**Ivy Jiang**

University of Pennsylvania

# Techno-Economic Optimization on Grid-Connected Electrolytic Ammonia Production

**RWTH Aachen University**
***SVT***
***Michael Rix***

**UROP – Undergraduate Research Opportunities Program**

**Aachen**
***21 July 2025***

## Introduction

Ammonia ($NH_3$) is a cornerstone of the chemical industry; not only is it critical for fertilizers feeding the global population, but it is also being explored as a carbon-free energy carrier and fuel. However, conventional ammonia synthesis via the Haber–Bosch process, coupled with fossil-based hydrogen production, is responsible for roughly 2% of global $CO_2$ emissions. Decarbonizing ammonia production is thus a key goal. Green ammonia production powered by renewable electricity is technologically feasible, but faces significant economic and operational challenges.

Against this backdrop, this report examines two system paradigms for electrochemical ammonia production using eNRR: 1. an off-grid, renewable-powered system sized to meet a fixed $NH_3$ output using only on-site wind/PV power, and 2. a grid-connected system that can dynamically purchase electricity from the grid (assumed to be partly renewable) to minimize cost.

We develop a Pyomo-based model that optimally designs and schedules a flexible ammonia electrolyzer system under a grid connected power supply mode. This is a study that builds on two recent master's theses that developed an optimization model for eNRR based ammonia (Bidaoui, 2024 and Schwindling, 2023). The objective is to meet a fixed annual $NH_3$ production target at minimal cost by balancing capital expenditure (CAPEX) on the electrolyzer against operating costs (OPEX) for electricity.

We compare this off-grid renewable scenario (with on-site wind/solar and curtailment of excess energy) to a grid-connected scenario (purchasing electricity from the wholesale market). Our findings show that a grid-connected electrolyzer can produce ammonia at slightly lower cost than an off-grid renewables-based system under 2023 price conditions, but only if ample low-cost power is available. Lying at the crux is the emphasis on operational flexibility; in this case, meaning the ability to ramp production up and down freely. It allows the system to exploit periods of abundant cheap energy and avoid running during unfavorable times. This flexibility dramatically reduces the LCOA by decoupling plant utilization from the intermittency of energy supply. The results also highlight how declining electricity prices (from 2023 to 2024) translate to lower ammonia costs, as the electrolyzer increases its per year utilization at full load. Overall, this work illustrates the promise in green ammonia production and provides insight into optimal design and operation strategies that could enable economically viable, sustainable ammonia in the future.

## Current Research

### Electrochemical Nitrogen Reduction Reaction

A promising alternative for ammonia production without the use of fossil fuels is the electrochemical nitrogen reduction reaction (eNRR). In this process, ammonia is directly produced from N2 (nitrogen) and H2O (water) in an electrochemical cell, which is powered by electricity (preferably by renewable electricity). Without necessitating H2 to be taken in directly (like in Haber Bosch), this process operates at mild temperatures and pressures and emits no $CO_2$ on-site when using renewable energy. eNRR offers significant environmental advantages. It avoids a massive greenhouse gas footprint, when compared to HB and can consume less energy by bypassing an intermediate step of hydrogen production.

### Significance and Limitations

Showing significant theoretical promise, eNRR is still far from techno-economic feasibility under current laboratory conditions, especially when compared to the economic feasibility of Haber Bosch. As emphasized by recent early-stage modeling efforts, including Rix et al., significant performance improvements in catalyst design and electrochemical parameters are required for eNRR to level with established ammonia synthesis methods such as Haber-Bosch. Two performance metrics, Faradaic efficiency (η_FE) and current density (j), are particularly important. The levelized cost of ammonia (LCOA) is highly sensitive to these parameters, given their direct influence on both energy consumption and system sizing.

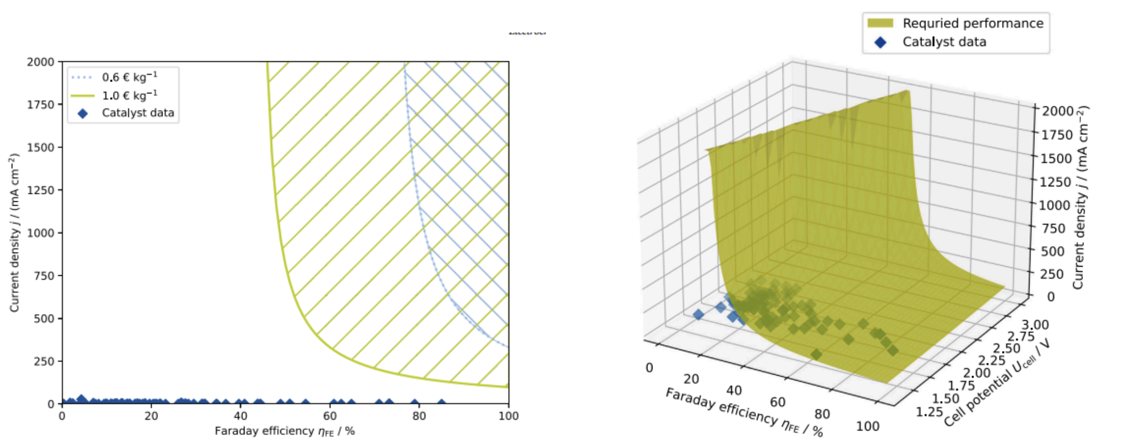


Figure 1: Required Performance vs. Current Performance [1]

Figure 1 demonstrates this dependency by mapping LCOA curves as a function of η_FE and j. The current current catalyst demonstrations (blue diamonds) lie overwhelmingly in a region orders of magnitude below the required performance for economic viability (looks to be around η_FE < 40% and j < 20 mA $cm^{-2}$). In contrast, achieving an LCOA below 1 €/kg $NH_3$ demands operating points exceeding 60% Faradaic efficiency and at least 500 mA $cm^{-2}$. The situation is visualized more robustly in the 3D visualization, which overlays current experimental data with a performance surface constrained by said economic targets. Here, the required zone shifts sharply toward higher current densities and moderate cell potentials, reinforcing the notion that both energy efficiency and rate capability must be improved in tandem.

These findings from Rix and others [1] underscore the significant research gap facing eNRR: there no existing catalyst simultaneously achieves high Faradaic efficiency, industrially relevant current densities, and stable operation at low overpotentials. Consequently, techno-economic modeling serves not only as a screen to test and retest feasibility but as a directional tool. Realistic benchmarks for catalyst development can be set and emphasis on which parameters become clear. All in all, dynamic and flexible system simulations, such as those that evaluate operation under real grid pricing or renewable fluctuations, become vital to assess whether grand bounds in future improvements in eNRR performance are feasible, translating the future of renewable ammonia production systems to be economically competitive and scalable.

**Renewable Energy-Powered Electrolyzer Model- Optimization of Design and Scheduling**

Emerging as a compelling concept in recent research has been the production of "green ammonia," further reducing the use of fossil fuels. This research builds off of an existing model [2] which ties the eNRR process to renewable energy sources. The aim is a renewable energy-powered electrolyzer model where dedicated wind and/or solar generators supply electricity to an ammonia electrolyzer in an isolated, off-grid system. However, designing such a system is complex because renewable power is intermittent, thus the available electricity varies hourly and seasonally. To address this, modeling and optimization frameworks are employed that coordinate the electrolyzer operation with the renewable power supply by managing the demand side. Operating under the assumption that the electrolyzer's response to changes in renewable energy supply is relatively fast, the time scale is modeled at hourly timesteps. In other words, the eNRR electrolyzer can ramp up or down almost instantaneously to match the input power, allowing the model robust flexibility based on energy availability or price [3].

In the renewable electrolyzer model, time-series data for one year (8760 hours) are used either from historical weather such as wind speeds and PV data. This includes the sizing of renewable generation assets and must ensure sufficient energy over the year to meet a target ammonia output. Energy supply is the driver of the electrolyzer's production through a set of constraints and decisions that ensure the required ammonia is produced at minimal cost. Specifically, the model is formulated as a mixed techno-economic optimization, simultaneously deciding the design (sizing of equipment) and the operation schedule (hourly production levels) to minimize a cost-based objective while meeting the annual production target.

The core of the electrolyzer model is based on fundamental electrochemical relationships. Though not the focus of this research project, these relationships govern the development of the original renewables model. Ammonia production rate (kg/h) is proportional to the total current passing through the cell dedicated to the NRR, which in turn equals the current density (A/m²) multiplied by the electrode active area (m²) and the Faradaic efficiency (since FE is the fraction of current going to ammonia). Thus, for a given catalyst performance (fixed FE and maximal current density), the electrolyzer area must be chosen such that at peak operation it can process enough current to achieve the desired ammonia output over the year. The renewable generation model provides the available power input each hour. Key parameters like the NRR Faradaic efficiency, current density limits, and cell voltage are set based on either current technology or future projections, since these have a strong influence on the outcomes (they determine how much area and energy are needed).

**Optimization of Design and Scheduling**

The renewable ammonia system is optimized to minimize the cost of ammonia production while meeting a specified production volume. This entails simultaneously determining the optimal design (equipment sizes) and the optimal operational strategy (scheduling of the electrolyzer). In mathematical terms, it is formulated as a constrained optimization problem (mixed-integer linear program after appropriate approximations). The objective function represents the levelized cost of ammonia (LCOA), which includes capital and operating costs. Below we outline the key decision variables, operational variables, and the objective function for this original renewables model:

**Key Decision Variables (Design)**

- Electrolyzer Size: This determines the maximum possible ammonia production rate. A larger electrode area allows more current and thus higher production, thus incurring a higher capital cost. The optimizer chooses the electrolyzer area to balance capacity vs. cost, constrained by a minimum size.

- Renewable Generation Capacity: In this model, the capacity of wind turbines and solar PV are decision variables. The model can invest in a certain kW of wind and kW of PV, up to some upper bound, to supply to the electrolyzer. These variables affect both capital cost (through the cost of installing physical renewables) and the energy available each hour (because generation = capacity × hourly capacity factor).

- Energy Storage Capacity: In the third mode, there is an extended model with a battery (BESS). The battery energy capacity (and possibly power capacity) is a design variable. Adding storage increases capital cost but provides more flexibility in operation, potentially allowing a smaller electrolyzer or less curtailment of renewables, reducing the volume of wasted generated energy.

**Key Operational Variables (Scheduling)**

- Electrolyzer Power Input / Production Rate: For each time step (hour), the model decides how much power the electrolyzer draws, and consequently how much ammonia is produced in that hour. This is represented as a fractional load factor, not ramping up to full capacity. It ranges from 0 (turned off) up to a maximum corresponding to the chosen maximum electrolyzer size and performance limits. These variables allow the electrolyzer to ramp down during low power availability and ramp up when abundant energy is generated from sunny or windy periods.

- Renewable Power Generation Used vs. Curtailed: The available renewable generation at each time step (given the installed capacities) can either be sent to the electrolyzer or curtailed. The model includes variables for the power generated by wind and PV that is actually utilized as well as the remainder is curtailed if generation exceeds what the electrolyzer (or battery) can handle. Curtailment is a valuable operational option that smartly avoids overloading the electrolyzer during peak renewable output and cuts down on costs.

- Battery Charge/Discharge Rates: If a battery is present, variables represent how much power is stored in the battery or drawn from the battery at each hour. The optimization will charge the battery during periods with surplus renewable energy and discharge it when the input from renewables is insufficient, subject to the battery's efficiency and capacity constraints. This smooths the power supply to the electrolyzer.

**Minimization of Objective Function**

The optimization's objective is to minimize the total annual cost of the system, thereby minimizing the LCOA. This objective function includes all capital expenditure costs. The CAPEX is the amortised and annualized cost of the electrolyzer, renewable generators, and any storage. Capital costs are converted to an equivalent annual cost using a capital recovery factor cost over the electrolyzer lifetime. Fixed O&M costs are a fraction of capital cost per year, significantly less than the CAPEX.

The model will incur higher CAPEX for larger capacities, so the optimizer tends to limit oversizing.

In summary, the objective function can be viewed as minimizing annualized total cost ($CAPEX_{El}$ + $CAPEX_{Wind}$ + $CAPEX_{PV}$). Since the annual ammonia output is fixed, this is equivalent to minimizing the LCOA in €/ton $NH_3$. By solving this optimization, researchers can determine the design that is least costly and the ideal operating schedule for a given intermittent period.

### Model Extensions

There are two extensions we will be referencing that improve the efficiency of energy usage from the electrolyzer: curtailment and energy storage.

- Curtailment Model: The base model assumes all available renewable energy must be used by the electrolyzer, which can force the electrolyzer to be very large to accommodate rare peak production hours. In the curtailed model, this constraint is discharged. To prevent oversized physical electrolyzer size, it is allowed to bypass peak hours and excess generated electricity by. As a result, the system can avoid high costs incurred by unnecessary capacity. The curtailed model finds an optimal balance between using renewable energy and spilling some of it, such that the annual production target is still met but the capital cost is lower.
- Battery Model: Another extension is adding a battery energy storage system (BESS) to store instead of dump this energy, store this surplus renewable electricity and supply it during deficits. The battery model builds on the curtailed model by providing a buffer for intermittency. During peak hours, such as overly sunny or windy periods, excess power can charge the battery and in periods when the renewable output is low, the stored energy from earlier can be discharged to keep the electrolyzer running at higher capacity than before. This smooths operation and increases the utilization of both the renewables and the electrolyzer. Smooth operation is desired because when the electrolyzer is utilized at max capacity for as often as possible, it produces ammonia more cheaply: the capital cost of the electrolyzer is amortised over more output. In our case, since the output is fixed, a smaller electrolyzer is needed to reach this same output. However, batteries incur their own capital costs and currently, storage costs are high, thus adding a battery may not be economically justified in all cases.

## Methodology/Conducted Research

Two components in electrochemical ammonia production contribute to the levelized cost of ammonia (LCOA): sourcing electricity and an electrolyzer that consumes the electricity. The goal is to build a Pyomo model that minimizes the LCOA, constrained by a fixed NH3 output. A piecewise-linear fit of $NH_3$ production rate vs. power input allows for computational efficiency while capturing load-dependent efficiency.

**Pyomo Optimization Model**

To effectively minimize the highly fluctuating costs that come with electrochemical ammonia production, we developed a conclusive optimization model using Pyomo. The model represents a single ammonia production facility set within a one year time frame at the hourly timestep (8760 hours). Two components contribute to the total cost. First is the electrolyzer system that does the physical ammonia synthesizing. Next is the electricity supply- whether it comes from renewable sources or the electric grid.

The high-level goal of the optimization is to minimize the levelized cost of ammonia (LCOA) while meeting a fixed annual production target. To be exact, by “levelized cost” we effectively mean the total annualized cost divided by annual $NH_3$ output, so minimizing total cost for a constant output is equivalent to minimizing LCOA.

A core feature of the model is a piecewise-linear production function for the electrolyzer. Already existing in the renewables model, the same fit model is used for the grid case. Rather than assume a constant efficiency, variable load factors may alter the efficiency to convert power to ammonia. At partial loads the system may, for instance, consume more kWh per kg of $NH_3$. To capture this without complex, nonlinear equations, the $NH_3$ production rate is approximated as a piecewise-linear function of the input power. This was done by fitting linear segments to data (or performance curves) relating power input to ammonia output. Essentially, the model can choose to operate the electrolyzer at various discrete load capacities , each with marginally better efficiency, allowing for a rough representation of diminishing returns at high load or efficiency drop at low load.

In order to improve the solver’s efficiency and performance all variables were scaled to unit ranges between 0 and 1. For example, power variables were normalized by its maximum possible value, and costs were normalized by a large reference cost. Practically, it helps avoid numerical issues in the solver (Gurobi).

The Pyomo model defines decision variables for the electrolyzer capacity and for the hourly operating power draw, which will be expanded further in the changed variables section. The key constraint that ties everything together is the annual production requirement: the sum of ammonia produced each hour (power input * efficiency in that hour) must equal the target annual production

**Grid Connected Model**

The primary new research conducted was transitioning the already existing ammonia plant model from an off-grid, renewables-powered system to a grid-connected system. In the grid-connected model, the electrolyzer purchases all its electricity from the public grid at prevailing market prices, instead of being tied to on-site generation. This paradigm shift requires slight changes to the variables and components to the function minimized, but remains similar in the fundamental model. The annual production requirement remained the same, and the overall goal of minimizing LCOA remained. However, what has changed includes the decision (design and operational) variables and constraints related to energy supply that we will outline shortly.

Because grid electricity acts as an infinite sink of electricity and operates under no supply constraints, the model's task boils down to when to buy and use electricity, and how much electrolyzer capacity to invest in. Once also blurred by the CAPEX of renewables capacities, the only capital expenditure is the electrolyzer. Intuitively, if electricity is cheap at a certain time, the model will favor running the electrolyzer then; if electricity is expensive, the model may choose to idle the electrolyzer to avoid high costs. The flexibility to turn down or shut off during high-price periods is key to minimizing OPEX. Importantly, since the annual $NH_3$ target is fixed, any hours of not producing will have to be compensated by producing more in other hours. This creates a trade-off: invest in a larger electrolyzer so that it can produce ammonia in fewer total hours (focusing only on the cheapest hours), or run more hours including moderately priced periods with a smaller electrolyzer. The optimizer evaluates this trade-off endogenously.

**Model Changes**

Fundamentally, the model remains a Pyomo model focused on minimizing LCOA. However some factors change to minimize LCOA in a grid scenario, where electricity is not generated in house. This first important variable remains. is $CAP_{El}$: Electrolyzer capacity (MW). This is a continuous variable. The model will choose an optimal size for the electrolyzer (scaling up capacity reduces the time needed to meet production but increases CAPEX). This next variable is a new added operational variable: $P_{grid}[t]$. The Grid power intake (MW or kW) each hour t = 8760 are continuous variables bounded between 0 and X. If $P_{grid}[t]=0$, it means the electrolyzer is off that hour. If $P_{grid}[t] = CAP_{el}$, the electrolyzer is running at full power. It can also take intermediate values if partially loaded and deemed optimal.

We retained the piecewise linear variables representing how much ammonia is produced as a function of power input. In practice, for each segment of the efficiency curve, a binary variable might indicate if that segment is active. The production constraint integrates $P_{grid}[t]$ with these segments to compute $NH_3$ output.

Various variables were also removed. In this mode, there are no wind or PV capacity variables, no curtailment variables, and no storage as it is not generating its own energy. The grid effectively acts as an infinite source. This simplification focuses the problem on the grid price variability aspect.

A key constraint is added. The energy balance constraint in off-grid mode (which said $P_{el}[t] \leq P_{wind}[t] + P_{PV}[t]$+ plus optional storage terms) was replaced by $P_{grid}[t]=P_{el}[t]$. In other words, all power consumed by the electrolyzer comes from the grid and all power imported is consumed. The annual production constraint still sums $P_{el}[t]$ times efficiency to equal the required $NH_3$ tonnage.

## New Minimized Objective Function

The new minimized objective function turns to:

$$min\,\{CAPEX_{El} + \sum_{t=1}^{T}(p_t * P_{grid,\,t})\}$$

This is equivalent to the representation of $CAPEX_{El}$ + $OPEX_{Elec}$

There are no costs for fuel or carbon since we assume grid electricity price already internalizes any fuel cost, and we are focusing on economic cost. The price data for each hour *t* is an input. In our study we took actual historical hourly prices for the German day-ahead market for the year 2023 and 2024. The year 2023 had generally higher prices (on average), whereas 2024 saw a reduction in price levels. This provided an interesting comparison for the model's behavior, analyzed in the results section.

Because the objective is linear in $p_t$ * $P_{grid}$ (with the coefficient being the price), avoiding higher prices is intrinsic and natural behavior of the optimizer. Essentially, for each hour, the decision to run or not run the electrolyzer becomes a question of whether the electricity cost is "worth it" in that hour in terms of producing $NH_3$. If an hour's price is extremely high (say 300 EUR/MWh), the model can skip that hour and make up production in some other cheaper hour, as long as capacity allows. The price of the production constraint will indicate the marginal cost of ammonia, and the model will equate that to running at times when electricity price is below that threshold. All in all, the electrolyzer acts as a price-responsive load.

## Solver

With the old model taking in weather data from four locations, including data on wind and PV, the new grid connected model runs on price data. Two separate cases from the grid-connected optimization were analyzed: 2023 hourly price data and 2024 data. The model was solved using Gurobi for each case, yielding an optimal electrolyzer capacity and an hourly operation schedule. New data was loaded in:

- Hourly Electricity prices: read from CSV files (strompreis_und_co2-emissionen_2023.csv and ..._2024.csv). These files provided hourly price data for electricity as well as their timestamp. We used the "Price in EUR/MWh" column, divided by 1000 to convert to EUR/kWh for internal calculations. For each run depending on the year, the input data is read from these files.

- Electrolyzer efficiency fit: a CSV containing parameters of the piecewise linear efficiency curve was included. In these fits were the fixed consumption (kWh) per kg $NH_3$ at zero load and the incremental consumption at full load, based on a given Faradaic efficiency (FE) and maximum current density. In our runs, we assumed a certain FE (66% for the baseline case, corresponding to the off-grid scenario performance) and we could adjust these parameters to analyze sensitivity. However, for comparing grid vs off-grid the electrolyzer performance characteristics were kept the same.

Because the optimization model was solving over a piecewise linear function, rather than the full model, the solver typically found the optimal solution in a matter of seconds. The optimal solution yields the minimum LCOA and provides the values of all decision variables at the optimum.

### Post Processing and Visualization

After solving the optimal values, a few key results were extracted, printed, or displayed for ease of analysis. First was the optimal electrolyzer capacity, which tell us how large the plant needs to be under varying energy supply scenarios. Next is Annual electricity consumption (MWh) and cost (EUR). From this we could derive the average price paid per MWh and the total OPEX. Most importantly is the LCOA (EUR per ton $NH_3$), minimizing the sum of electrolyzed CAPEX and operating costs of energy sourcing. Along similar lines is the LCOE (EUR/kWh or €/MWh). Computed as annual OPEX/total kWh for the grid case, it is effectively the weighted-average electricity price the electrolyzer ended up paying. Finally is the load capacity factor (efficiency), indicating what fraction of time the electrolyzer was running at full. A lower capacity factor means the electrolyzer is idle or at low load much of the time (which might be optimal if electricity is expensive or unavailable); a higher factor means it's used more continuously.

All of these values are rescaled back to their original and accurate form from their current 0-1 state with their respective maximum values.

We then plotted a time-series of operations for sample weeks to illustrate behavior. For one, we plotted the duration curve of electrolyzer power, which shows the distribution of operating levels. This highlighted whether the electrolyzer mostly operated at full power or at intermediate loads. The cost breakdown (CAPEX vs OPEX per ton) was also charted for comparison between scenarios.

## Results

### LCOA Comparison

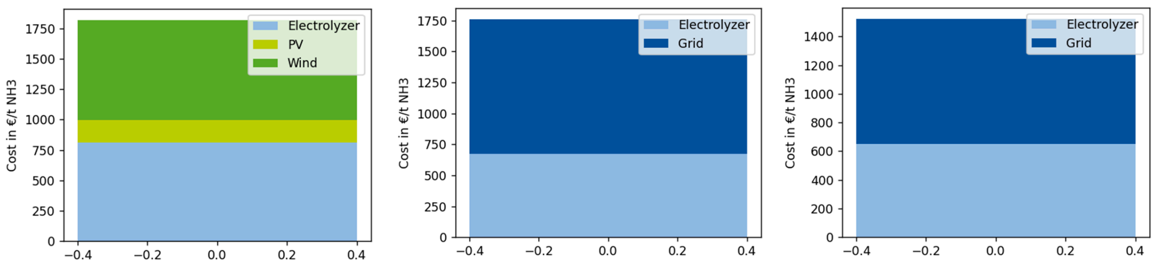


Figure 3: LCOA Cost Comparison and Distribution for curtailment model (Aachen 2023), grid mode 2023, and grid mode 2024

A natural observation of the optimization is that the grid-connected electrolyzer achieves a slightly lower LCOA than the off-grid renewables system under the conditions studied (Aachen 2023 weather data and 2023, 2024 Germany electricity prices). For the case of Aachen in 2023, the model yielded an LCOA of approximately €1,817.5 per ton $NH_3$ for the off-grid (renewable curtailment) scenario, versus €1,759 per ton $NH_3$ for the grid-connected scenario using 2023 electricity prices.

In other words, using wholesale electricity brought the ammonia production cost down by about 3.2% relative to relying solely on local wind and solar in 2023. 2023 saw around 300 hours of negatively priced electricity due to excess renewables, so this result makes intuitive sense: the flexible grid connected plant could exploit these lower-than-0 priced hours. By contrast, an off-grid plant is limited to its own generation; it must be built with excess capacity to meet the $NH_3$ target in sub-optimal weather, leads to higher capital costs and some wasted energy in best-case conditions[2].

To further make sense of this difference in electrolyzer capital costs, it makes sense to break down the cost distribution of the two models.

In the off-grid scenario, the total cost per ton $NH_3$ balances the electrolyzer CAPEX and the renewables CAPEX (wind and PV). The electrolyzer in that scenario was sized larger (compared to grid case) because it needed to absorb a lot of energy during windy/sunny periods to meet the annual production and to counteract the low energy times, in which it would sit idle or at low load (for example, during low-wind nights or cloudy winter days). In essence, the renewable energy is free, but effectively some of it goes unused (curtailment) and wasted when generation exceeds what the electrolyzer can handle. This overcapacity increases the electrolyzer CAPEX allocation per ton of product.

Naturally, in the grid scenario, there is no wind/PV investment at all. The electrolyzer CAPEX is the only capital cost. Thus the electrolyzer doesn't need to accommodate seasonal fluctuations by oversizing, because it can draw from the electric grid sink at any time at a given price. Thus, the CAPEX per ton and capacity is lower. Here, the tradeoff is the electricity OPEX. In the grid case, each MWh of energy used has a market cost. But since the model smartly avoids the most expensive hours, the *average* price paid per MWh is moderate. Conclusively, the total OPEX per ton NH3 came out lower than the implicit cost of energy in the off-grid case. In short, the grid mode optimally saved enough CAPEX and utilized cheap energy such that even after paying for electricity, the net LCOA was slightly less.

These findings are contingent on finding sufficient cheap, low-carbon electricity being available on the grid. It can be visualized from 2023 German price data, the tight correlation between grid prices and carbon emissions. Higher prices correlated to higher carbon emissions and lower prices correlated with lower carbon emissions.

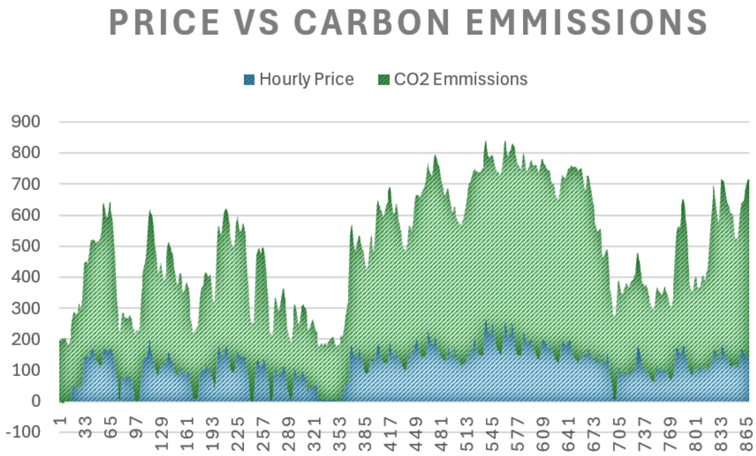


Figure 4: Hourly price in 2023 vs CO2 emissions

If grid prices were uniformly high and intensive in $CO_2$ emissions, this CAPEX advantage could diminish or even reverse. In the 2023 scenario, Germany's grid had a lot of volatility with many low-price periods, enabling the electrolyzer to run during optimal hours at full capacity quite often. However, if the plant had to run during high-price, fossil-fueled peak hours, the economics and emissions would be worse. Hence, one key conclusion is that grid-connected green ammonia can only be cost-competitive when wholesale electricity is sufficiently cheap and preferably renewable. Otherwise, a renewable system might be competitive and preferable despite being a higher upfront investment.

**Electrolyzer Power Comparison**

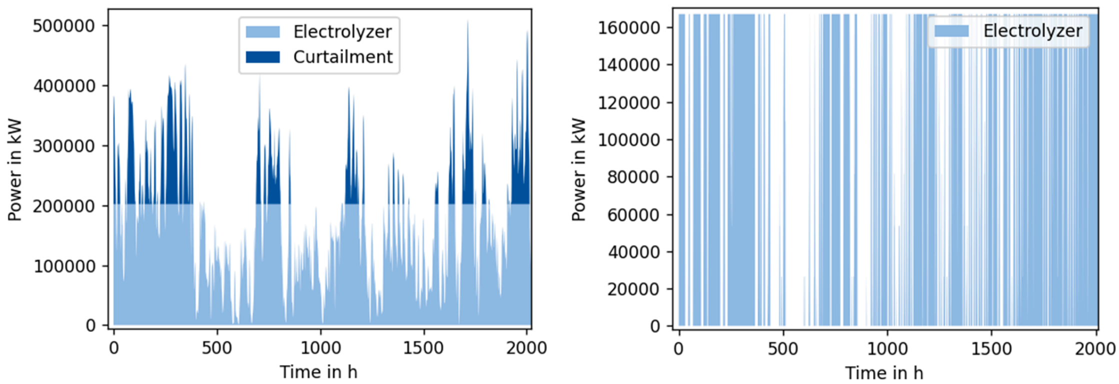


Figure 5: Load factor comparison between 2023 curtailment node and 2023 grid mode

Between the renewable energy powered electrolyzer and grid-connected electrolyzer, two different operational strategies emerge when minimizing the LCOA. Figure 3 illustrates the electrolyzer's distribution of their respective operating levels over the year (zoomed in to the first 2000 hours).

In the off-grid mode, the electrolyzer frequently operated at intermediate loads. It rarely operated at 100% capacity continuously because more often than not, the renewable input did not exactly equal the electrolyzer's max capacity. There were many hours where only a fraction of the capacity could be utilized due to limited wind/PV generation. Additionally, there were some hours of near full power (during simultaneous periods of strong wind and sun) and some hours of zero power (when neither resource was producing energy). This variable profile resulted in an average load capacity factor/efficiency of around 66.3% in the RE model. To interpret 66.3% efficiency, the electrolyzer operated at full capacity for 66.3% of the time or operated at 66.3% capacity for the entire duration of the year. In this case, it is somewhere in between these two extremes.

In contrast, the operation of the grid-connected electrolyzer resembled that of binary behavior and was either operating at 100% or 0% for the majority of the time. The optimization of this model found it ideal to avoid wasting capacity on partial, intermediate loads. If the model found that the electricity was cheap enough to run on, it made sense to utilize the full capacity of the electrolyzer to maximize $NH_3$ output in that hour (since buying a kWh at a cheap price yields more $NH_3$ when the electrolyzer is at its best efficiency, which is usually near full load). If power was expensive or not needed, the electrolyzer would rather sit at a 0% load than run at a low load. The result is that the electrolyzer operates at 100% of its capacity during almost all the hours it is on, which according to the model was 80.8% of the year in 2023 83.3% of the year in 2024.

This means the electrolyzer in grid mode achieved a higher effective 'efficiency' and a more uniform output when it was running. Essentially during favorable price hours, the electrolyzer ramps to full production and completely shuts down in non favorable hours.

Conceptually, this differing operating strategy correlates with other outcomes. While the renewable plant was dictated by the availability of energy (from wind and PV sources), leading to many suboptimal partial-power hours, the grid plant had greater luxury to choose which hours to operate at. As a result, the specific energy consumption (kWh/kg $NH_3$) in the grid scenario was lower than in the off-grid scenario (no energy is curtailed since all of the imported electricity is used).

Another outcome is the correlation with capital usage. Since the grid plant's capacity is active mainly in those high-output hours, it has a higher utilization rate, whereas the off-grid plant's capacity sits underutilized for more time.

By design of the curtailment model and now visualized, the renewable curtailment scenario offers significant periods of excess renewable energy were curtailed because the electrolyzer was already at max and could not use all the available wind/PV. In the 'curt' model, dumping of excess energy occurred whenever generation exceeded electrolyzer capacity. This is an inefficiency forced by the inequality of energy supply and demand. In the grid scenario, theoretically no curtailment occurs in the plant. Theoretically, the plant takes as much as it wants, when it wants, limited exclusively by its own size. Thus, the grid scenario avoids curtailment by optimization of design and power balance. Instead it sometimes curtails demand (by turning off during sub-optimal hours). In a way both strategies curtail in their respective ways: curtailing supply vs curtailing demand. These are different forms of smart operating flexibility that attempt to reduce electrolyzer cost.

In short, the electrolyzer in grid mode behaves in a bi-modal manner: heavily utilized when power is cheap, completely off when power is expensive. Whereas, the off-grid electrolyzer behaves as a taker of its environment: running whenever energy is available, even at low load capacities. The former is a more efficient and cost-effective usage of the electrolyzer capacity. The results confirm this conclusion with the specific CAPEX/t NH3 and electrolyzer capacity being lower in the grid case (as concluded earlier in the LCOA comparison section).

## Energy Price Sensitivity

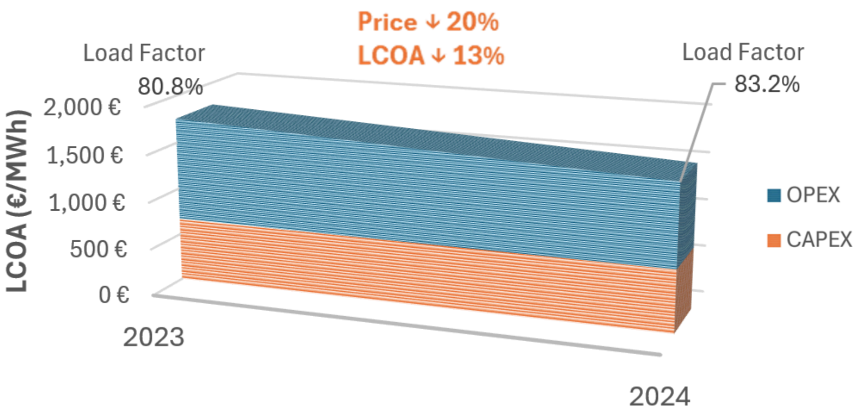


Figure 6: Grid mode– price change between 2023 and 2024 effect on LCOA and load factor

The wholesale price of electricity has a direct impact on the LCOA. Comparing the optimization results for years 2023 and 2024 using German price data, we can see how changed power prices translates to LCOA and operational decisions.

From 2023 to 2024, wholesale electricity prices in Germany fell significantly; the average base-load price dropped from 95.2 €/MWh in 2023 to 78.5 €/MWh in 2024. According to our model, the average price of electricity actually *consumed* by the electrolyzer dropped from 81 €/MWh in the 2023 case to 64 €/MWh in the 2024 case. These values correspond to the model's calculated LCOE of 0.081 €/kWh vs 0.064 €/kWh.

Consequently, the LCOA fell from 1,759.4 €/t to about 1523.6 €/t NH3 (a 13% reduction) from 2023 to 2024. This is an insightful sensitivity: roughly a 20% drop in electricity price yielded a 13% drop in ammonia cost. The relationship is not 1-to-1 because a portion of the LCOA is near-fixed from minimally changed CAPEX, and because the optimizer in 2023 was already avoiding the worst prices. Following intuition, because of 2024's lower-price average, the electrolyzer can run more often at full capacity and doesn't need to be as overbuilt to skip expensive hours. Consistently, the model suggested a slight decrease in optimal electrolyzer capacity for 2024, coupled with a higher annual capacity factor.

With falling prices came the electrolyzer's capacity factor increasing from about 80.8% in 2023 to 83.3% in 2024. This means the electrolyzer was idle only ~16.7% of the time in 2024, versus ~19.2% in 2023. With more hours of affordable power, the plant found it more economical to produce for a greater fraction of the year, thus spreading its fixed cost over more output; that is why the CAPEX contribution/ton came down slightly in 2024. The OPEX per ton, naturally, dropped more substantially because of the lower prices paid for energy. The overall result is that grid-based ammonia production becomes more attractive as electricity prices fall when compared to renewable based ammonia production.

The 2023 vs 2024 comparison underscores a strategic benefit of operational flexibility under volatile market conditions. In both years, the plant adjusted its behavior optimally to respond. In 2023, it was a bit more conservative (limiting run hours to avoid high prices), whereas in 2024 it could be more aggressive (running more since prices were less prohibitive). The optimization automatically finds the new equilibrium. This robust adaptability means a well-designed green ammonia facility could take advantage of trends in the power market. Hypothetically, if we anticipate that as more renewables come online, power will become cheaper and more volatile, a flexible ammonia plant stands to gain from these dips in price. We saw a modest increase in capacity factor from 2023 to 2024, so there will only be further increases if 2025 and beyond brings greener, more low-cost renewable hours.

### Motivations for Operational Flexibility

All of the above results hint at the overarching conclusion that operational flexibility is a key in cost-effective green ammonia production. Seen in both contexts, being able to adjust the production rate in real-time to match resource availability or price signals is crucial.

In the off-grid scenario, the plant must handle the intermittency of wind and solar energy. Flexibility here means that an electrolyzer can operate at variable load capacities, pause when renewable input is low, and ramp up when renewable input is high, without too many reparations. If the system lacked flexibility, required steady operation, and used all of its input energy, it would need massive energy storage or backup generators to maintain continuous running. This would not only be wasteful but would also inflate costs drastically. Therefore, by designing the system to tolerate lower load capacities and variable output, we can size the electrolyzer and renewables more economically and accept that at times no ammonia is produced (and sometimes curtailed).

In the on-grid scenario, flexibility corresponds to its demand response capability. The ability to quickly reduce or fully shut off power draw during high-price periods avoids buying expensive electricity (which would raise OPEX and LCOA). Conversely, the ability to run at full load during cheap periods maximizes output when costs are low. If the plant were inflexible and could not exploit price arbitrage, it would end up paying a higher LCOA. Our model's optimal strategy essentially acts as an ideal responsive load. It emulates the classical supply-demand 'price-taker' in perfect competition: operate at the optimal, market set price, or don't operate at all. It "consumes" electricity when the market price is below the plant's marginal value of ammonia, and stops when above. This flexibility was responsible for the substantial cost savings observed (a 13% LCOA reduction with 20% price drop from 2023 to 2024, as the plant adapted its consumption). In both scenarios, flexibility significantly lowers the cost to produce ammonia by ensuring that neither expensive energy nor idle capital dominate the cost structure. This sets the stage for future green ammonia systems to operate very differently from the rigid fossil-based plants of the 20th century.

## Conclusion

### Key Points

This research examined the techno-economic optimization of green ammonia production through an electrochemical process subject to two different electricity supply modes: renewable (off-grid) setup and a grid-connected setup. Through a Pyomo-based MILP model, the levelized cost of ammonia (LCOA) was minimized and constrained by meeting a fixed annual output.

Despite the differing energy sourcing strategies, the optimal, minimized LCOA reached similar values. With the Germany 2023 energy prices, they were roughly 4–5 times the typical cost of ammonia from natural gas, highlighting the continued necessity for cost reductions in electrolyzer technology and electricity supply. With grid mode supplying a cheaper production of ammonia than from renewables, it marginally improves production cost if grid prices are sufficiently inexpensive and low-carbon.

However, similar LCOAs meant different electrolyzer load factors between renewable and grid-connected systems. The off-grid system required a larger electrolyzer (per unit of output) with a lower utilization due to the variability of on-site renewables. The grid-based system could use a smaller electrolyzer running at a higher utilization (~80-83% capacity factor) by matching its production schedule with the electricity market. Economically beneficial and efficient, utilizing the electrolyzer more fully allowed the grid case to spread its capital cost over more production. Thus, electrolyzer CAPEX/ton ammonia was higher in the off-grid case and slightly lower in the grid case.

More insight into cost distribution is gained by comparing year to year price fluctuations on grid mode. A direct but nonlinear influence of electricity price on ammonia cost occurs: a ~20% drop in average power price from 2023 to 2024 translated to about a 13% reduction in LCOA. With OPEX taking the greater effect of the declining market price (one of two values used to calculate OPEX), CAPEX remains relatively more fixed. However, because electrolyzer utilization increases for a cheaper cost, CAPEX still decreases slightly. As grids incorporate more renewables and prices decline further, we can expect green ammonia costs to come down.

Both scenarios underscore that operational flexibility is crucial. The ability for the grid connected system to ramp production down to zero or up to full capacity on an hourly basis was assumed in the model and was key to cost minimization. In the renewable scenario, flexibility meant the plant could curtail excess imported power and operate at intermediate loads, not bimodally. In short, one curtailed supply (renewables) and one curtailed demand (grid). Rigid operation would lead to much higher costs in either of these contexts, either forcing buying of expensive power or needing massive storage to buffer renewables.

The optimal approach to green ammonia in this analysis is a plant that is right-sized and highly flexible, taking full advantage of either dedicated renewables or grid power when it is cheap and green. This leads to a lower LCOA than a comparably inflexible design. A promising path in which smart scheduling and price-opimization can enable for more cost-competitive renewable ammonia, without full reliance on oversized renewable infrastructure.

**Future Work and Broader Impact**

This work can be extended in several directions. Future research could incorporate a hybrid model or option for energy storage. A hybrid scenario with both on-site renewables and a grid connection, potentially plus storage (battery or hydrogen), takes advantages from both worlds: use cheap grid power when available, fall back on owned renewables when grid is expensive or during grid outages, and store surplus renewable energy for non-peak hours. Optimizing such a hybrid system would add complexity but might further reduce costs and cushion against aggressive price fluctuations and guarantee continuity of $NH_3$ supply.

Zooming out into a broader impact perspective, our findings of necessary operational adaptability reinforce that flexible demand like ammonia synthesis can play a pivotal role in a renewable-based energy economy. Ammonia plants act as a controllable load that can soak up excess power and completely shut off during spikes of generated electricity, thus facilitating more efficient use of already-generated: functioning as part-time energy storage or demand response. This mutually beneficial tradeoff of grid–electrolyer demand and pull of cheap electricity makes the whole system cheaper and greener and thus, policy makers and energy planners should consider incentivizing industrial users to be flexible. For example, special electricity tariffs could be offered to green ammonia producers to encourage them to consume electricity when it’s abundant and back off when it’s scarce, aligning economic signals with grid needs.

In conclusion, achieving economically viable green ammonia will require system-level optimization, not letting one component, in isolation, improve alone. This study contributes to that understanding by quantifying how design and scheduling choices impact cost. The path forward for a greener future lies in integrating engineering, economics, and policy to steadily close the gap between green and traditional ammonia.

## References

[1] M. J. Rix and A. Mitsos, “Early-stage techno-economic evaluation of electrochemical nitrogen reduction to ammonia based on catalyst performance,” Electrochimica Acta, vol. 523, p. 145893, 2025, doi: 10.1016/j.electacta.2025.145893.

[2] K. Bidaoui, Optimales Design und Betrieb der elektrochemischen Stickstoffreduktion zur Ammoniakproduktion mit erneuerbarer Energiezufuhr, M.S. thesis, Lehrstuhl für Systemverfahrenstechnik, RWTH Aachen Univ., Aachen, Germany, 2024.

[3] Y. He et al., “Ammonia electrolyzer design for high efficiency and product concentration,” Nature Catalysis, Jul. 2025. [Online]. Available: https://www.nature.com/articles/s44286-025-00207-9