox between -200 and 0 mV is taken as the baseline. The definition of the measure of AI use differs indifferent columns. Fixed effects at the day and id levels and SE clustering at the id level are included. CI are at the 95% level.

Figure I.2: AI heterogeneous impact by Oxido-Reduction Potential (ORP) level



Notes. Electricity Consumption and Carbon footprint are expressed in log values. The definition of the measure of AI use differs in different columns. Fixed effects at the day and id levels and SE clustering at the id level are included. Confidence intervals are expressed at the 95 % level.

J Effluent Quality Analysis

J.1 The Nitrogen Removal Pathway

Nitrogen removal in wastewater treatment plants (WWTPs) relies on a two-stage biological process: nitrification followed by denitrification. These two processes are biochemically sequential yet operationally antagonistic, as they require fundamentally different redox conditions. Understanding their interplay is essential to rationalize both the levels of ammonium and nitrate observed in WWTP effluents and the channels through which AI-based process control can improve outcomes (Tchobanoglous and Eddy., 2014).

Nitrification. Nitrification is an aerobic, autotrophic process carried out by specialized bacteria that oxidize ammonium (NH_4^+) to nitrite (NO_2^-) and then to nitrate (NO_3^-):

$$NH_4^+ + \frac{3}{2}\,O_2 \xrightarrow{\textit{specialized bacteria}} NO_2^- + H_2O + 2\,H^+ \tag{16}$$

$$NO_2^- + \frac{1}{2}\,O_2 \xrightarrow{\textit{specialized bacteria}} NO_3^- \tag{17}$$

This process requires a sustained supply of dissolved oxygen (DO) and is sensitive to temperature, pH, and organic loading. The net reaction converts NH_4^+ into NO_3^-, so incomplete nitrification leaves residual ammonium in the effluent, while complete nitrification produces nitrate, which must in turn be removed by denitrification.

Denitrification. Denitrification is an anoxic, heterotrophic process by which facultative bacteria use nitrate as a terminal electron acceptor in the absence of oxygen, reducing it to dinitrogen gas:

$$NO_3^- \xrightarrow{-2e^-} NO_2^- \xrightarrow{-2e^-} NO \xrightarrow{-e^-} N_2O \xrightarrow{-e^-} N_2\uparrow \tag{18}$$

This process requires a carbon source (typically the organic matter present in the influent) and, critically, the absence of dissolved oxygen, which would otherwise be preferentially used as the electron acceptor, outcompeting nitrate and shutting down denitrification.

The nitrification-denitrification trade-off. The core operational challenge in biological nitrogen removal is that these two processes require opposite aeration conditions. Nitrification demands sustained oxygenation, while denitrification requires anoxic zones. In practice, WWTPs implement alternating aerobic and anoxic phases through intermittent aeration strategies to satisfy both requirements within the same treatment train. A key challenge for the system is to access a fine dynamic optimum: enough aeration to drive complete nitrification of incoming NH_4^+, but not so much as to suppress the anoxic zones needed for subsequent denitrification.

J.2 AI and Nitrogen Removal

Under low-to-moderate ammonium loads, conventional rule-based controllers can maintain this balance without great difficulty. The nitrification capacity of the sludge is not saturated, and the anoxic volume in the reactor is sufficient to handle the nitrate produced. Effluent concentrations of both NH_4^+ and NO_3^- remain within acceptable bounds with relatively coarse aeration control.

The challenge of High Ammonium Loads. The management of activated sludge systems under high ammonium loads is inherently difficult, as conventional control systems typically respond by increasing aeration to prevent ammonium breakthrough (Xiong et al., 2025). While increasing aeration is a rational local response to rising influent loads, it is often counterproductive at the system level because over-aeration extends the aerobic phase at the expense of the anoxic phase, thereby suppressing the denitrification process (Raboni et al., 2020). The inhibitory effect of dissolved oxygen (DO) on denitrification is extensively documented (Oh and Silverstein, 1999; Raboni et al., 2020). Under conditions of excess aeration, oxygen carries over through recirculated nitrified water, where it acts as a more energetically favorable electron acceptor than nitrate, effectively shutting down denitrification (Qin et al., 2023). This leads to a scenario where nitrate accumulates in the effluent even if nitrification remains incomplete, resulting in the simultaneous elevation of both ammonium and nitrate levels. Conventional DO feedback control has been shown to perform poorly under high and variable ammonium loading because it is responding only after oxygen depletion is observed, rather than anticipating nitrogen dynamics (Åmand, Olsson and Carlsson, 2013). In full-scale plants, this issue is exacerbated by chemical analysis delays that can reach several hours or days, forcing operators to rely on manual adjustments or conservative, energy-intensive strategies. These limitations have driven a growing body of literature focused on advanced aeration control and provide the operational context in which AI-based systems can yield significant performance gains and reductions in operational expenditure (Xiong et al., 2025; Croll et al., 2023; Chen et al., 2021).

The dynamic aeration optimization channel. AI-based process control can be expected to improve effluent nitrogen quality through dynamic aeration optimization. Rather than responding reactively to a single sensor reading, AI controllers can integrate information from multiple sensors, hence AI controllers can learn the optimal phase durations as a function of real-time influent characteristics, temperature, and sludge activity, adapting dynamically rather than following fixed time-based cycles. It can also have a pro-active rather than simply reactive attitude toward nitrogen loads to anticipate loading variations and adjust aeration setpoints ahead of time. This allows the system to maintain the aerobic-anoxic balance across a wider range of operating conditions, including high-load episodes during which conventional controllers fail.

J.3 Empirical Results

The empirical findings reported below are consistent with this biochemical picture. The negative effect of AI adoption on both NH_4^+ and NO_3^- concentrations in effluents is the expected signature

of improved joint optimization of the nitrification-denitrification cycle: AI reduces ammonium by improving nitrification completeness, and simultaneously reduces nitrate by preserving the anoxic conditions necessary for denitrification outcomes that are mechanistically linked rather than independent. The concentration of this effect in high-ammonium regimes further corroborates the mechanism: it is precisely when conventional controllers over-aerate and compromise denitrification that AI-based dynamic optimization yields the largest gains.

Table J.1: Effect of AI pilotage on effluent quality

	(1) Nitrate	(2) Ammonium	(3) Phosphate
AI usage	-0.508***	-0.292**	-0.030
	(0.114)	(0.117)	(0.031)
Constant	1.977***	1.738***	0.616***
	(0.049)	(0.051)	(0.012)
Plant FE	147	148	74
Date FE	969	969	604
Observations	20,269	20,636	8,484

Notes. This table presents the effect of AI usage on Nitrate (NO_3^-), Ammonium (NH_4^+), and Phosphate (PO_4^{3-}) concentrations using a standard two-way fixed effects regression model. Plant and date fixed effects are included in all specifications. Standard errors appear in parentheses and are clustered at plant level. Observations are at the day level. $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table J.2: Impact of AI on compliance with rejection limits

	(1) NH_4^+ compliance	(2) PO_4^{3-} compliance
AI usage	0.585***	-0.093
	(0.146)	(0.147)
Plant FE	116	52
Observations	15,754	5,713

Notes. This table reports estimates from conditional (fixed-effects) logistic regressions. Coefficients are expressed in log-odds. The dependent variable is a dummy equal to one when the samples respect the quality thresholds. Standard errors are obtained by bootstrap with 200 replications and are clustered at the plant level. The estimation relies on within-plant variation only. Plants with no within-plant variation in the compliance status are not included in the estimation. Observations are at the day level. $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table J.3: Quantile regressions of ammonium concentration on AI usage

Quantile	.1	.2	.3	.4	.5	.6	.7	.8	.9
AI usage	0.29***	0.24***	0.19***	0.15***	0.11**	0.05	-0.06	-0.29*	-0.71**
	(0.04)	(0.04)	(0.04)	(0.051)	(0.05)	(0.06)	(0.09)	(0.16)	(0.28)
Observations	20,815								

Notes. This table presents the effect of AI usage on ammonium (NH_4^+) concentrations using quantile regressions with plant fixed effects using the estimator of Machado and Silva (2019). Bootstrap standard errors clustered at plant level appear in parentheses. Observations are at the day level. $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

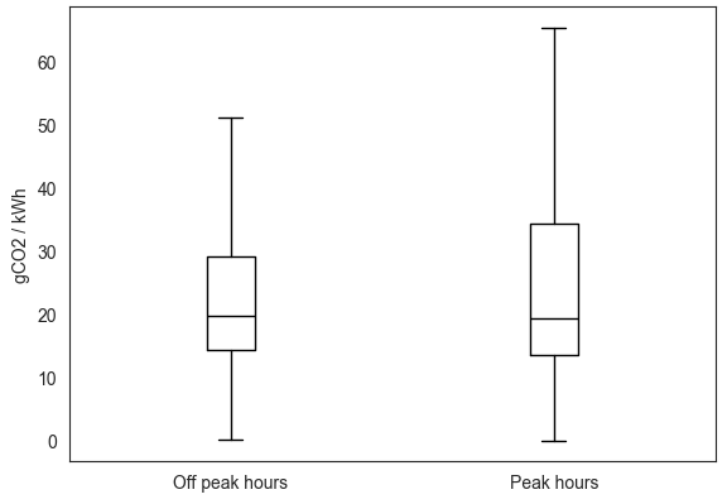
Table J.4: Effect of AI depending on ammonium concentration heterogeneity

	Dependent variable		
Ammonium decile	Nitrate (mg/L)	Redox (mV)	Power (kW/h)
D1 (low)	0.184	-7.810	-0.440
	(0.229)	(13.458)	(1.580)
D2	0.199	-17.802**	-1.465
	(0.232)	(8.949)	(1.966)
D3	-0.326*	-20.125***	-0.230
	(0.171)	(6.676)	(1.692)
D4	-0.599**	-31.411***	-1.655
	(0.236)	(8.239)	(1.233)
D5	-0.542***	-39.733***	-2.476**
	(0.143)	(7.148)	(1.088)
D6	-0.622***	-42.749***	-2.057*
	(0.192)	(8.309)	(1.185)
D7	-0.615***	-36.095***	-2.834**
	(0.147)	(5.365)	(1.430)
D8	-0.597***	-37.637***	-1.877
	(0.168)	(6.747)	(1.392)
D9	-0.872***	-39.411***	-2.613**
	(0.144)	(6.290)	(1.168)
D10 (high)	-1.271***	-57.903***	-2.149
	(0.167)	(6.878)	(1.517)
R^2	0.225	0.486	0.884
Plant FE	147	143	148
Date FE	966	956	969
Observations	20,024	19,009	20,636

Notes. This table reports the marginal effect of AI usage on Nitrate concentration (mg/L), redox (mV) and Power (kWh) depending on the concentration of ammonium. We use a standard two-way fixed effects regression, and report the coefficients of the interaction terms between AI usage and ammonium concentrations deciles. Clustered standard errors at the plant level are in parentheses. Observations are at the day level. $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

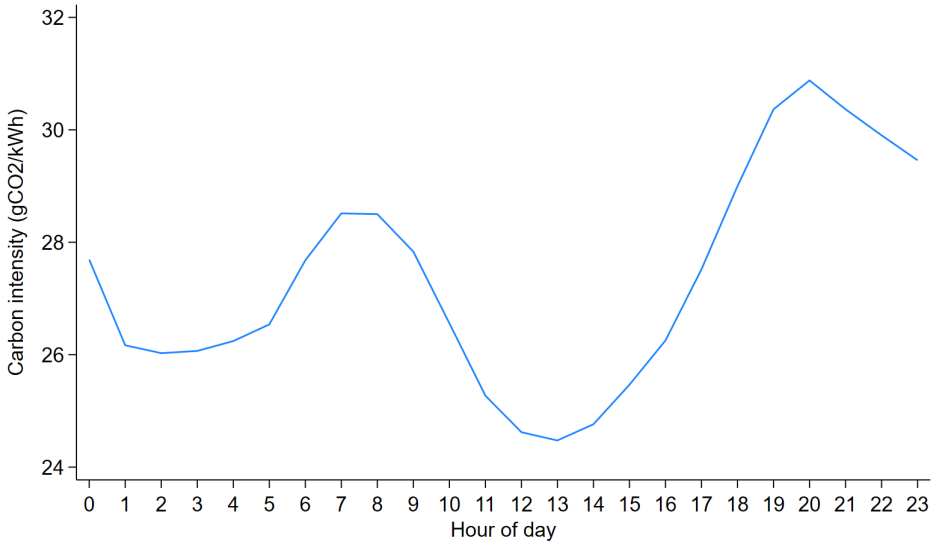
K Power reallocation between peak and off-peak hours

Figure K.1: Carbon intensity, Peak and Off-peak hours



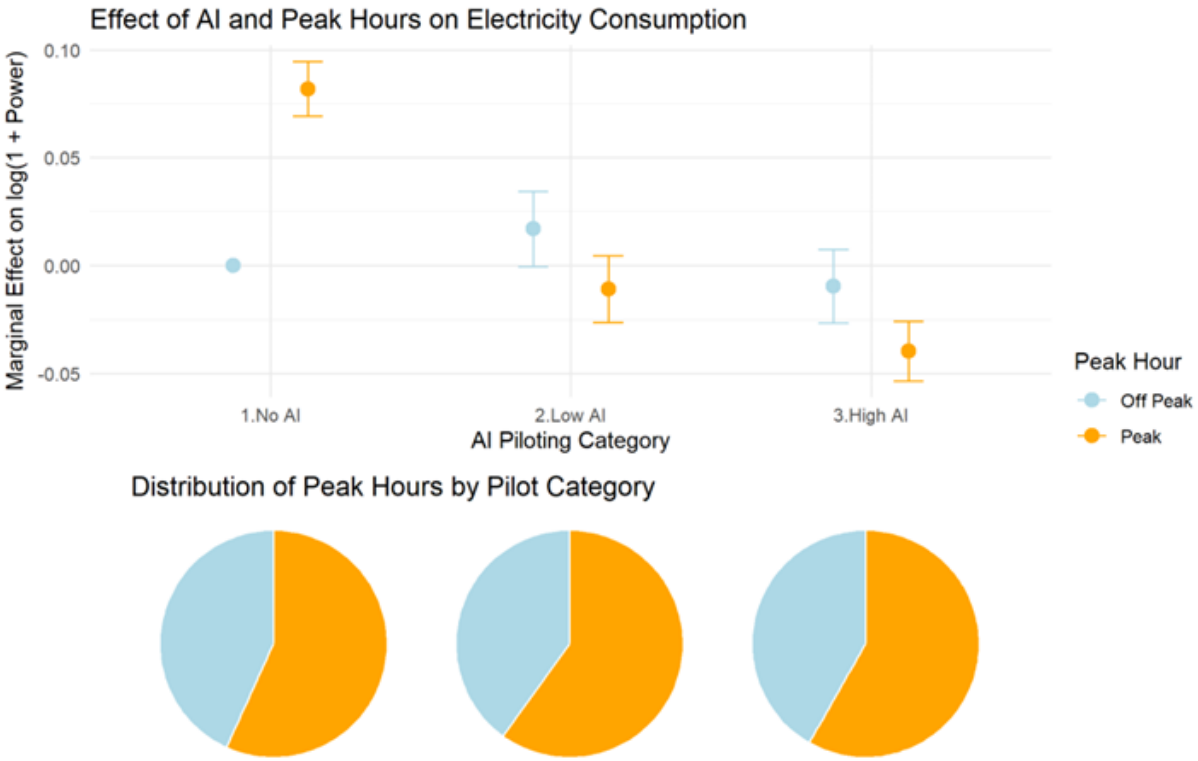
Notes. This figure shows boxplots of hourly carbon intensity (CO2 equivalent emissions in g/kWh) for off-peak and peak hours. Peak hours are hours from 5 AM to 10 PM with a higher electricity price compared to off-peak hours.

Figure K.2: Carbon intensity during the day



Notes.

Figure K.3: TWFE: Peak vs. Off-peak Hours

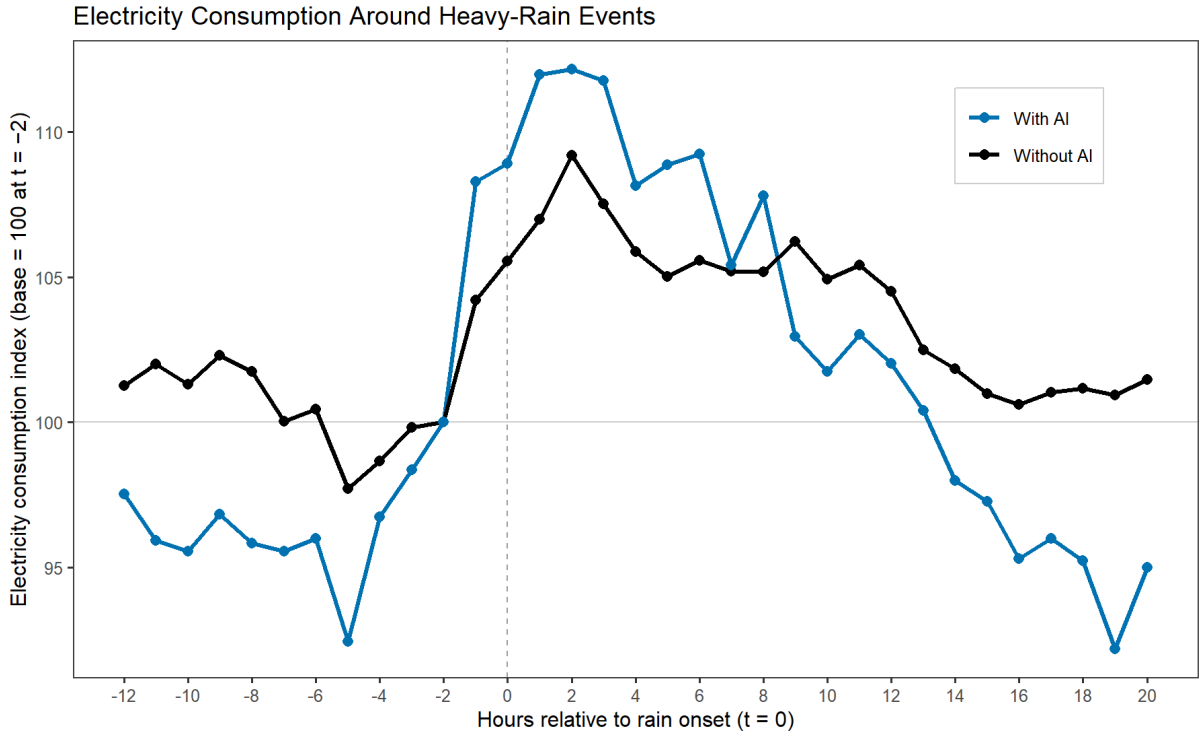


Notes. This figure shows boxplots of hourly carbon intensity (CO2 equivalent emissions in g/kWh) for off-peak and peak hours.

L Extreme Meteorological Events

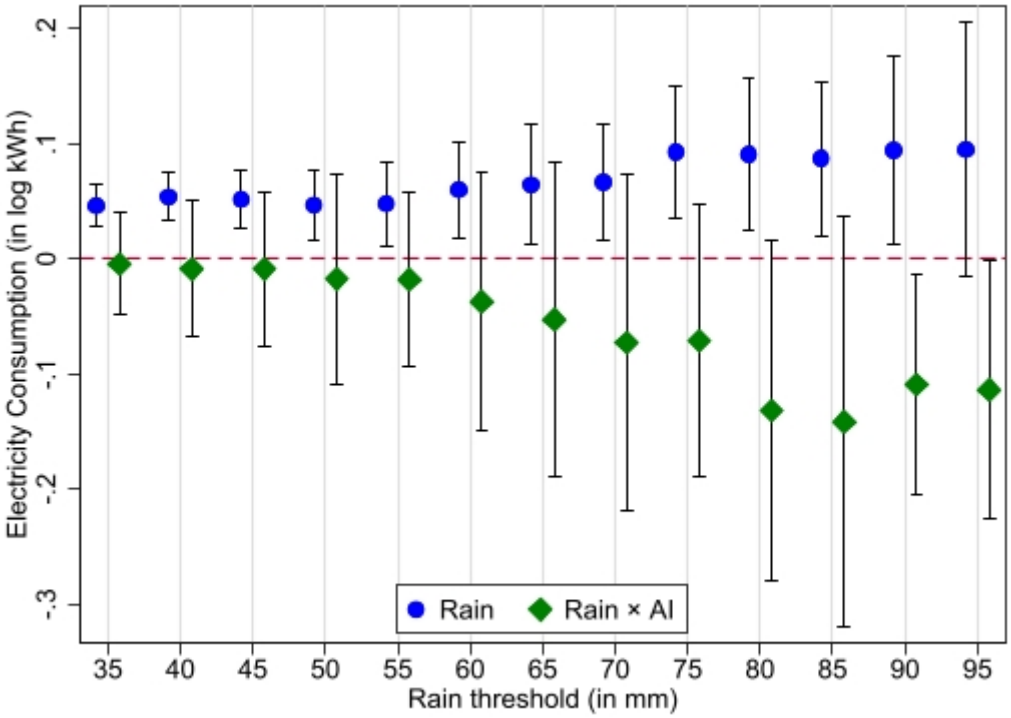
L.1 Consumption Profiles With and Without AI

Figure L.1: Electricity Consumption Around Heavy-Rain Events: AI vs. Non-AI Plants



Notes. This figure plots average electricity consumption around rain onset events (30-55 mm/h), separately for plants operating with and without AI monitoring. Each event is normalized to 100 at $t = -2$ hours before the high rainfall event. Residuals are purged of plant × hour-of-day, plant × day-of-week, and plant × rain-bin fixed effects to control for systematic differences in consumption patterns and concurrent rainfall intensity. "With AI" denotes plants for which more than 80% of hours in the 30 days prior to the event were under AI inference mode (P2); "Without AI" denotes plants with no AI exposure.

L.2 Intense Rainy Episode at the day-level



Thr.	35	40	45	50	55	60	65	70	75	80	85	90	95
IPW	10^2	10^2	10^2	10^2	10^2	10^3	10^3	10^3	10^3	10^3	10^3	10^3	10^3

Figure L.2: Hourly Concentration Volume (hour-level)

References

Åmand, Lena, Gustaf Olsson, and Bengt Carlsson. 2013. "Aeration Control - A Review." *Water Science & Technology*, 67(11): 2374–2398.

Blackwell, David. 1953. "Equivalent Comparisons of Experiments." *The Annals of Mathematical Statistics*, 24(2): 265–272.

CGDD. 2019. "L'assainissement collectif des eaux usées." *notre-environnement.gouv.fr*, Accessed: March 2026.

Chen, Kehua, Hongcheng Wang, Borja Valverde-Pérez, Siyuan Zhai, Luca Vezzaro, and Aijie Wang. 2021. "Optimal control towards sustainable wastewater treatment plants based on multi-agent reinforcement learning." *Chemosphere*, 279: 130498.

Croll, Henry C., Kaoru Ikuma, Say Kee Ong, and Soumik Sarkar. 2023. "Reinforcement learning applied to wastewater treatment process control optimization: Approaches, challenges, and path forward."

de Chaisemartin, Clément, and Xavier D'Haultfoeuille. 2024. "Difference-in-Differences Estimators of Intertemporal Treatment Effects." *The Review of Economics and Statistics*, 1–45.

Oh, Jae-Seung, and JoAnn Silverstein. 1999. "Oxygen Inhibition of Activated Sludge Denitrification." *Water Research*, 33(8): 1925–1937.

Qin, Lu, Dongqi Wang, Zhe Zhang, Xiaoxiao Li, Guodong Chai, Yishan Lin, Cong Liu, Rui Cao, Yuxin Song, Haiyu Meng, Zhe Wang, Hui Wang, Chunbo Jiang, Yuan Guo, Jiake Li, and Xing Zheng. 2023. "Impact of Dissolved Oxygen on the Performance and Microbial Dynamics in Side-Stream Activated Sludge Hydrolysis Process." *Water*, 15(11).

Raboni, Massimo, Paolo Viotti, Elena Cristina Rada, Fabio Conti, and Maria Rosaria Boni. 2020. "The Sensitivity of a Specific Denitrification Rate under the Dissolved Oxygen Pressure." *International Journal of Environmental Research and Public Health*, 17(24): 9366.

Tchobanoglous, George., and Metcalf Eddy. 2014. *Wastewater engineering treatment and resource recovery. McGraw-Hill series in civil and environmental engineering.* Fifth edition. ed., New York, NY:McGraw-Hill Education.

Xiong, Zhaoyang, Leiming Cheng, Ahmed Ibrahim Yunus, Dewu Jian, Xinping Zhang, Thomas Igou, Sha Tong, and Yongsheng Chen. 2025. "Machine Learning-Assisted Expert Control in Wastewater Aeration Treatment Process Application." *ACS ESamp;T Engineering*, 5(11): 3140â3152.