

Solar Power as an Optimal Portfolio

David Kraljic
COMCOM CAPITAL d.o.o.,
Faculty of Mathematics and Physics,
University of Ljubljana
Ljubljana, Slovenia

Jernej Katanec, Matej Logar,
Blaz Sobocan, Miha Troha
COMCOM CAPITAL d.o.o.
Ljubljana, Slovenia

Abstract—Conventional solar power plants are typically oriented southward to maximize energy production. However, maximizing energy generation does not necessarily translate to maximizing revenue, as power plants or offtakers are subject to fluctuating market prices with varying volatility. In this study, we investigate optimal orientations of power plants to maximize revenue and how these have evolved over time. Additionally, we incorporate the impact of solar production on market prices by analyzing Day-Ahead supply and demand curves. Our findings illustrate how the optimal configuration adapts when accounting for the “cannibalization” effect, where increased solar capacity drives down market prices. Furthermore, we consider each panel of a solar power station as an asset that can be “invested” in different sky directions. By aligning panels in diverse orientations, they are exposed to varying price and volatility profiles. Using a mean-variance portfolio framework, we construct an optimal portfolio of panel orientations that balances profit and risk, demonstrating how the revenue volatility of a solar power plant can be reduced.

Index Terms—Solar power, Optimal portfolio theory, Solar cannibalization

I. INTRODUCTION

The rapid expansion of solar photovoltaic (PV) capacity has significantly influenced electricity markets. However, financial returns for solar generators do not always scale proportionally with energy production. Traditional PV power stations are typically designed to maximize total energy generation by orienting panels southward in the northern hemisphere (see, e.g., [1], [2]).

In liberalized electricity markets, however, maximizing generation does not necessarily equate to maximizing revenue. Due to the merit-order effect, increased solar output during peak production hours can suppress electricity prices, sometimes leading to negative pricing events, as observed in various markets across the USA and Europe [4]. This price cannibalization effect reduces the economic value of additional solar capacity and highlights the importance of designing PV systems that optimize financial performance rather than energy output alone [3].

The impact of price cannibalization grows as solar penetration increases. Rising PV capacity drives midday electricity prices downward relative to the evening peak, creating an incentive to shift solar production later in the day. Recent

research has shown that optimizing panel orientation can improve financial outcomes by aligning production with higher electricity prices [3], [4]. For example, west-facing panels can generate more revenue by capturing higher-priced afternoon and evening hours.

Beyond individual system optimization, the deployment strategy of solar capacity itself plays a crucial role in determining revenue outcomes. If new capacity is installed sequentially and optimized independently by different investors, market dynamics change over time, often leading to suboptimal orientations and excessive price cannibalization. A more coordinated optimization approach, accounting for future installations and their impact on electricity prices, could mitigate these inefficiencies and improve long-term financial returns.

This study builds on previous work by examining how optimal PV orientations evolve in response to market conditions and increasing solar capacity. Using historical Day-Ahead market data from Germany, we analyze how additional solar deployment shifts supply-demand equilibrium and assess the revenue-maximizing panel configurations under different market scenarios. We explicitly model price feedback effects, where additional solar capacity influences market-clearing prices.

Additionally, we explore strategies for managing revenue volatility. A solar power plant’s revenue depends not only on average electricity prices but also on exposure to price fluctuations. Using a mean-variance portfolio framework, we formulate an optimal portfolio of panel orientations that balances expected revenue with financial risk.

Instead of treating a single solar farm as a monolithic entity, we consider each panel as an individual asset that can be “invested” in different orientations, thereby diversifying exposure to price fluctuations. Research using Modern Portfolio Theory (MPT) has so far focussed on geographical diversification and demonstrated that treating renewable energy within the MPT framework can mitigate revenue volatility while improving overall financial returns [5], [6].

Our findings reveal that a diversified panel orientation strategy, rather than a uniform south-facing configuration, can reduce revenue volatility. We also demonstrate that price cannibalization effects are exacerbated when solar capacity is deployed sequentially without market-wide coordination, leading to inefficient system orientations and lower long-term

revenues.

This research contributes to the ongoing discourse on solar investment strategies, providing actionable insights for developers, investors, and policymakers seeking to maximize the economic value of solar assets in competitive electricity markets.

II. DATA AND TOOLS

A. Weather Data

For historical solar data—including shortwave, diffuse, and direct normal solar radiation—we use outputs from the ECMWF IFS¹ model, obtained from [7]. The data is available at an hourly resolution. To illustrate our analysis, we selected a test location in southern Germany (lat: 48.65, lon: 11.28).

B. Market Data

We obtained historical Day-Ahead prices, along with supply and demand curves for each auction hour, from EPEX [8]. Our analysis focuses on hourly Day-Ahead prices for the German market.

While solar power plant output is also exposed to intra-day and imbalance prices due to the inherent uncertainty of solar generation, Day-Ahead forecasts are generally accurate enough that the bulk of solar production can be reasonably priced using Day-Ahead prices as a proxy [9]. We use data from 2019 onwards but exclude 2022, as extreme price fluctuations due to natural gas shortage during that year are unrepresentative of the typical market.

C. Tools

We use *pvl* to convert solar irradiance data to power [10]. Most of the optimization problems are solved using the DIRECT algorithm [11], except for the optimal portfolio problem, which is solved using GUROBI [12].

III. OPTIMIZATION OBJECTIVES

A. Maximizing Energy Production

The simplest objective for optimizing a solar power plant at a given location is maximizing total energy output. If we denote the power generated at time t , with surface tilt θ and azimuth ϕ , as $P(t, \theta, \phi)$, the optimization problem is formulated as:

$$\max_{\theta, \phi} \sum_{t=t_0}^{t_f} P(t, \theta, \phi) \quad (1)$$

where t_0 and t_f denote the start and end times of the optimization period. In general, ignoring specific weather patterns and terrain constraints, the optimal orientation is south-facing. Using realistic solar irradiance data for the period 2019–2024 at our test location, we solve the optimization problem defined in Eq. 1. Results are shown in Fig. 1.

¹European Centre for Medium-Range Weather Forecasts Integrated Forecasting System

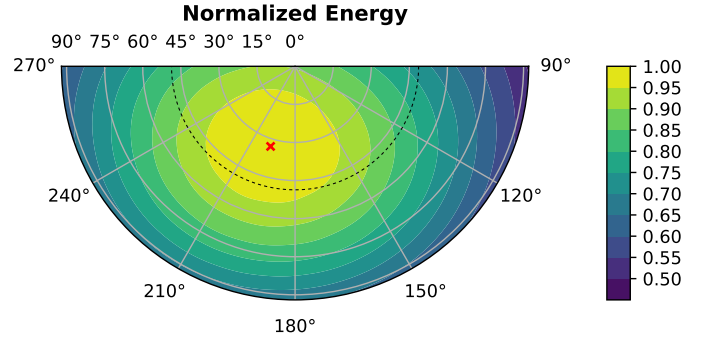


Fig. 1: Energy production represented on a polar plot. The values are normalized by the maximum energy output, which occurs at the orientation marked by the red cross. The radial axis represents panel tilt, while the angular axis represents azimuth (90° is East, 180° is South, etc.). The dotted line indicates the latitude line.

B. Maximizing Revenue

If the goal is to maximize revenue rather than just energy production, the optimization problem is given by:

$$\max_{\theta, \phi} \sum_{t=t_0}^{t_f} P(t, \theta, \phi) \cdot \pi(t) \quad (2)$$

where $\pi(t)$ represents the electricity price at time t . This formulation ensures that the optimization accounts for both energy output and its economic value at different times of the day. We solve the optimization problem in Eq. 2 for the period 2019–2025 and illustrate the result in Fig. 2.




Fig. 2: The optimal panel orientation maximizing revenue (red cross) is compared with the orientation maximizing energy production (black cross). The optimal revenue orientation shifts westward to capture evening peak prices.

The precise orientation that maximizes annual revenue depends on the specific solar irradiance and market price conditions in a given year. However, there is a general trend primarily driven by the increasing installation of solar capacity (see Fig. 3) that pushes midday prices lower relative to the




Fig. 3: Growth of installed solar capacity in Germany [13].

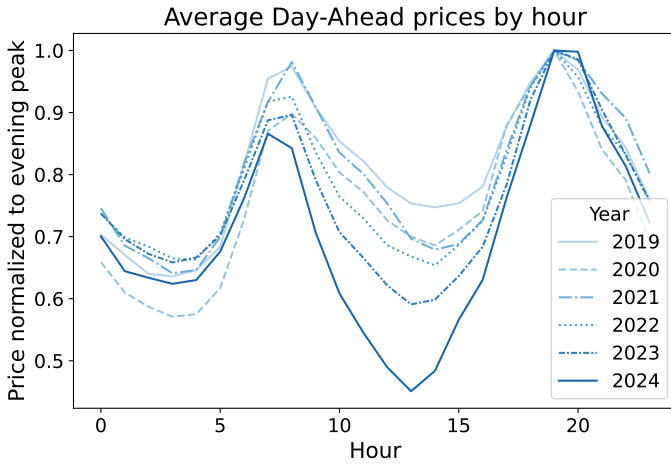


Fig. 4: Daily profile of Day-Ahead prices over the years, normalized to the evening peak.

evening peak (see Fig. 4). Consequently, when optimizing panel orientation to maximize revenue (Eq. 2) each year, we observe a shift toward westward-facing configurations (to capture higher evening prices) and an increase in tilt angle (to capture more revenue during winter months). This evolution is illustrated in Fig. 5.

C. Maximizing Revenue Through Scenario Analysis of Increased Solar Capacity

The above formulation assumes historical market prices. However, when additional solar power plants are built, the supply-demand equilibrium changes. To account for this feedback effect, the optimization problem is reformulated as:

$$\max_{\theta, \phi} \sum_{t=t_0}^{t_f} P(t, \theta, \phi) \cdot \pi(\hat{P}(t, \theta, \phi), t) \quad (3)$$

where $\pi(\hat{P}(t, \theta, \phi), t)$ represents the electricity price as a function of both time and the output $\hat{P}(t, \theta, \phi)$ of all new solar installations. This formulation captures the impact of increased solar production on market-clearing prices.




Fig. 5: Orientation of panels to maximize revenue for each year. The effect of solar cannibalization drives an increasingly westward-facing configurations over time.

To quantify the effect of additional solar capacity on prices, we calculate the expected output of a new solar installation and insert it at the beginning of the supply curve. We then recalculate the supply-demand equilibrium to determine the resulting market price shift. This method implicitly assumes that solar power is always "in-the-money", meaning that the cleared price remains higher than the bid price of a solar power plant.

For the period 2019–2025, we assume the addition of 4 GW of solar capacity to the German Day-ahead market while keeping other market conditions unchanged. The optimal panel orientation shifts significantly, and the captured revenue is reduced by approximately 50% (see Fig. 6).

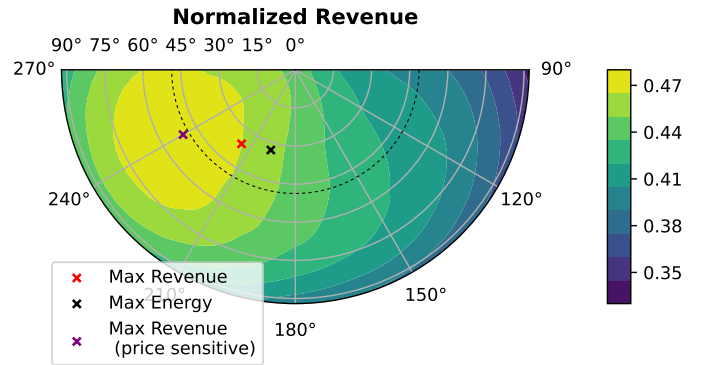


Fig. 6: Optimal panel orientation accounting for price cannibalization of additional solar installations. The optimal tilt and azimuth differ significantly from the previous maximum revenue assumption, reflecting the impact of additional 4 GW solar capacity on market prices. The revenues are normalized to the case where prices do not respond to additional solar capacity (red cross).

New solar capacity introduces additional complexity when determining optimal orientations. For instance, if the 4 GW of new solar capacity is deployed in sequential 1 GW increments,

each stage being optimized based only on prior installations, the final system configuration is suboptimal compared to the case where all 4 GW are optimized collectively. This fragmented decision-making process results in 4.6% lower hypothetical revenues over the 2019–2025 period.

This inefficiency arises due to a lack of full market transparency regarding the amount, location, and timeline of new solar installations. As a result, early-stage power plants are oriented based on outdated price expectations, leading to increasingly suboptimal configurations as market prices adjust to rising solar penetration. Figures 7 and 8 illustrate these dynamics.




Fig. 7: Impact of additional solar capacity on Day-Ahead prices. The plot shows average hourly Day-Ahead prices for the period 2019–2025 (excluding 2022). Two scenarios are compared: (1) optimizing all 4 GW of additional solar capacity collectively and (2) sequentially optimizing 1 GW at a time, based only on prior installations. In the second case, we observe greater midday price suppression due to suboptimal orientation choices, whereas in the first case, more capacity is optimized to capture the evening peak.

D. Reducing Volatility

Panels oriented in different directions are exposed to distinct price profiles, leading to variations in both revenue levels and revenue volatility. To model this, we denote the fraction of a solar power plant oriented in a given direction as $f(\theta, \phi)$. The total revenue of the plant is then given by:

$$R = \sum_{t=t_0}^{t_f} \sum_{\theta} \sum_{\phi} f(\theta, \phi) P(t, \theta, \phi) \pi(t) \quad (4)$$

The corresponding revenue volatility is:

$$\sigma_R^2 = \sum_{\theta, \phi, \theta', \phi'} f(\theta, \phi) \Sigma_R(\theta, \phi, \theta', \phi') f(\theta', \phi') \quad (5)$$

where $\Sigma_R(\theta, \phi, \theta', \phi')$ represents the covariance matrix of revenues between panel orientations (θ, ϕ) and (θ', ϕ') .




Fig. 8: Comparison of solar generation profiles under different optimization strategies. The collective optimization of all 4 GW of solar capacity results in a greater shift toward evening production, reducing midday cannibalization. Conversely, sequential investments without market-wide coordination lead to excessive midday generation, reinforcing price suppression (see Fig. 7).

Instead of orienting all panels in the same direction to maximize expected revenue, some revenue can be sacrificed to achieve lower volatility. To balance these competing objectives, we introduce $\lambda \geq 0$ as a parameter controlling the trade-off between revenue and volatility. The optimization problem is then formulated as:

$$\max_{f(\theta, \phi)} R - \lambda \sigma_R^2 \quad (6)$$

This is a classic mean-variance portfolio optimization framework [14], where each panel orientation is treated as an "asset" with associated expected revenue and revenue volatility. Given a specific risk aversion parameter λ , we solve for $f(\theta, \phi)$ in Eq. 6 to determine the optimal distribution of panel orientations.

For example, setting λ in Eq. 6 such that the optimized revenue volatility is at most 85% of the volatility of a revenue-maximizing plant results in a 4.5% reduction in expected revenue. The corresponding power plant consists of panels distributed across multiple orientations, as illustrated in Fig. 9

IV. FUTURE WORK

In this work we analyzed the impact of price cannibalization resulting from additional solar capacity using historical supply and demand curves. As this approach is retrospective, our analysis primarily serves as a what-if scenario assessment. A more comprehensive approach would involve forward-looking price modelling that accounts for projected market dynamics, including changes in other generation sources within the supply stack, not just solar.

Specifically, we assumed that the additional solar capacity is introduced solely into the Day-Ahead market, added at

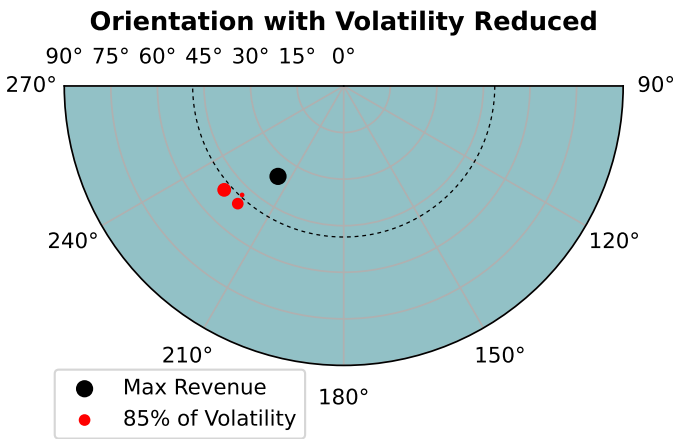


Fig. 9: Reducing revenue volatility by diversifying panel orientations. The black dot represents the revenue-maximizing orientation, while the red dots indicate a diversified set of orientations that reduce volatility to 85% of the original value. The size of each red dot is proportional to the fraction of panels oriented in that direction, ensuring that the total area of red dots equals the area of the black dot.

the beginning of the supply stack, and does not respond dynamically to market conditions. We also assumed that all the additional capacity is added into the German supply stack. In reality, due to market coupling, some of this capacity would be distributed across other countries. As a result, our modeled addition of 4 GW to the German Day-Ahead stack corresponds to a higher level of physically installed capacity in practice.

Additionally, our analysis focused on a single location for illustrative purposes. A more complete study would extend this framework to multiple locations, considering the spatial distribution of new solar installations and their aggregate effect on market prices. Future work should also explore how geographic diversification influences price cannibalization and revenue optimization, providing a more realistic perspective on the deployment of solar power in competitive electricity markets. A separate study, beyond the scope of this work, could focus on identifying the optimal portfolio of solar installations, potentially including axis-tracking systems, combined with energy storage solutions.

V. CONCLUSIONS

This study analyzes the trade-off between optimizing a solar power plant for maximum energy production versus maximum revenue. We demonstrated that the increasing deployment of solar capacity has been driving midday electricity prices lower relative to the evening peak. As a result, the optimal panel orientation has shifted westward to capture higher prices in the evening while avoiding excessive cannibalization during midday hours.

Using supply and demand curves from the Day-Ahead market, we examined the impact of additional solar capacity on market prices. Our findings indicate that if new solar instal-

lations are optimized independently, such as through sequential investments by different investors, resulting orientations may be suboptimal, leading to increased price cannibalization and reduced overall profitability. In contrast, a coordinated optimization approach could mitigate these inefficiencies and better align generation with market dynamics.

Finally, we showed that a single solar power plant can reduce its revenue volatility by diversifying panel orientations instead of adopting a uniform direction. The ability to control volatility can be an essential factor in structuring Power Purchase Agreements, particularly as solar cannibalization continues to depress revenues and lower the revenue-to-volatility ratio. As market conditions evolve, incorporating volatility management strategies will become increasingly important for ensuring the financial viability of solar investments.

REFERENCES

- [1] Mehleri, E. D., Zervas, P. L., Sarimveis, H., Palyvos, J. A., & Markatos, N. C. (2010). Determination of the optimal tilt angle and orientation for solar photovoltaic arrays. *Renewable Energy*, 35(11), 2468–2475.
- [2] Naraghi, M. H., & Atefi, E. (2022). Optimum solar panel orientation and performance: A climatic data-driven metaheuristic approach. *Energies*, 15(624).
- [3] Durán-Castillo, G. E., Weis, T., Leach, A., & Fleck, B. A. (2023). Revenue-optimized photovoltaic orientation in a northern competitive electricity market with carbon offsets. *Energy Reports*, 10, 3133–3145.
- [4] Brown, P. R., & O’Sullivan, F. M. (2019). Shaping photovoltaic array output to align with changing wholesale electricity price profiles. *Applied Energy*, 256, 113734.
- [5] Hu, J., Harmsen, R., Crijns-Graus, W., & Worrell, E. (2019). Geographical optimization of variable renewable energy capacity using modern portfolio theory. *Applied Energy*, 253, 113614.
- [6] Schmalensee, R. (2017). The performance of U.S. wind and solar generators. *Energy Policy*, 108, 742-750.
- [7] Zippenfenig, P. Open-Meteo.com Weather API. (2023), <https://open-meteo.com/>
- [8] EPEX SPOT SE Day-Ahead market data. (2025), <ftp.epexspot.com>, Accessed: 2025-02-06
- [9] Kulakova, S. & Zielb, F. The impact of renewable energy forecasts on intraday electricity prices. *Economics Of Energy & Environmental Policy*. 10, pp. 79-104 (2021), <https://www.jstor.org/stable/27030680>
- [10] Anderson, K., Hansen, C., Holmgren, W., Jensen, A., Mikofski, M. & Driesse, A. pvlb python: 2023 project update. *Journal Of Open Source Software*. 8, 5994 (2023,12), <http://dx.doi.org/10.21105/joss.05994>
- [11] Jones, D., Perntunen, C. & Stuckman, B. Lipschitzian optimization without the Lipschitz constant. *Journal Of Optimization Theory And Applications*. 79, 157-181 (1993,10), <http://dx.doi.org/10.1007/BF00941892>
- [12] Gurobi Optimization, Version 9.0.2, LLC Gurobi Optimizer Reference Manual. (2023), <https://www.gurobi.com>
- [13] Bundesnetzagentur für Elektrizität, Gas, Telekommunikation, Post und Eisenbahnen SMARD: Electricity market data platform. (2025), <https://www.smard.de/>, Accessed: 2025-02-07
- [14] Markowitz, H. Portfolio selection. *The Journal Of Finance*. 7, 77 (1952,3), <http://dx.doi.org/10.2307/2975974>