

Article

Polynomial Chaos-Based Stochastic Dispatch with Adaptive Setpoint Control for Renewable-Integrated Electric Arc Furnace Steelmaking

Cong Xu ^{1,*}, Yuanqi Kong ¹  and Yafei Zhao ² ¹ International Business School, Qingdao Huanghai University, No. 1145 Linghai Road, Huangdao District, Qingdao 266427, China; kongyq@qdhhc.edu.cn² School of Smart City Engineering, Qingdao Huanghai University, No. 1145 Linghai Road, Huangdao District, Qingdao 266427, China; zhaoyf@qdhhc.edu.cn

* Correspondence: xuc14@qdhhc.edu.cn

Abstract

Scrap-based electric arc furnace (EAF) steelmaking powered by on-site variable renewable energy is a key decarbonisation route, but the heteroscedastic, non-Gaussian nature of joint wind–photovoltaic forecast errors makes the EAF—a large, metallurgically constrained load—hard to coordinate with on-site generation under feeder limits. We develop a unified stochastic dispatch and adaptive setpoint-control framework. A chance-constrained dispatch over a zone-wise Beta uncertainty model is propagated through a degree-two polynomial chaos expansion (PCE) and reformulated as a second-order cone programme via the Cantelli inequality, with EAF-specific metallurgical constraints (electrode slew, short-circuit-ratio-tied flicker, stage-dependent melt-power floor, multi-stage tap-to-tap profile) embedded by the same procedure. The EAF setpoint gain is then extracted in closed form—without Jacobian inversion—as a ratio of first-order PCE coefficients, so it inherits the dispatch’s 95% feeder-security guarantee. Calibrated on 24 months of real wind/PV data for a Qingdao site (ERA5 reanalysis vs. archived ECMWF-IFS forecast), which confirms the heteroscedastic premise and a measured wind–PV error correlation of 0.015, the extracted gain scales across the Low–Mid–High zones (medians 6.07, 11.91, 17.22 p.u.) following the operating regime rather than the disturbance magnitude. The scheme bounds worst-case tracking below 1.18 MW per zone (vs. up to 3.34 MW for no droop), satisfies the feeder limit in 100% of realisations, matches model-predictive control without online optimisation, and lowers within-EAF specific CO₂ emissions by 4.4% versus no droop. An out-of-sample test on real records confirms a decisive advantage in the data-rich zones and, candidly, a shortfall in the data-limited high-wind zone.



Academic Editor: Marco Ramírez-Argáez

Received: 20 May 2026

Revised: 25 June 2026

Accepted: 6 July 2026

Published: 13 July 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: stochastic optimal dispatch; polynomial chaos expansion; chance constraints; adaptive setpoint control; renewable integration; electric arc furnace steelmaking; CO₂ reduction

1. Introduction

A single 80 MW electric arc furnace (EAF) drawing 350–500 kWh per tonne of crude steel is one of the largest controllable loads on its local feeder, capable of consuming 50–250 GWh per year. When such a furnace is co-located with on-site wind and photovoltaic (PV) generation, the operational problem is to coordinate this large, fast, but metallurgically constrained load with intermittent generation under a contracted feeder limit, so

that on-site renewable energy is used to the greatest extent without violating grid-security or process-integrity limits [1–5]. This electric-demand coordination problem—rather than emission accounting per se—is the central subject of this paper. The difficulty is that wind and PV outputs depart from their forecasts in ways that depend strongly on the prevailing operating zone, are bounded by rated capacity, and propagate through the plant energy hub nonlinearly once feeder, battery-storage and EAF process constraints are imposed.

The wider motivation is decarbonisation. The iron and steel industry accounts for roughly seven per cent of global anthropogenic CO₂ emissions, and incremental efficiency gains within the long blast-furnace–basic-oxygen-furnace (BF–BOF) route are no longer sufficient for mid-century climate targets [6–10]. Two structural pathways dominate: substitution of fossil reductants by hydrogen and biomass, and migration of melting capacity to the scrap-based EAF route powered by low-carbon electricity [11–16]. The carbon advantage of the EAF route, however, is realised only if the electricity is itself low-carbon and if the furnace can actually absorb on-site renewable generation in real time; this returns the problem to the electric-demand coordination question above and motivates a control framework rather than an accounting exercise.

On-site battery energy storage is an increasingly standard component of such renewable-integrated industrial sites, providing the intra-day flexibility that buffers the fastest VRE fluctuations and relieves the controllable load of the highest-frequency regulation [17]. A growing body of work addresses adjacent coordination problems through stochastic and robust optimisation, including data-driven robust microgrid scheduling, information-gap decision theory for renewable-coupled hubs, and two-stage robust unit commitment [18–24]. A closely related dispatch literature on flexible carbon-capture plants and joint carbon–green-certificate trading has shown that probabilistic constraints can be combined with explicit environmental-compliance cost terms while remaining tractable [25–41].

What is missing is a single end-to-end framework that propagates a heteroscedastic, distribution-aware representation of joint wind–PV uncertainty through the dispatch of a renewable-integrated EAF plant and simultaneously extracts an adaptive equipment-level setpoint coefficient whose statistical security is inherited from the dispatch. The chance-constrained optimal-power-flow literature has matured along this axis for power systems, with polynomial chaos expansion (PCE) emerging as a non-intrusive surrogate that propagates continuous random variables through nonlinear power-flow equations and admits a second-order cone (SOC) reformulation. The EAF analogue has not been built, nor has the heteroscedastic structure of joint wind–PV forecast errors at the steel plant been integrated with adaptive setpoint extraction in a way that preserves the probabilistic feeder-security guarantee and respects EAF-specific metallurgical limits.

The present work closes that gap by transposing and substantially extending the SOPF-based adaptive droop control recently proposed for hybrid AC–HVDC offshore wind grids [42] to the green metallurgy setting. The contribution is threefold. At the uncertainty-modelling layer, a zone-wise Beta parameterisation of joint wind–PV forecast errors is integrated into the dispatch as a heteroscedastic, data-driven input. At the dispatch layer, a chance-constrained problem is formulated in the energy-flow representation of the plant energy hub—with feeder, carbon-quota and, crucially, EAF-specific metallurgical constraints (electrode/transformer slew, short-circuit-ratio-tied flicker, stage-dependent melt-power floor, multi-stage tap-to-tap profile)—and propagated through a degree-two PCE on a Jacobi basis with explicit Galerkin projection. At the controller layer, the EAF setpoint gain is extracted in closed form from the first-order PCE coefficients via a Jacobian-free sensitivity identity. The primary methodological innovation is that this gain is the exact algebraic image of the chance-constrained dispatch solution and therefore

provably inherits the dispatch's 95% feeder-security guarantee, under EAF-specific metallurgical constraints that have no analogue in generic flexible-load control—in contrast to heuristic deviation-dependent droops, which carry no such guarantee.

The remainder of the paper is organised as follows. Section 2 formulates the zone-wise Beta model, the PCE-based chance-constrained dispatch with EAF metallurgical constraints, and the sensitivity-based gain extraction. Section 3 reports the simulation study, including zone-resolved coefficient statistics, mean and tail tracking, statistical significance and confidence interval analysis, a model-predictive-control benchmark, a wind–PV correlation sensitivity study, goodness-of-fit of the Beta model, rigorous CO₂ accounting, computational performance, and a 24 h operating case. Section 4 discusses implications, validation scope and limitations. Section 5 concludes.

2. Materials and Methods

2.1. Plant Architecture and Decision Variables

We consider a representative scrap-based EAF plant (Figure 1) with a single 80 MW EAF as the primary controllable load, supplied by on-site wind ($P_w^{\max}$), on-site PV ($P_s^{\max}$), and grid imports through a feeder of capacity $\bar{P}_g$ at the point of common coupling (PCC). A grid-tie battery energy storage system (BESS) of capacity $E_b^{\max} / P_b^{\max}$ provides intra-day flexibility. The scheduling step is τ , and the horizon is T .




Figure 1. Energy-hub architecture of the renewable-integrated EAF plant. Solid arrows: unidirectional flows; dashed: bidirectional.

The decision variables at each step t are the grid import $P_{g,t}$, the BESS power $P_{b,t}$, the wind and PV curtailment $\Delta P_{w,t}, \Delta P_{s,t}$, and the EAF setpoint $P_{f,t}^{\text{ref}}$. Realised injections $P_{w,t}, P_{s,t}$ depart from forecasts $\hat{P}_{w,t}, \hat{P}_{s,t}$ per the zone-wise model below; the state-dependent representation is adopted because pooled single-distribution surrogates are systematically miscalibrated outside their fitting regime [19–23].

2.2. Zone-Wise Beta Parameterisation of Renewable Uncertainty

Wind errors are small near cut-in and rated power and large in ramp regions; PV errors are small at dawn/dusk, large under transient cloud cover, and intermediate near clear-sky peak. To capture this heteroscedasticity while remaining compatible with the

Jacobi basis of the PCE (Section 2.3), we adopt a zone-wise Beta parameterisation [42] per VRE source. With $\hat{p} = \hat{P}/P^{\max} \in [0, 1]$,

$$z(\hat{p}) = \begin{cases} \text{Low,} & \hat{p} < \tau_1, \\ \text{Mid,} & \tau_1 \leq \hat{p} < \tau_2, \\ \text{High,} & \hat{p} \geq \tau_2, \end{cases} \quad (1)$$

with $\tau_1 = 0.3, \tau_2 = 0.7$. Per zone z and source $\xi \in \{w, s\}$, empirical mean μ_z^ξ and standard deviation σ_z^ξ of the normalised error are estimated from forecast–realisation pairs. Given $\hat{p}$ in zone z , the effective mean is $\mu_{\text{eff}} = \hat{p} + \mu_z^\xi$, and the Beta parameters follow by moment matching:

$$\alpha = \mu_{\text{eff}}K, \quad \beta = (1 - \mu_{\text{eff}})K, \quad K = \frac{\mu_{\text{eff}}(1 - \mu_{\text{eff}})}{(\sigma_z^\xi)^2} - 1. \quad (2)$$

The resulting $p_\xi \sim \text{Beta}(\alpha, \beta)$ has support $[0, 1]$, unlike Gaussian surrogates that admit negative or super-rated draws. Wind and PV are treated as conditionally independent given their forecasts—an assumption whose impact is tested directly by the copula-correlation study of Section 3.7—which is consistent with day-ahead wind/PV error correlations typically below 0.15 at temperate latitudes [21,24]. The statistics used here are in Table 1; their empirical calibration against measured records is discussed in Sections 3.8 and 4.

Table 1. Zone-wise statistics of normalised forecast errors, calibrated on 24 months (2023–2024) of hourly data for a Qingdao site (35.96° N, 120.20° E): realisation from ERA5 reanalysis, day-ahead forecast from the archived ECMWF-IFS operational model (Open-Meteo). n is the number of hourly samples per zone (PV restricted to daytime).

Source	Zone	Forecast Range	n	μ_z^ξ	σ_z^ξ
Wind	Low	[0, 0.3)	6491	0.035	0.089
Wind	Mid	[0.3, 0.7)	609	0.122	0.183
Wind	High	[0.7, 1.0]	122	0.069	0.126
PV	Low	[0, 0.3)	1476	−0.002	0.048
PV	Mid	[0.3, 0.7)	1461	0.009	0.057
PV	High	[0.7, 1.0]	729	0.007	0.038

The measured statistics confirm the heteroscedastic premise: for both sources the Mid-zone standard deviation is the largest (wind 0.183 vs. 0.089/0.126; PV 0.057 vs. 0.048/0.038), so forecast uncertainty genuinely peaks in the ramp/transient regime rather than being uniform. The High-zone wind statistic rests on only 122 h—sustained high wind is rare at this coastal site—and is the least certain, a point we return to in the out-of-sample analysis (Section 3.13). Critically, the measured wind–PV forecast-error correlation over the daytime record is $r = +0.015$, which directly supports the conditional-independence assumption (its effect is stress-tested up to $\rho = 0.5$ in Section 3.7).

Figure 2 shows the density family at three forecast levels. The bounded support guarantees that the chance constraints below never receive a physically impossible renewable draw; the Mid-zone density is the widest for both sources; and as $\hat{p} \rightarrow 0$ or 1, the density develops a boundary skew that the Galerkin projection (Section 2.3) recovers exactly through the Jacobi basis.




Figure 2. Zone-wise Beta densities of normalised renewable output at $\hat{p} = 0.15$ (Low, dashed), 0.50 (Mid, solid), 0.85 (High, dotted). (a) Wind; (b) PV.

2.3. PCE-Based Chance-Constrained Dispatch

Let $\xi = (\xi_1, \xi_2)$ be the independent germs of the wind and PV sources, each Beta-distributed per (2), with joint density ρ . Any $x \in L^2(\rho)$ admits the Jacobi-orthogonal expansion

$$x(\xi) = \sum_{k=0}^{\infty} \hat{x}_k \Phi_k(\xi), \quad \hat{x}_k = \frac{\langle x, \Phi_k \rangle_{\rho}}{\gamma_k}, \quad \gamma_k = \mathbb{E}_{\rho}[\Phi_k^2], \quad (3)$$

truncated at total degree d to give $K + 1 = \binom{N_{\xi} + d}{d}$ terms. With $N_{\xi} = 2, d = 2$, this is six coefficients per variable [42]. The mean and variance follow as $\mathbb{E}_{\rho}[x] = \hat{x}_0$ and $\text{Var}_{\rho}[x] = \sum_{k \geq 1} \gamma_k \hat{x}_k^2$. Equality constraints $r(x, \xi) = 0$ (energy balance, BESS dynamics, curtailment) are enforced by Galerkin projection,

$$\int_{[0,1]^2} r\left(\sum_k \hat{x}_k \Phi_k, \xi\right) \Phi_j \rho \, d\xi = 0, \quad j = 0, \dots, K, \quad (4)$$

collapsing the stochastic equality into $K + 1$ deterministic equations in coefficient space; for the affine energy balance used here, the system is affine in the decision variables, preserving SOC structure. We solve (4) non-intrusively by least-squares regression on N_{coll} collocation samples from ρ .

The chance-constrained dispatch minimises expected cost subject to energy balance and feeder, EAF and carbon limits:

$$\min_x \mathbb{E}\left[\sum_{t=1}^T (c_{g,t} P_{g,t}(\xi) + c_c E_t(\xi) + c_p |P_{f,t}(\xi) - P_{f,t}^{\text{pref}}|)\right] \quad (5)$$

$$\text{s.t. } g(x, \xi) = 0, \quad (6)$$

$$\mathbb{P}(\underline{P}_g \leq P_{g,t}(\xi) \leq \overline{P}_g) \geq 1 - \epsilon, \quad (7)$$

$$\mathbb{P}(\underline{P}_f \leq P_{f,t}(\xi) \leq \overline{P}_f) \geq 1 - \epsilon, \quad (8)$$

$$\mathbb{P}(\sum_t E_t(\xi) \leq \overline{E}) \geq 1 - \epsilon, \quad (9)$$

with $E_t = \rho_{g,t} P_{g,t} \tau$ and carbon quota $\overline{E}$. The objective aggregates import cost, carbon cost (c_c) and a process-deviation penalty (c_p).

2.3.1. EAF-Specific Metallurgical Constraints

An EAF is not a generic flexible load: it cannot be ramped arbitrarily fast; its power quality couples to the bus short-circuit strength, and it carries a thermodynamic (melt-pool) state that degrades if perturbed. These are encoded as additional constraints on the random EAF power $P_{f,t}(\xi)$ and reformulated to SOC form by the

same Cantelli procedure used below, so the SOCP structure is unchanged. First, the electrode-regulation/transformer ramp limits the per-step change to a slew R_f ,

$$\mathbb{P}(|P_{f,t}(\xi) - P_{f,t-1}(\xi)| \leq R_f \tau) \geq 1 - \epsilon. \quad (10)$$

Second, EAF flicker/harmonic emission (IEC 61000-3-7 P_{st}) scales inversely with the short-circuit ratio $SCR = S_{sc}/S_f$, giving an SCR-tied power ceiling

$$\mathbb{P}(P_{f,t}(\xi) \leq P_f^{flk}) \geq 1 - \epsilon, \quad P_f^{flk} = \eta_{flk} SCR S_f / k_{st}, \quad (11)$$

with flicker-emission coefficient k_{st} and contractual headroom η_{flk} . Third, melt-pool thermal stability imposes a stage-dependent minimum-power floor and a bounded deviation from the reference profile,

$$\mathbb{P}(P_{f,t}(\xi) \geq \underline{P}_f^{(s(t))}) \geq 1 - \epsilon, \quad \mathbb{P}(|P_{f,t}(\xi) - P_{f,t}^{ref,(s(t))}| \leq \delta_f^{(s(t))}) \geq 1 - \epsilon, \quad (12)$$

where $s(t)$ is the metallurgical stage. The 45 min tap-to-tap cycle is partitioned into bore-down, main-melting and refining stages with distinct reference power and floors,

$$s(t) \in \{\text{bore-down, main melting, refining}\}, \quad \underline{P}_f^{(rf)} < \underline{P}_f^{(mm)}, \quad (13)$$

illustrated in Figure 3. The stage-varying floor is the metallurgical reason the gain must change across the heat: the same feeder excursion is admissible in refining but not during bore-down, so a single fixed gain cannot be metallurgically correct across stages—a distinction absent for a homogeneous flexible load. The process-deviation penalty c_p in (5) prices residual deviation inside the band $\delta_f^{(s)}$.




Figure 3. Stage-resolved EAF power profile over a 45 min tap-to-tap cycle (bore-down/main melting/refining), with stage floors $\underline{P}_f^{(s)}$ and transformer rating $\bar{P}_f$.

2.3.2. Cantelli SOC Reformulation

Substituting the PCE into a one-sided chance constraint and applying the distribution-free Cantelli inequality gives, for any y with mean μ_y and variance σ_y^2 ,

$$\mathbb{P}(y \leq \bar{y}) \geq 1 - \epsilon \Leftrightarrow \mu_y + \kappa(\epsilon)\sigma_y \leq \bar{y}, \quad \kappa(\epsilon) = \sqrt{\frac{1-\epsilon}{\epsilon}}. \quad (14)$$

For $P_{g,t}$, this yields the SOC constraint

$$\hat{g}_{0,t} + \kappa(\epsilon)\sqrt{\sum_{k \geq 1} \gamma_k \hat{g}_{k,t}^2} \leq \bar{P}_g, \quad (15)$$

and constraints (8)–(12) reformulate identically (the slew constraint (10) acts on the increment coefficients $\hat{f}_{k,t} - \hat{f}_{k,t-1}$). The whole problem is an SOCP with $\mathcal{O}(TN_{\text{var}}(K+1))$ coefficients; its solution returns not only the expected optimum $\hat{x}_0$ but the full coefficient set used next.

2.4. Sensitivity-Based Adaptive Setpoint Extraction

Proposition 1 (First-order PCE coefficient as analytical sensitivity). For $x(\xi)$ in (3) on a graded Jacobi basis with $\Phi_{1,i}(\xi) = \xi_i - \mathbb{E}_\rho[\xi_i]$,

$$\left. \frac{\partial x}{\partial \xi_i} \right|_{\xi^0} = \hat{x}_{1,i} \Phi'_{1,i}(\xi^0), \quad (16)$$

i.e., the partial sensitivity is proportional to the first-order PCE coefficient.

Proof Sketch. Term-by-term differentiation of (3) and evaluation at ξ^0 ; the graded basis annihilates all higher-order contributions, leaving (16) [42]. $\square$

We verified (16) post hoc on the dispatch surrogate: for the grid-import response, the first-order PCE coefficient and a centred finite-difference derivative agree to $\sim 10^{-12}$ in every zone, and the extracted sensitivities recover the closed-form values $\hat{g}_{1,w} = -P_w^{\max} = -50$, $\hat{g}_{1,s} = -P_s^{\max} = -30$. Because the energy balance is affine, the identity is exact here; for mildly nonlinear extensions, it remains an analytical sensitivity of the degree- d surrogate with a controllable residual.

What Is New

Proposition 1 is used elsewhere only internally, to assemble the variance term in (15); the first-order coefficients are then discarded. We instead read them out as the deliverable. The conventional route to a sensitivity-based gain forms and inverts the energy-hub Jacobian at every step; ours forms no Jacobian—the gain is a ratio of two numbers already in the coefficient vector. And because that ratio is built only from coefficients of a dispatch that already enforces (15) at level $1 - \epsilon$, the gain inherits the same security guarantee, unlike heuristic deviation-dependent droops [43], whose gains carry no probabilistic certificate.

Consider an EAF regulated about its reference by the droop law

$$P_{f,t} = P_{f,t}^{\text{ref}} - k_{\text{opt}}(P_{g,t}^{\text{meas}} - P_{g,t}^{\text{ref}}). \quad (17)$$

For a single source, the sensitivity ratio is $\tilde{k}_i = |\hat{f}_{1,i}/\hat{g}_{1,i}|$; for two independent sources, a grid-sensitivity-weighted aggregation reduces to the ℓ_1 form

$$\tilde{k} = \frac{\sum_i |\hat{f}_{1,i}|}{\sum_i |\hat{g}_{1,i}|}, \quad (18)$$

the unique aggregation invariant to sign permutation of the coefficients and minimising the expected ℓ_1 command penalty in (5) [42]. The dimensionless ratio is mapped to engineering scale by a zone-specific constant α_z , $k_{\text{opt}} = \alpha_z \tilde{k}$, calibrated once per zone so the median matches an engineering reference. The complete procedure is Algorithm 1.

Algorithm 1 PCE dispatch with adaptive setpoint extraction.

Require: Forecasts $\{\hat{P}_{w,t}, \hat{P}_{s,t}\}$, zone statistics, prices, quota, plant + metallurgical parameters, degree d , level $1 - \epsilon$

- 1: **for** $t = 1, \dots, T$ **do**
- 2: Identify zones (1); compute Beta parameters (2)
- 3: Build Jacobi basis; fit PCE by collocation regression
- 4: Form and solve the SOCP (5)–(13)
- 5: Extract $\hat{f}_{1,i}, \hat{g}_{1,i}$; compute $\tilde{k}$ (18), $k_{\text{opt}} = \alpha_z \tilde{k}$
- 6: Dispatch $(P_{f,t}^{\text{ref}}, k_{\text{opt}}, \text{BESS})$ to the plant
- 7: **end for**

3. Results**3.1. Test Plant, Sizing Rationale and Setup**

The plant has a single 80 MW EAF, a 50 MW wind farm, a 30 MWp PV array, a 20 MWh/10 MW grid-tie BESS, and an $\bar{P}_g = 80$ MW feeder, producing ≈ 1.0 Mt of crude steel per year at 420 kWh/t. These ratings follow mini-mill and renewable-integrated steel-plant design practice [2–5]: the 80 MW EAF and 420 kWh/t are mid-range for a single-furnace mini-mill, and the deliberately balanced 80 MW EAF/80 MW PCC/80 MW combined VRE sizing isolates the controller comparison. As discussed in Section 4, this balance is also what produces the zero-curtailment result, and a tighter-feeder configuration is identified as future work. The scheduling step is $\tau = 0.25$ h, the chance-constraint level $1 - \epsilon = 0.95$, and the PCE degree $d = 2$.

The Beta model is calibrated on the 24-month real record of Table 1. The single engineering constant per zone, α_z , is then set once so that the median extracted gain matches a feeder-regulation reference $k_{\text{base}} = \{6, 12, 17\}$ for Low/Mid/High (rising regulation duty with output); applied to the real σ_g , this gives $\alpha_z = \{39.0, 38.8, 80.9\}$ and replaces the hand-set values of earlier drafts—there are no remaining tuned parameters in the gain extraction. Two complementary evaluations are reported: (i) a controlled 90-case study (30 per zone) in which realisations are sampled from the real-calibrated zone Beta so that the four controllers are compared on identical disturbances, and (ii) an out-of-sample test (Section 3.13) in which the test scenarios are real held-out hourly forecast–realisation records rather than Beta re-samples. In the controlled study, the disturbance magnitude is $\Delta = \|\mathbf{p}^{\text{real}} - \mathbf{p}^{\text{pred}}\|_2$ (Table 2). The marginal grid factor is $\rho_g = 0.55$ tCO₂/MWh [16,44] (relaxed to a time-varying factor in Section 3.10); the carbon price is 120 CNY/tCO₂ [34,36,39]; the time-of-use tariff is 1.20/0.85/0.45 CNY/kWh (peak/shoulder/valley); and $c_p = 200$ CNY/MWh. Three benchmarks accompany the adaptive scheme: fixed $k = 20$ and $k = 15$ (engineering settings against transformer headroom) and a no-droop ($k = 0$) baseline that holds the EAF at the dispatched reference, taken as a proxy for conventional constant-power operation.

Table 2. Realised disturbance magnitude Δ (MW) by zone, 30 cases each.

Zone	Mean	Std	Min	Max
Low	5.10	0.28	4.54	5.73
Mid	11.22	0.44	10.31	12.05
High	6.98	1.14	3.60	8.61

3.2. Adaptive Setpoint Characteristics

Figure 4 and Table 3 report the extracted k_{opt} . The medians scale across zones as 6.07:11.91:17.22 ($\approx 1:2:3$), a separation no single fixed coefficient can reproduce. Notably, this ordering is not that of the disturbance magnitude: with the real statistics, the Mid zone

carries the largest absolute disturbance (11.22 MW, Table 2) yet not the largest gain, while the High zone has a smaller disturbance (6.98 MW) but the largest gain—so the optimal gain follows the operating regime, not the disturbance scale.




Figure 4. Extracted k_{opt} versus joint forecast $(\hat{p}_w, \hat{p}_s)$ across 90 cases. $\circ$ Low, $\triangle$ Mid, $\square$ High.

Table 3. Distribution of k_{opt} by zone (30 cases each). CV is the coefficient of variation; the last column is the within-zone Pearson correlation between k_{opt} and the realised disturbance Δ .

Zone	Mean	Std	Median	CV (%)	$r(k_{\text{opt}}, \Delta)$
Low	6.07	0.20	6.07	3.4	+0.16
Mid	12.00	0.41	11.91	3.4	+0.07
High	16.76	3.02	17.22	18.0	+0.59

The intra-zone coefficient of variation is small in the data-rich Low and Mid zones (3.4%), so the output is a stable engineering parameter there; it is larger in the High zone (18.0%), a direct consequence of the thin 122 h high-wind record (Table 1). The within-zone correlation with the raw disturbance is weak in Low and Mid ($r \leq 0.16$) but appreciable in High ($r = 0.59$), again reflecting the small High sample in which σ_g and Δ co-move. Read together with the cross-zone ordering above—largest disturbance in Mid, largest gain in High—the gain tracks the operating regime rather than the disturbance scale, so a deviation-dependent droop [43] keyed to observed disturbance would select the wrong gain.

3.3. Mean Tracking Performance

Tracking error is $\Delta P = |P_{f,t}^{\text{droop}} - P_{f,t}^{\text{truth}}|$, where $P_{f,t}^{\text{truth}}$ is the ideal post-realisation setpoint. Table 4 reports winner rates and mean errors.

In the Low zone, the optimum (≈ 6.1) is far below the fixed candidates, which over-react and are beaten even by the passive no-droop baseline (0.620 MW vs. 1.160/0.725); the adaptive scheme attains 0.164 MW (a 73.6% reduction vs. no droop). In the Mid zone, the optimum sits between the fixed candidates and $k = 15$ comes closest; adaptive gives the smallest mean (0.567 vs. 0.586 MW) but, as Section 3.4 shows, the Mid adaptive-vs- $k = 15$ gap is not statistically significant. In the High zone, the two fixed gains are both close to the optimum, so adaptive (0.540 MW) leads $k = 20$ by only 5.4% and $k = 15$ by 3.6%. The recurring pattern is that the best fixed gain is zone-specific and unknown in advance, which the adaptive scheme resolves automatically.

Table 4. Performance by zone, 30 cases each. “Winner rate” is the percentage of cases in which a strategy attains the lowest tracking error; mean error in MW. Bold indicates the best value in each row.

Band	Metric	Adaptive	$k = 20$	$k = 15$	No Droop
Low	Winner rate (%)	96.7	0.0	0.0	3.3
	Mean err. (MW)	0.164	1.160	0.725	0.620
Mid	Winner rate (%)	56.7	0.0	43.3	0.0
	Mean err. (MW)	0.567	0.890	0.586	2.403
High	Winner rate (%)	46.7	23.3	30.0	0.0
	Mean err. (MW)	0.540	0.571	0.561	1.710

3.4. Confidence Intervals and Statistical Significance

To avoid over-reading point estimates, we report bootstrap 95% confidence intervals on the mean tracking errors (5000 resamples, Table 5) and paired Wilcoxon signed-rank tests of adaptive vs. each benchmark across the 30 paired cases per zone (Table 6).

Seven of the nine comparisons are significant at the 5% level; the two exceptions are the Mid and High zones against $k = 15$ ($p = 0.56$ and $p = 0.24$), where the schemes are statistically indistinguishable because $k = 15$ happens to sit close to the optimum in those zones. We therefore do not claim universal superiority: the advantage is large and significant where the fixed gain is mistuned (the Low zone, and all comparisons against no droop), and modest or indistinguishable where a fixed gain is accidentally well-matched to the prevailing zone.

Table 5. Mean tracking error with bootstrap 95% CI (MW).

Zone	Strategy	Mean	CI Lower	CI Upper
Low	Adaptive	0.164	0.127	0.201
Low	$k = 20$	1.160	1.089	1.234
Low	No droop	0.620	0.552	0.685
Mid	Adaptive	0.567	0.478	0.657
Mid	$k = 15$	0.586	0.518	0.656
Mid	No droop	2.403	2.215	2.587
High	Adaptive	0.540	0.446	0.637
High	$k = 20$	0.571	0.474	0.669
High	No droop	1.710	1.419	1.994

Table 6. Paired Wilcoxon signed-rank tests, adaptive vs. benchmark (30 cases per zone). Positive median difference favours the adaptive scheme. Bold “no” marks the comparisons that are not statistically significant at the 5% level.

Zone	Comparison	Median Diff. (MW)	p	Sig. (5%)
Low	vs. $k = 20$	+0.990	1.9×10^{-9}	yes
Low	vs. $k = 15$	+0.582	1.9×10^{-9}	yes
Low	vs. no droop	+0.520	9.3×10^{-9}	yes
Mid	vs. $k = 20$	+0.341	4.7×10^{-8}	yes
Mid	vs. $k = 15$	+0.043	0.56	no
Mid	vs. no droop	+1.909	1.9×10^{-9}	yes
High	vs. $k = 20$	+0.014	1.9×10^{-2}	yes
High	vs. $k = 15$	+0.017	0.24	no
High	vs. no droop	+1.400	1.9×10^{-9}	yes

3.5. Tail Behaviour

Per-heat metallurgical cost is dominated by rare large excursions, so we report the 95th-percentile and worst-case tracking error (Table 7); the full disturbance-tracking scatter is in Figure 5.

Table 7. Tail tracking error: 95th percentile $q_{0.95}$ and worst case $\Delta P^{\max}$ (MW). Bold indicates the best value in each row.

Band	Quantile	Adaptive	$k = 20$	$k = 15$	No Droop
Low	$q_{0.95}$	0.320	1.585	1.048	0.856
	$\Delta P^{\max}$	0.377	1.711	1.251	0.940
Mid	$q_{0.95}$	1.038	1.226	0.917	3.149
	$\Delta P^{\max}$	1.176	1.252	0.960	3.341
High	$q_{0.95}$	1.058	1.052	1.114	2.623
	$\Delta P^{\max}$	1.127	1.140	1.150	2.798




Figure 5. EAF tracking error vs. disturbance magnitude by zone, in a three-row layout. Markers: one per case; lines: within-zone regression. (a) Low; (b) Mid; (c) High.

The decisive tail benefit is the bounding of the worst case: the adaptive worst-case tracking error stays ≤ 1.18 MW in every zone, against up to 3.34 MW for no droop ($\approx 4\%$ of EAF rating in a single 15 min interval, enough to disturb slag-foaming stability), and is the lowest of all strategies in the Low and High zones. In the Mid zone, the well-

matched $k = 15$ edges the adaptive scheme at the 95th percentile (0.917 vs. 1.038 MW), consistent with the non-significant Mid mean difference; the adaptive scheme is never the worst and never diverges the way a mistuned fixed gain does. The bounded tail is the empirical face of the inherited 95% feeder-security guarantee.

3.6. Benchmarking Against Model Predictive Control

We compare the closed-form adaptive gain with a receding-horizon/certainty-equivalent model-predictive controller (MPC) on the same 90 cases. A quadratic certainty-equivalent MPC that trades feeder deviation against setpoint movement reduces, in this affine setting, to a droop with gain $g = a_{fe}/(a_{fe} + a_{mv})$; with constant cost weights, this is a fixed, zone-blind gain (“MPC const”), whereas scaling the feeder weight by the PCE-predicted σ_g gives a σ -aware variant (“MPC σ ”). Table 8 reports the result.

Table 8. Mean tracking error (MW) and effective gain (p.u.) of the adaptive scheme vs. MPC variants and the best fixed gain. Bold indicates the best value in each row.

Zone	Adapt.	Mean Tracking Error (MW)			$k = 20$	Effective Gain (p.u.)		
		MPC σ	MPC Const			Adapt.	MPC σ	MPC Const
Low	0.164	0.852	0.810	1.160	6.07	16.39	15.92	
Mid	0.567	0.771	0.630	0.890	12.00	17.82	15.62	
High	0.540	0.533	0.553	0.571	16.76	17.03	15.92	

The constant-weight MPC collapses to a near-constant gain (≈ 15.6 – 15.9 across all zones) and behaves like a fixed gain, performing poorly in the Low zone (0.810 MW). The σ -aware MPC improves on it and is essentially tied with the adaptive scheme in the High zone (0.533 vs. 0.540 MW), but trails it in the Low and Mid zones because it lacks the head-room/zone normalisation and—unlike the closed-form ratio (18)—requires an explicit on-line weight model. The adaptive scheme thus attains MPC-class tracking with no online optimisation and with the inherited security guarantee that neither MPC variant carries (see Figure 6).





Figure 6. Mean tracking error by zone for the adaptive scheme, the two MPC variants, and the best fixed gain.

3.7. Sensitivity to Wind–PV Correlation

The conditional-independence assumption is tested by drawing the realised wind and PV outcomes from a Gaussian copula with correlation $\rho \in \{0, 0.15, 0.30, 0.50\}$, while the dispatch model retains independence (Table 9, Figure 7).

Table 9. Effect of wind–PV copula correlation on adaptive tracking and feeder-limit violation (model assumes independence).

ρ	Adaptive Track. (MW)	Adaptive Feeder Viol. (%)	$k = 20$ Track. (MW)
0.00	0.444	0.0	0.942
0.15	0.486	0.0	0.919
0.30	0.502	0.0	0.931
0.50	0.522	0.0	0.956

**Figure 7.** Adaptive and fixed-gain tracking, and adaptive feeder-violation rate, versus imposed wind–PV correlation ρ .

Although the measured correlation is only $r = 0.015$ (Section 2.2), we stress-test far beyond it: adaptive tracking degrades gracefully ($0.444 \rightarrow 0.522$ MW as $\rho : 0 \rightarrow 0.5$, a 17% rise) and the feeder limit is never violated, so the framework is robust to a mis-specified independence assumption even at correlations an order of magnitude above the observed value; a copula extension of the model is nonetheless noted as future work for synoptically correlated sites.

3.8. Goodness-of-Fit of the Beta Model

Table 10 reports the goodness-of-fit of the Beta family against a truncated-normal alternative, computed on the measured realised-output records (Section 2.2, Open-Meteo ERA5/IFS) by the Kolmogorov–Smirnov (KS) statistic. The Beta clearly wins in the data-rich Low zones (wind 0.127 vs. 0.291; PV 0.060 vs. 0.162) and is competitive in the Mid and High PV zones; the truncated normal is marginally better only for Mid-zone wind, and both fits are poor for High-zone wind (KS 0.50/0.79), a direct consequence of the 122-sample limit there. The Beta family is thus the better bounded-support model overall on real data, while the High-wind fit is flagged as data-limited rather than over-claimed.

Table 10. Goodness-of-fit on measured realised-output records: KS statistic for Beta vs. truncated-normal fits per zone (n from Table 1). Bold indicates the best value in each row.

Zone	Source	KS (Beta)	KS (Trunc. Normal)
Low	Wind	0.127	0.291
Mid	Wind	0.185	0.056
High	Wind	0.500	0.795
Low	PV	0.060	0.162
Mid	PV	0.046	0.053
High	PV	0.075	0.060

3.9. Energy, Carbon and Feeder-Security Outcomes

Table 11 gives shift-averaged outcomes across all 90 cases. The adaptive scheme empirically satisfies the feeder limit in 100% of realisations (0.0% violation), confirming that the 95% chance-constraint design target is met with margin and inherited by the controller. Renewable curtailment is zero for all strategies because the balanced sizing lets the EAF absorb all on-site VRE.

Table 11. Shift-averaged outcomes (90 cases). Specific emissions per tonne of crude steel (≈ 880 t/shift); feeder-violation rate is the share of realisations breaching $\bar{P}_g/\underline{P}_g$. Bold indicates the best value in each row.

Metric	Adaptive	$k = 20$	$k = 15$	No Droop
Grid imports (MWh)	206.8	201.8	204.2	216.3
Operational cost (10^3 CNY)	179.0	177.0	177.8	182.2
Emissions (tCO ₂)	113.8	111.0	112.3	119.0
Specific emissions (kgCO ₂ /t)	129.3	126.1	127.6	135.2
Feeder violation (%)	0.0	0.0	0.0	0.0

Relative to no droop, the adaptive scheme cuts specific emissions by 4.4% ($135.2 \rightarrow 129.3$ kgCO₂/t) and grid imports by 4.4%; relative to the best fixed gain ($k = 20$), it is 2.5% higher on both, a gap distributed across zones rather than concentrated in any regime. We are explicit about attribution: the dominant carbon advantage—the $\approx 93\%$ gap to the BF–BOF benchmark of ≈ 1760 kgCO₂/t (Figure 8)—is a structural property of the renewable-integrated EAF route and the electricity emission factor, not of the proposed controller; this electricity-only figure is on a narrower boundary than the cradle-to-gate BF–BOF value, and the comparison should be read as route-level. What the controller contributes is the residual within-EAF layer: ≈ 6 kgCO₂/t (≈ 6 kt CO₂/yr for a 1 Mt plant) relative to no droop, plus the tail-tracking and feeder-security benefits that do not appear in shift averages [1,8,11,15].




Figure 8. Shift-averaged specific CO₂ emissions across strategies, with the long-route BF–BOF benchmark (≈ 1760 kgCO₂/t). Broken vertical axis preserves intra-strategy resolution.

3.10. Rigorous Time-Varying CO₂ Accounting

The flat factor $\rho_g = 0.55$ is a simplification. Table 12 recomputes specific emissions with an hour-dependent marginal factor (0.32 at the solar-marginal midday, 0.82 at the evening coal-marginal peak, 0.58–0.62 otherwise) applied to each case by its dispatch hour.

Table 12. Specific emissions (kgCO₂/t) under a flat vs. a time-varying marginal grid emission factor.

Strategy	Flat Factor	Time-Varying Factor
Adaptive	129.3	136.4
$k = 20$	126.1	133.1
$k = 15$	127.6	134.6
No droop	135.2	142.7

The time-varying factor raises all figures by $\approx 5\%$ because residual imports cluster in higher-carbon peak hours; the ranking and the adaptive-vs-no-droop gap (now 6.3 kgCO₂/t) are preserved, showing the flat factor understates absolute emissions but does not bias the controller comparison.

3.11. Computational Performance

The SOCP is solved with the CLARABEL 0.11.1 interior-point solver through CVXPY 1.8.2 (Python 3.14, NumPy 2.4.3, SciPy 1.17.1). Over 40 randomly drawn dispatch instances, the per-step solve time is 12.9 ms on average, 33.5 ms at the 95th percentile, and 46.7 ms worst case on a commodity laptop; a full 24 h horizon (96 steps) solves in 1.41 s. Against the 900 s real-time budget of a 15 min step, this is nearly five orders of magnitude of headroom, and the degree-two, two-germ expansion needs only six coefficients per variable, so the framework is amply real-time-capable and scales linearly in the horizon.

3.12. Twenty-Four-Hour Operating Case

Following the reviewers, we extend the evaluation to a full 24 h day (96 steps) with a diurnal PV profile and a mean-reverting wind profile. Table 13 reports the daily outcome. The adaptive scheme keeps mean tracking at 0.52 MW—far below the no-droop 2.27 MW—while specific emissions (161.5 kgCO₂/t over a realistic diurnally weighted day, higher than the balanced-zone average because night hours carry no PV) remain within 2.2% of the best fixed gain. As in the controlled experiment, a fixed gain can appear competitive on a particular day only because that day's regime mix happens to match its tuning—which cannot be known in advance.

Table 13. Twenty-four-hour operating outcome (96 steps).

Strategy	Grid (MWh)	Mean Track. (MW)	Spec. Emis. (kgCO ₂ /t)
Adaptive	775.3	0.52	161.5
$k = 20$	758.7	0.87	158.1
$k = 15$	766.8	0.55	159.8
No droop	820.2	2.27	170.9

3.13. Out-of-Sample Validation on Real Records

The preceding subsections compare controllers on disturbances sampled from the (real-calibrated) Beta model. As a genuine out-of-sample test, we now hold the Beta model fixed inside the dispatch but draw the test scenarios directly from real held-out hourly forecast–realisation records (Open-Meteo ERA5/IFS, daytime, up to 150 records per zone): the forecast ($\hat{p}_w, \hat{p}_s$) comes from the archived IFS forecast and the realisation (p_w, p_s) from ERA5, so neither the disturbance nor the feeder excursion is a Beta re-sample. Table 14 reports the result and Figure 9 visualises the mean.

The out-of-sample test corroborates the controlled findings in the data-rich regimes and is candid about the data-limited one. In the Low zone, the adaptive scheme is decisively best (mean 0.146 vs. 0.455–0.846; $q_{0.95}$ 0.533 vs. 1.53–2.44), and in the Mid zone, it leads on the mean (0.423, essentially tied with $k = 15$) and on the 95th percentile

(1.840 vs. 2.455 for $k = 15$ and 8.337 for no droop). In the High zone, however, the adaptive scheme is outperformed by both fixed gains (mean 0.391 vs. 0.150/0.177): the gain calibrated on only 122 high-wind hours over-reacts to the few genuine high-wind realisations. The feeder limit is satisfied in 100% of real-record realisations for all strategies. We read this honestly: the framework generalises out of sample where the calibration data are plentiful, and the High-zone shortfall is a data-availability problem—more high-wind records, or a site with a higher wind regime, would be needed to calibrate that zone reliably—rather than a failure of the method itself.

Table 14. Out-of-sample tracking error (MW) on real ERA5/IFS records: mean and 95th percentile per zone. Bold indicates the best value in each row.

Band	Metric	Adaptive	$k = 20$	$k = 15$	No Droop
Low	Mean	0.146	0.846	0.536	0.455
	$q_{0.95}$	0.533	2.440	1.528	1.633
Mid	Mean	0.423	0.773	0.438	1.806
	$q_{0.95}$	1.840	3.200	2.455	8.337
High	Mean	0.391	0.150	0.177	0.775
	$q_{0.95}$	2.333	0.477	0.873	4.608




Figure 9. Out-of-sample mean tracking error on real ERA5/IFS forecast–realisation records, by zone.

4. Discussion

The $\approx 1:2:3$ separation of k_{opt} across zones (Table 3) shows that the optimal EAF–feeder coupling is not approximately constant under heteroscedastic uncertainty but varies multiplicatively with the prevailing forecast level. A fixed-gain controller can be optimally tuned for at most one zone; in the others, its mis-tuning produces quantifiable degradation, up to a near seven-fold increase in mean tracking error in the Low zone for $k = 20$. Because the within-zone correlation between the gain and the raw disturbance is weak ($|r| \leq 0.21$) while the Mid and High zones overlap in absolute disturbance, the optimal gain is a function of the operating regime, not the disturbance scale; deviation-dependent droops, which key on observed disturbance, are therefore calibrated to the wrong signal.

At the same time, we are careful not to overclaim. The significance analysis (Table 6) shows the adaptive advantage is large and significant where a fixed gain is mistuned but statistically indistinguishable from $k = 15$ in the Mid zone, and on shift-averaged grid imports and cost, the post hoc-optimal $k = 20$ is actually 1.5% better than the adaptive scheme. The honest reading is therefore not that the adaptive scheme dominates on every metric, but that it is robust: it removes the wagered commitment to a single zone that fixed tuning requires, it bounds the tail in every regime, and it never collapses the way a mis-

tuned fixed gain does in an off-design zone. The value is a mean–variance improvement (negligible expected cost, large reduction in operational and tail variance) rather than uniform superiority.

The strongest reviewer challenge was whether the method is merely a recombination of mature power-system tools applied to a load that is not materially different from a generic flexible load. We address this on two levels. Methodologically, the contribution is the structural identity at the dispatch–controller seam: the first-order PCE coefficients that PCE-SOPF uses internally to tighten a variance term are instead read out as the controller gain (Proposition 1), so the gain is obtained with no Jacobian inversion and provably inherits the dispatch’s chance-constraint guarantee—a property that heuristic droops lack. Physically, the EAF is distinguished from a generic flexible load by constraints with no flexible-load analogue: a hard electrode/transformer slew (10), a flicker ceiling tied to the bus short-circuit ratio (11), a melt-pool stability floor and band (12), and a stage-varying tap-to-tap profile (13) that makes the metallurgically admissible gain change within a single heat (Figure 3). We note honestly that in the present balanced-sizing case study, these constraints are latent—curtailment is zero and tracking stays ≤ 1.18 MW, so they are not numerically binding—and that their role is to make the formulation a genuine EAF model and to define the envelope inside which the inherited guarantee is meaningful. They become binding for plants with tighter feeder/PCC limits or weaker buses, which we identify as the natural next case study.

The uncertainty model is now grounded in real data rather than asserted: the zone statistics, the wind–PV error correlation ($r = 0.015$) and the goodness-of-fit (Table 10) are all measured on a 24-month ERA5/IFS record, and Section 3.13 validates the controller out-of-sample on real held-out forecast–realisation records. This addresses the most serious concern about a purely synthetic study, with one honest qualification: at the chosen coastal site, sustained high wind is rare (122 h in 24 months), so the High-zone calibration is thin, and the out-of-sample High-zone performance is correspondingly weaker. A site with a stronger wind regime, or a longer record, would tighten that zone—a data-availability matter rather than a limitation of the method.

Two classes of claim nonetheless still require measured data for definitive closure and are scoped as future work. The present study uses no measured EAF operating data: the heat profile, stage floors and transformer/electrode limits follow design practice and the cited literature, and the controller’s closed-loop behaviour is evaluated against a re-solved deterministic dispatch rather than a physical furnace. Whether a real EAF transformer and electrode regulator realise k_{opt} within a 15 min step, and whether bounded MW tracking preserves slag-foaming stability, can be established only by hardware-in-the-loop tests on a scaled EAF testbed, which we plan as the validation step beyond this proof of concept; the idealised assumptions of fast communication and immediate actuation likewise abstract away latency and measurement noise that a pilot study must quantify.

Finally, the carbon results should be read with the right attribution. The $\approx 93\%$ reduction relative to the BF–BOF route is structural, arising from the renewable-integrated EAF route and the grid emission factor, and is on an electricity-only boundary narrower than the cradle-to-gate BF–BOF figure; it is not a consequence of the controller. The controller’s own contribution is the within-EAF residual of $\approx 4.4\%$ (vs. no droop) together with the bounded tail and guaranteed feeder security, achieved with no additional capital—an instance of the broader argument that near-term industrial decarbonisation is driven as much by intelligent coordination of installed assets as by new abatement hardware [9,10,14–16].

5. Conclusions

This paper developed a PCE-based chance-constrained dispatch and closed-form adaptive setpoint-control framework for renewable-integrated EAF steelmaking. Joint wind–PV forecast errors are modelled by a zone-wise Beta parameterisation; the dispatch is propagated through a degree-two PCE on a Jacobi basis and reformulated as a second-order cone programme by the Cantelli inequality, with EAF-specific metallurgical constraints—electrode/transformer slew, short-circuit-ratio-tied flicker, stage-dependent melt-power floor, and a multi-stage tap-to-tap profile—embedded by the same procedure. The controller gain is extracted in closed form from first-order PCE coefficients, with no Jacobian inversion, and therefore inherits the dispatch’s 95% feeder-security guarantee; the sensitivity identity underlying this extraction is verified to machine precision.

Calibrated on 24 months of real ERA5/IFS data, the extracted gain scales across the Low–Mid–High zones (medians 6.07, 11.91, 17.22 p.u.) and follows the operating regime rather than the disturbance magnitude, so no fixed gain is jointly optimal across the operating envelope. The adaptive scheme keeps worst-case tracking below 1.18 MW in every zone against up to 3.34 MW for a no-droop baseline, satisfies the feeder limit in 100% of realisations, matches model-predictive control while needing no online optimisation, and remains robust to wind–PV correlation up to $\rho = 0.5$. Its improvement over the benchmarks is statistically significant in seven of nine zone-versus-benchmark comparisons, the exceptions being the Mid and High zones, where a fixed gain happens to be well-matched; an out-of-sample test on real held-out records confirms a decisive advantage in the data-rich Low and Mid zones and, candidly, a shortfall in the data-limited High zone. The framework’s value is thus regime-robustness and bounded tail risk rather than uniform superiority. On the carbon side, it lowers within-EAF specific emissions by 4.4% relative to no droop, while the order-of-magnitude advantage over the long route is structural to the renewable-integrated EAF rather than controller-driven.

The uncertainty model is calibrated and validated out-of-sample on real wind/PV records, but the controller-performance and metallurgical-integrity findings remain simulation-based and are intended as a proof of concept. The priorities for definitive validation are richer high-wind data to strengthen the High-zone calibration, incorporation of measured EAF heat data, a tighter-feeder case study that exercises the metallurgical constraints at their bounds, and hardware-in-the-loop experiments on a scaled EAF testbed.

Author Contributions: Conceptualisation, C.X.; methodology, C.X.; software, C.X.; validation, C.X., Y.K. and Y.Z.; formal analysis, C.X. and Y.Z.; investigation, C.X. and Y.K.; resources, C.X. and Y.Z.; data curation, C.X. and Y.K.; writing—original draft preparation, C.X.; writing—review and editing, C.X., Y.K. and Y.Z.; visualisation, Y.K.; supervision, C.X.; project administration, C.X.; funding acquisition, C.X. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported in part by the 2026 Decision Consulting Research Project of the Regional Institutional Research Centre, Hebei GEO University, “Identification of green development potential and spatial optimisation in mountainous regions based on remote-sensing data and deep learning” (grant No. QYZDYJZX202618); and by the 2025 Shandong Provincial Philosophy and Social Science Innovation and Development Research Project “Mechanism of the impact of artificial intelligence applications on the gender structure of the labour market: an empirical analysis based on Shandong Province” (grant no. Z-FZY25026).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The zone-wise Beta statistics (Table 1), the wind–PV error correlation, the goodness-of-fit (Table 10) and the out-of-sample test (Section 3.13) are derived from openly available data retrieved through the Open-Meteo API for 35.96° N, 120.20° E over 2023–2024 at hourly resolution: the ERA5 reanalysis (realisation) via the Open-Meteo Historical Weather (Archive) API and the archived ECMWF-IFS day-ahead forecast via the Open-Meteo Historical Forecast API (<https://open-meteo.com>). The numerical case-study data, the companion Python notebook that reproduces every table and figure, and the data-retrieval and processing scripts are available from the corresponding author upon reasonable request.

Acknowledgments: The authors acknowledge the supportive research environment of the International Business School and the School of Smart City Engineering at Qingdao Huanghai University.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

BESS	Battery energy storage system
BF–BOF	Blast furnace–basic oxygen furnace
CI	Confidence interval
CV	Coefficient of variation
EAF	Electric arc furnace
KS	Kolmogorov–Smirnov
MPC	Model predictive control
PCC	Point of common coupling
PCE	Polynomial chaos expansion
PV	Photovoltaic
SCR	Short-circuit ratio
SOC(P)	Second-order cone (programme)
SOPF	Stochastic optimal power flow
VRE	Variable renewable energy

References

1. Milford, R.L.; Pauliuk, S.; Allwood, J.M.; Müller, D.B. The roles of energy and material efficiency in meeting steel industry CO₂ targets. *Environ. Sci. Technol.* **2013**, *47*, 3455–3462. [[CrossRef](#)] [[PubMed](#)]
2. Yi, X.; Yue, Q.; Dou, Z.; Bu, Q. Enhancing efficiency in electric arc furnace steelmaking: A multi-objective optimization approach using the non-dominated sorting genetic algorithm II. *Steel Res. Int.* **2025**, *96*, 2400370.
3. Wang, J.; Fang, Q.; Lu, P.; Huang, Z.; Zhang, J.; Zhang, H.; Ni, H. An integrated model for coordinating multi-step scrap allocation with economy in long-process steelmaking. *J. Sustain. Metall.* **2025**, *11*, 3220–3239. [[CrossRef](#)]
4. Yuan, F.; Liu, X.; Xu, A.; Li, X. Effects of the scrap steel ratio and bottom-blowing process parameters on the fluid flow characteristics in a physical model of a steelmaking converter. *Metals* **2026**, *16*, 160. [[CrossRef](#)]
5. Zawodnik, V.; Gruber, A.; Kienberger, T. Fueling industrial flexibility: Discrete-time dispatch optimization of electric arc furnaces. *Energies* **2025**, *18*, 4838. [[CrossRef](#)]
6. Quader, M.A.; Ahmed, S.; Ghazilla, R.A.R.; Ahmed, S.; Dahari, M. A comprehensive review on energy efficient CO₂ breakthrough technologies for sustainable green iron and steel manufacturing. *Renew. Sustain. Energy Rev.* **2015**, *50*, 594–614. [[CrossRef](#)]
7. Quader, M.A.; Ahmed, S.; Dawal, S.Z.; Nukman, Y. Present needs, recent progress and future trends of energy-efficient ultra-low carbon dioxide (CO₂) steelmaking (ULCOS) program. *Renew. Sustain. Energy Rev.* **2016**, *55*, 537–549. [[CrossRef](#)]
8. Conejo, A.N.; Birat, J.-P.; Dutta, A. A review of the current environmental challenges of the steel industry and its value chain. *J. Environ. Manag.* **2020**, *259*, 109782. [[CrossRef](#)]
9. Bataille, C.; Åhman, M.; Neuhoﬀ, K.; Nilsson, L.; Fischedick, M.; Lechtenböhmer, S.; Solano-Rodriguez, B.; Denis-Ryan, A.; Stiebert, S.; Waisman, H.; et al. A review of technology and policy deep decarbonization pathway options for making energy-intensive industry production consistent with the Paris agreement. *J. Clean. Prod.* **2018**, *187*, 960–973. [[CrossRef](#)]
10. Cavaliere, P. Barriers to steel decarbonization. *Steel Res. Int.* **2026**, 2501317.

11. Karakaya, E.; Nuur, A.; Assbring, L. Potential transitions in the iron and steel industry in Sweden: Towards a hydrogen-based future? *J. Clean. Prod.* **2018**, *195*, 651–663. [\[CrossRef\]](#)
12. Rehfeldt, M.; Worrell, E.; Eichhammer, W.; Fleiter, T. A review of the emission reduction potential of fuel switch towards biomass and electricity in European basic materials industry until 2030. *Renew. Sustain. Energy Rev.* **2020**, *120*, 109672. [\[CrossRef\]](#)
13. Suopajarvi, H.; Umeki, K.; Mousa, E.; Hedayati, A.; Romar, H.; Kemppainen, A.; Wang, C.; Phounglamcheik, A.; Tuomikoski, S.; Norberg, N.; et al. Use of biomass in integrated steelmaking—Status quo, future needs and comparison to other low-CO₂ steel production technologies. *Appl. Energy* **2018**, *213*, 384–407. [\[CrossRef\]](#)
14. Lo Basso, G.; Pastore, L.M.; Sgaramella, A.; Mojtahed, A.; Ciancio, A.; Massulli, A.R.; Vespasiano, D.; Santoli, L. Status and perspectives of hydrogen role for decarbonising industry: A comprehensive review. *Renew. Sustain. Energy Rev.* **2025**, *224*, 116083. [\[CrossRef\]](#)
15. Putera, A.D.P.; Putra, H.P.; Karuana, F.; Ghazidin, H.; Kuswa, F.M.; Suyatno, S.; Aziz, M.; Prismantoko, A. Toward decarbonization in steelmaking: A review on direct reduced iron using biomass-based syngas. *Sustain. Energy Technol. Assess.* **2025**, *82*, 104554. [\[CrossRef\]](#)
16. Tong, H.; Liu, J.; Wang, Y.; An, H.; Song, G.; Hou, Y.; Tang, X.; Yi, H. The transformation of China's iron and steel industry under carbon constraints: A balance between economy, emission reduction and technical feasibility. *J. Environ. Sci.* **2026**, *165*, 860–876. [\[CrossRef\]](#)
17. Ma, N.; Zhao, P.; Xu, W.; Liu, A.; Wang, J. Assessment on the coupling effect and economic viability of a novel hydrogen-electric coupled energy storage system combining hydrogen-fueled CAES and oxy-combustion CCES driven by a power-to-gas device: A pathway for synergistic enhancement. *J. Energy Storage* **2025**, *135*, 118178. [\[CrossRef\]](#)
18. Ju, C.; Ding, T.; Jia, W.; Mu, C.; Zhang, H.; Sun, Y. Two-stage robust unit commitment with the cascade hydropower stations retrofitted with pump stations. *Appl. Energy* **2023**, *334*, 120675. [\[CrossRef\]](#)
19. Phu, P.V.; Huy, T.H.B.; Park, S.; Kim, D. An IGDT approach for the multi-objective framework of integrated energy hub with renewable energy sources, hybrid energy storage systems, and biomass-to-hydrogen technology. *J. Energy Storage* **2024**, *89*, 111488. [\[CrossRef\]](#)
20. Lu, G.; Bai, X.; Qi, X.; Zhang, F.; Yan, X.; Feng, J.; Ju, L. Data-model fusing driven robust configuration optimization model and decision-making method for renewable energy generation base considering electric–hydrogen conversion. *Renew. Energy* **2025**, *241*, 122320. [\[CrossRef\]](#)
21. Yang, M.; Wang, J.; Chen, Y.; Zeng, Y.; Su, X. Data-driven robust optimization scheduling for microgrid day-ahead to intra-day operations based on renewable energy interval prediction. *Energy* **2024**, *313*, 134058. [\[CrossRef\]](#)
22. Xue, Q.; Wang, A.; Jiang, S.; Wang, Z.; Yang, Y.; Cheng, Y.; Zheng, Z. Stochastic optimization of energy systems configuration for nearly-zero energy buildings considering load uncertainties. *Renew. Energy* **2025**, *243*, 122610. [\[CrossRef\]](#)
23. Alamir, N.; Kamel, S.; Abdelkader, S.M. Stochastic multi-layer optimization for cooperative multi-microgrid systems with hydrogen storage and demand response. *Int. J. Hydrogen Energy* **2025**, *100*, 688–703. [\[CrossRef\]](#)
24. Han, M.; Sun, R.; Feng, P.; Hua, E. Unveiling characteristics and determinants of China's wind power geographies towards low-carbon transition. *J. Environ. Manag.* **2023**, *331*, 117215. [\[CrossRef\]](#)
25. Abdilahi, A.M.; Mustafa, M.W.; Abujarad, S.Y.; Mustapha, M. Harnessing flexibility potential of flexible carbon capture power plants for future low carbon power systems: Review. *Renew. Sustain. Energy Rev.* **2018**, *81*, 3101–3110. [\[CrossRef\]](#)
26. Li, J.; Zhang, C.; Davidson, M.R.; Lu, X. Assessing the synergies of flexibly-operated carbon capture power plants with variable renewable energy in large-scale power systems. *Appl. Energy* **2025**, *377*, 124459. [\[CrossRef\]](#)
27. Yu, X.; Zhang, S.; Liu, L.; Du, J. Optimal design and scheduling of carbon capture power plant based on uncertainty decision-making methods. *J. Clean. Prod.* **2022**, *380*, 134852. [\[CrossRef\]](#)
28. Yu, G.; Zhang, Z.; Cui, G.; Dong, Q.; Wang, S.; Li, X.; Shen, L.; Yan, H. Low-carbon economic dispatching strategy based on feasible region of cooperative interaction between wind-storage system and carbon capture power plant. *Renew. Energy* **2024**, *228*, 120706. [\[CrossRef\]](#)
29. Liang, S.; Lou, S.; Zhu, Z.; Sun, Y. Flexible operation of solar-assisted carbon capture power plants considering interval-enhanced CVaR. *Appl. Energy* **2025**, *378*, 124850. [\[CrossRef\]](#)
30. Guo, X.; Liao, Y.; Li, G.; Han, Y.; Ma, X. Low-carbon economic dispatch of photovoltaic–carbon capture power plant considering deep peak regulation. *J. Clean. Prod.* **2023**, *420*, 138418. [\[CrossRef\]](#)
31. Yi, T.; Zhang, C. Scheduling optimization of a wind power-containing power system considering the integrated and flexible carbon capture power plant and P2G equipment under demand response and reward and punishment ladder-type carbon trading. *Int. J. Greenh. Gas Control* **2023**, *128*, 103955. [\[CrossRef\]](#)
32. Wang, H.; Feng, T.; Kong, J.; Cui, M.; Xu, M. Grappling with the trade-offs of carbon emission trading and green certificate: Achieving carbon neutrality in China. *J. Environ. Manag.* **2024**, *360*, 121101. [\[CrossRef\]](#)
33. Li, J.; He, X.; Niu, W.; Liu, X. Analysis of the joint trading of local green power certificates, carbon emissions rights, and electricity considering demand flexibility. *Int. J. Electr. Power Energy Syst.* **2024**, *155*, 109653. [\[CrossRef\]](#)

34. Zhang, X.; Guo, X.; Zhang, X. Assessing the policy synergy among power, carbon emissions trading and tradable green certificate market mechanisms on strategic GENCOs in China. *Energy* **2023**, *278*, 127833. [[CrossRef](#)]
35. Yan, S.; Wang, W.; Li, X.; Lv, H.; Fan, T.; Aikepaer, S. Stochastic optimal scheduling strategy of cross-regional carbon emissions trading and green certificate trading market based on Stackelberg game. *Renew. Energy* **2023**, *219*, 119268. [[CrossRef](#)]
36. Guo, X.; Wang, L.; Ren, D. Optimal scheduling model for virtual power plant combining carbon trading and green certificate trading. *Energy* **2025**, *318*, 134750. [[CrossRef](#)]
37. Chen, H.; Wu, H.; Li, H.; Chen, F.; Shui, S.; Li, Z. Bi-level optimal scheduling of integrated energy systems considering incentive-based demand response and green certificate-carbon trading mechanisms. *Energy Rep.* **2025**, *13*, 330–344. [[CrossRef](#)]
38. Liu, D.; Luo, Z.; Qin, J.; Wang, H.; Wang, G.; Li, Z.; Zhao, W.; Shen, X. Low-carbon dispatch of multi-district integrated energy systems considering carbon emission trading and green certificate trading. *Renew. Energy* **2023**, *218*, 119312. [[CrossRef](#)]
39. Wang, X.; Chen, H.; Tong, X.; Gao, Y.; Pan, P.; Liu, W. Optimal scheduling of a multi-energy complementary system simultaneously considering the trading of carbon emission and green certificate. *Energy* **2024**, *310*, 133212. [[CrossRef](#)]
40. Ju, L.; Yin, Z.; Lu, X.; Yang, S.; Li, P.; Rao, R.; Tan, Z. A tri-dimensional equilibrium-based stochastic optimal dispatching model for a novel virtual power plant incorporating carbon capture, power-to-gas and electric vehicle aggregator. *Appl. Energy* **2022**, *324*, 119776. [[CrossRef](#)]
41. Ju, L.-W.; Li, H.-H.; Yu, X.; Dong, F.-G.; Song, X.-Y.; Tan, Z. A combined external delivery optimization model for wind power and thermal power considering carbon trading and tradable green certificates based on credibility theory. *Int. J. Electr. Power Energy Syst.* **2016**, *78*, 51–60. [[CrossRef](#)]
42. Du, H.; Lekić, A. SOPF-based adaptive droop control for hybrid AC–HVDC grids under offshore wind uncertainty. *arXiv* **2026**, arXiv:2605.05992.
43. Wang, Y.; Li, B.; Zhou, Z.; Chen, Z.; Wen, W.; Li, X.; Wang, C. DC voltage deviation-dependent voltage droop control method for VSC-MTDC systems under large disturbances. *IET Renew. Power Gener.* **2020**, *14*, 891–896. [[CrossRef](#)]
44. Zhu, J.; Li, Z.; Jiang, X.; Jiang, P. Synergistic effects of air pollution and carbon reduction policies in China's iron and steel industry. *Energies* **2025**, *18*, 5379. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.