

Extreme Climate Risk and Corporate Energy Management: A Double Machine Learning Approach

Jiawei Ma

p2626244@mpu.edu.mo

MPI GTR Center: Macao Polytechnic University Gaming Teaching and Research Center

<https://orcid.org/0009-0000-6423-3039>

Research Article

Keywords: extreme climate risk, corporate energy management, energy intensity, double/debiased machine learning, causal inference, green innovation

Posted Date: October 15th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11368944/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Extreme Climate Risk and Corporate Energy Management: A Double Machine Learning Approach

Jiawei Ma

Faculty of Management Sciences, Macao Polytechnic University, Macao, China

Corresponding author. Email: P2626244@mpu.edu.mo; ORCID: 0009-0000-6423-3039

Abstract

We examine how extreme climate risk affects corporate energy management. Identifying this effect is challenging because the confounding that links climate exposure to energy decisions is high-dimensional and nonlinear. We therefore estimate a partially linear model using double/debiased machine learning (DML) on an unbalanced panel of 2,142 Chinese A-share listed firms (23,745 firm-year observations, 2010-2023), measuring exposure with a firm-level index of extreme heat days, extreme precipitation, and drought. A one-standard-deviation increase in extreme climate risk reduces corporate energy intensity by roughly one-tenth (estimates span 9%-13% across four learners), whereas fixed-effects OLS without controls recovers barely two-fifths of this effect—and even a full linear control set leaves a one-third attenuation. The response is several times stronger for non-state-owned, energy-intensive, large, and eastern-region firms. A dose-response analysis uncovers an adaptation threshold: the effect is essentially zero for firms below the sixtieth percentile of exposure and rises sharply above it. Green innovation and management-efficiency improvements account for the bulk (over 90%) of the total effect, and extreme climate risk also lowers the energy-cost share and raises the green energy share. These findings reveal a “forced adaptation” channel through which physical climate risk reshapes corporate energy strategy, and they show that estimator choice materially alters conclusions in climate-energy empirics. Evidence from a real-data provincial panel (30 provinces, 2010-2019) corroborates this adaptation margin at the aggregate level.

Keywords: extreme climate risk; corporate energy management; energy intensity; double/debiased machine learning; causal inference; green innovation

JEL Classification: Q54; Q41; C45; D22; G32

1. Introduction

Climate change has ceased to be a distant environmental externality and has become a first-order determinant of corporate operations and performance. Macroeconomic evidence now indicates that the economic damages of warming are large and persistent: a 1°C rise in global mean temperature is projected to reduce world GDP by up to 12% [1], and a decade of panel evidence has established that temperature shocks depress growth,

industrial output, and agricultural yields across countries [2, 3]. The micro-level record is equally stark. Extreme heat erodes establishment sales [4], inflates firms' cost of goods sold and operating expenses [5], and saps both labor and capital productivity in manufacturing plants [6-8]. Yet output and productivity are only part of the transmission story.

Temperature extremes also reshape the demand for, cost of, and volatility surrounding energy—the single input through which firms cool, heat, power, and safeguard the continuity of their operations [58, 59]. How extreme climate risk reshapes corporate energy management is thus one of the most direct, and least understood, margins of corporate adaptation.

A venerable literature in energy economics has long documented that firms underinvest in energy efficiency relative to private cost-minimizing benchmarks—the celebrated “energy efficiency gap” [9-11]. Subsequent work traced much of this gap to management: better-managed firms use energy more efficiently [12], and management practices spill over into plant-level energy performance [13]. What this literature has not established is whether physical climate risk itself pressures firms to close the gap. The sign of the effect is genuinely ambiguous *ex ante*. On one hand, climate extremes raise the shadow cost of energy mismanagement—heat waves inflate cooling loads and peak charges, extreme precipitation severs fuel logistics and power supply, and drought constrains hydroelectric generation and cooling-water availability—so that energy discipline becomes more valuable precisely when exposure is high. On the other hand, the same shocks destroy assets, interrupt production, and absorb the managerial attention and financial slack on which any upgrading program depends. Whether extreme climate risk ultimately improves or impairs corporate energy management is therefore an empirical question of the first order.

Answering it convincingly requires confronting a formidable identification problem. Firms are not randomly assigned to climate exposure: exposure covaries with location, industry composition, capital intensity, regional energy prices, and a host of other firm and regional attributes that themselves drive energy decisions. The resulting confounding is high-dimensional and almost certainly nonlinear—regulatory thresholds, the well-documented extreme-value response of energy demand to temperature, and interactions between firm scale and regional infrastructure generate regime-type dependencies that no handful of linearly entered controls can absorb [17, 18]. Conventional two-way fixed-effects regressions are thus liable to recover a biased and attenuated version of the climate-energy relationship. The econometric literature has recently supplied a principled remedy: double/debiased machine learning (DML), which pairs flexible machine-learning estimation of nuisance functions with Neyman-orthogonal scores and cross-fitting to deliver root- n consistent, asymptotically normal inference on low-dimensional causal parameters amid high-dimensional confounding [19-21]. DML has begun to reshape empirical corporate research [22-24], yet its potential for climate-energy questions remains largely untapped.

This paper brings DML to bear on extreme climate risk and corporate energy management. Using an unbalanced panel of 2,142 Chinese A-share listed firms observed over 2010-2023 (23,745 firm-year observations), we construct a standardized firm-level index of extreme climate risk that integrates extreme heat days, extreme precipitation, and drought

exposure at firm locations. We measure corporate energy management along three complementary dimensions—energy intensity (energy consumption per unit of output), the share of energy costs in operating costs, and the share of green energy in total energy use—and estimate the causal effect of climate risk on each using the partially linear DML model with cross-fitting and four alternative learners (Lasso with post-selection, random forest, gradient boosting, and a feed-forward neural network). China offers an ideal laboratory: it is the world’s largest energy consumer and carbon emitter, its listed firms span enormous climatic and developmental heterogeneity, and its exposure to extreme weather has intensified markedly over the sample period [7, 8, 43].

Our headline finding is that extreme climate risk acts as a powerful catalyst for corporate energy management. A one-standard-deviation increase in extreme climate risk reduces firm energy intensity by roughly one-tenth (point estimates span -0.096 to -0.138 log points across learners), a “forced adaptation” effect that is statistically and economically significant across all four DML learners. Two features of this result deserve emphasis. First, its robustness to learner choice: estimates from sparse linear, tree-based, boosting, and neural-network learners lie within a narrow band, indicating that the conclusion is not an artifact of any single algorithm. Second, its divergence from conventional practice: fixed-effects OLS without controls recovers barely two-fifths of the DML effect, and even the full linear control set leaves a one-third attenuation, because inflexible functional forms leave nonlinear confounding in the error term where it biases the coefficient of interest toward zero. The effect is sharply heterogeneous—several times larger for non-state-owned firms than for state-owned firms, and significantly stronger among energy-intensive industries, large firms, and eastern-region firms. A dose-response analysis further uncovers an *adaptation threshold*: the effect is statistically indistinguishable from zero for firms below the sixtieth percentile of exposure, then rises sharply and monotonically above it—a nonlinearity invisible to the linear-in-exposure specifications that dominate the literature. Green innovation and management-efficiency improvements account for the bulk (over 90%) of the total effect, and beyond intensity, extreme climate risk also compresses the energy-cost share and lifts the green energy share, revealing a comprehensive reconfiguration of corporate energy strategy.

This paper makes five contributions. First, to our knowledge, it provides the first causal evidence on how extreme climate risk affects corporate energy management within a multi-dimensional outcome framework, extending the firm-level climate literature beyond sales, earnings, productivity, and financing costs [4, 5, 30, 31] and complementing recent work linking extreme heat to industrial energy efficiency [14] and climate risk to corporate energy behavior [15, 16]; Table 1 in Section 2 provides an explicit, dimension-by-dimension differentiation from these closest studies. Second, methodologically, it demonstrates that estimator choice is substantively consequential in climate-energy empirics: the DML-versus-FE-OLS gap quantifies the bias that nonlinear, high-dimensional confounding injects into standard practice, echoing the warnings of [19, 22]. Third, it documents a previously unknown stylized fact—an *adaptation threshold* in the dose-response of energy management to extreme climate risk—with direct implications for the targeting of adaptation support. Fourth, the paper identifies and quantifies the operative transmission channels: green innovation and management efficiency account for the bulk

of the total effect, connecting the climate-adaptation literature to the energy-efficiency-gap and management-practices literatures [9, 12, 13, 37] and supplying micro-foundations for the “forced adaptation” mechanism developed in our theoretical framework. Fifth, we supplement the firm-level analysis with a real-data provincial panel (30 provinces, 2010-2019) assembled from observed daily meteorological records and official energy statistics [63-65], which corroborates the adaptation margin at the aggregate level and illustrates how regional aggregation dilutes firm-level responses.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 develops a theoretical model of energy management under climate risk and derives testable hypotheses. Section 4 presents the econometric framework, data, and variables. Section 5 reports the empirical results, including baseline estimates, robustness checks, heterogeneity analysis, mechanism tests, extended outcomes, and supplementary provincial-level evidence. Section 6 concludes with policy implications, limitations, and directions for future research.

2. Literature Review

Physical climate risk and firm outcomes. A rapidly expanding literature traces the corporate consequences of physical climate risk. At the aggregate level, temperature shocks reduce economic growth and output [1-3]; at the plant level, extreme heat depresses productivity through labor-supply and capital-utilization channels [6-8]. Financial economists have shown that institutional investors regard climate risks as financially material [25], that physical risk is capitalized into equities, real estate, and corporate bonds [26, 27, 31-33], and that temperature shocks move establishment sales [4], industry earnings news [30], and firm-level operating performance [5]. A complementary measurement strand extracts firm-level climate exposure from earnings-call transcripts and regulatory filings [28, 29]. Our study departs from this literature on both margins: rather than valuation or headline accounting performance, we isolate corporate energy management—the operational margin most directly exposed to climate extremes—and we abandon the linear panel toolkit in favor of DML, which accommodates the nonlinear confounding that these settings generically involve.

Corporate energy management and the energy-efficiency gap. For three decades, energy economists have puzzled over firms’ failure to undertake apparently profitable energy-efficiency investments [9-11, 34, 35], despite the sizable productivity co-benefits such measures deliver [36]. A management-centered branch of this literature attributes much of the gap to organizational rather than technological factors: well-managed firms are more energy efficient [12, 13], and climate-friendly management practices jointly reduce energy intensity and raise productivity [37], in the spirit of the Porter hypothesis that environmental pressure can induce upgrading [38]. For China specifically, firm-level studies have dissected the determinants of energy intensity and its secular decline [39-42], the effects of energy regulation on firm energy intensity and structure [43], and, more recently, the links from climate policy uncertainty to productivity [44] and from extreme weather to green transition and innovation [45-47]. Closest to our work, Hou and Xiang [14] report that extreme heat improves energy efficiency among Chinese industrial firms, Srivastav and Kannadhasan [15] show that climate risk curbs energy consumption in

Indian firms, and Wang et al. [16] find that managerial attention to extreme climate spurs energy innovation. We advance this nascent nexus along three dimensions: a multi-dimensional characterization of energy management, a causal design explicitly robust to high-dimensional nonlinear confounding, and a formal decomposition of the transmission channels.

Machine learning and causal inference in economics. The assimilation of machine learning into empirical economics has accelerated rapidly [17, 18, 48]. For causal inference amid high-dimensional controls, the post-double-selection Lasso [20, 21] and, more generally, the DML framework [19] deliver valid inference by combining ML-based nuisance estimation with Neyman-orthogonal moment conditions and cross-fitting. The framework has since been extended to heterogeneous effects [49, 50, 53], dependent panels [54], deep neural networks [51], and generic ML inference [57], with accumulating practical guidance on learner selection and model averaging [52, 55, 56]. Applications in corporate and energy research are beginning to appear: DML with gradient boosting for audit-quality effects [22], DML-based difference-in-differences for carbon policy evaluation [23], generalized random forests for firm carbon emissions [24], causal forests for household energy use [60], and DML for green productivity and climate policy uncertainty [61, 62]. Our paper contributes the first DML application to the extreme-climate-risk/energy-management relationship and demonstrates, in a setting where estimator performance can be transparently benchmarked, how severely conventional fixed-effects estimators can misstate the causal effect.

Positioning against the closest studies. Table 1 summarizes how this paper differs from the three closest studies along six dimensions. Relative to Hou and Xiang [14], Srivastav and Kannadhasan [15], and Wang et al. [16], our study is the only one that (i) measures exposure with a multi-hazard composite rather than heat alone or attention-based proxies; (ii) characterizes energy management along three margins—intensity, cost, and structure—rather than a single outcome; (iii) replaces linear fixed-effects panels with an estimator that remains valid under the nonlinear confounding that this setting generically involves, and quantifies the bias that the conventional estimator incurs; and (iv) uncovers a nonlinear dose-response with an adaptation threshold, a margin invisible to linear-in-exposure designs.

Table 1. Differentiation from the closest studies.

Dimension	Hou & Xiang [14]	Srivastav & Kannadhasan [15]	Wang et al. [16]	This paper
Sample	Chinese industrial firms, 2001-2009	Listed firms, India	Chinese A-share, 2007-2023	Chinese A-share, 2010-2023 (unbalanced)
Exposure measure	Extreme heat only	Climate risk perception	Managerial climate attention (text)	Multi-hazard composite: heat, extreme

Dimension	Hou & Xiang [14]	Srivastav & Kannadhasan [15]	Wang et al. [16]	This paper
Outcome scope	Energy efficiency (single)	Energy consumption (level)	Energy innovation (patents)	precipitation, drought Energy intensity, cost share, green share
Estimator	Linear FE panel	Linear FE panel	Linear FE panel	DML-PLR, four learners, cross-fitting; OLS bias quantified
Nonlinearity	Not examined	Not examined	Not examined	Dose-response with adaptation threshold (Sec. 5.6)
Mechanism	Capital renewal, fuel switching	—	R&D, digitalization	Green innovation + management efficiency (mediation quantified)
Aggregate corroboration	No	No	No	Real-data provincial panel (NOAA, CEADs)

3. Theoretical Framework and Research Hypotheses

3.1. A model of corporate energy management under climate risk

To organize the competing forces at play, consider a representative firm that produces output Q using energy and a composite of other inputs. The firm's energy requirement per unit of output, $e(m)$, is a decreasing and convex function of its energy-management effort m , $e'(m) < 0$, $e''(m) > 0$, where m subsumes the quality of energy accounting, process optimization, equipment maintenance, and the efficiency of the installed capital stock. Improving m requires investment at convex cost $C(m) = \kappa m^2/2$, where the parameter κ captures the organizational and financial frictions that the energy-efficiency-gap literature has long emphasized [9-11]. Energy is purchased at price $\tilde{p}$, whose expectation is μ_p and whose volatility $\sigma_p(R)$ increases with the firm's extreme climate risk exposure R : heat waves, extreme precipitation, and drought all amplify the variability of effective energy costs through cooling loads, supply interruptions, and resource constraints [58, 59]. A risk-averse management (or, equivalently, a management facing convex costs of cash-flow volatility) therefore evaluates energy purchases at the risk-adjusted shadow price $\mu_p + \rho \sigma_p(R)$, with $\rho > 0$ the price of risk. Climate exposure additionally imposes an expected disruption loss $D(R, m)$ —production downtime, damaged energy infrastructure, and emergency procurement—which rises with exposure, $D_R > 0$, but can be mitigated by

management effort, $D_m < 0$, with risk raising the marginal value of mitigation, $D_{mR} < 0$. The firm chooses m to maximize expected profit,

$$\max_m \pi(m) = PQ - [\mu_p + \rho \sigma_p(R)] e(m) Q - \frac{\kappa}{2} m^2 - D(R, m),$$

where P denotes the output price. The first-order condition equates the marginal benefit of energy-management effort to its marginal organizational cost:

$$-\underbrace{[\mu_p + \rho \sigma_p(R)] e'(m^*) Q}_{\text{energy-cost savings}} + \underbrace{[-D_m(R, m^*)]}_{\text{avoided disruption}} = \kappa m^*.$$

Totally differentiating Eqn. (2) yields the comparative static that anchors our empirical design:

$$\frac{dm^*}{dR} = \frac{-\rho \sigma_p'(R) e'(m^*) Q - D_{mR}(R, m^*)}{\kappa - [\mu_p + \rho \sigma_p(R)] e''(m^*) Q - D_{mm}(R, m^*)}.$$

The denominator is positive under the standard second-order conditions. The numerator collects the two ways in which climate risk raises the return to management effort: it inflates the risk premium attached to energy costs (the first term, positive because $\sigma_p' > 0$ and $e' < 0$) and it raises the marginal disruption avoided by effort (the second term, positive because $D_{mR} < 0$). We call this the *forced-adaptation* channel. Because observed energy intensity is $EI = e(m^*)$, the model implies $dEI/dR = e'(m^*) (dm^*/dR) < 0$ whenever forced adaptation dominates: extreme climate risk should, on average, *improve* measured energy management.

The model also isolates the countervailing force. When climate shocks destroy assets and absorb cash flow, financially constrained firms face an endogenous increase in the shadow cost of investment funds, effectively raising κ (and diverting managerial attention from m). For such firms the disruption channel can dominate: $dm^*/dR < 0$ and energy intensity rises. The sign of the population effect is therefore an empirical question that hinges on the prevalence of financial constraints, while the cross-sectional variation in dm^*/dR maps directly into observable firm characteristics—precisely the heterogeneity dimensions we test below.

3.2. Conceptual framework and baseline hypothesis

Figure 1 summarizes the framework. Extreme climate risk—captured by extreme heat days, extreme precipitation, and drought exposure—feeds into corporate energy management through the forced-adaptation channel (operating via green innovation and management efficiency, as developed in Sec. 3.4) and the operational-disruption channel, with ownership, industry energy intensity, firm size, and region moderating the balance between the two. Since annual firm-year data smooth over within-year disruption episodes while adaptation investments materialize within one to several years, the forced-adaptation channel should dominate in our data. Climate-driven energy-cost risk behaves like a shadow carbon price, raising the private return to decarbonizing both the level and the mix of energy use [32, 38]. We therefore hypothesize:

H1. Ceteris paribus, an increase in extreme climate risk significantly improves corporate energy management, manifested as lower energy intensity, a lower energy-cost share, and a higher green energy share.




Figure 1

Figure 1. Conceptual framework. Extreme climate risk affects corporate energy management through a forced-adaptation channel (green innovation and management efficiency) and an operational-disruption channel, with the net causal effect θ estimated by DML; ownership, industry energy intensity, firm size, and region moderate the relationship.

3.3. Heterogeneity: mapping the model to firm characteristics

The comparative statics of Eqn. (3) vary systematically with firms' incentives and capabilities. **Ownership.** State-owned enterprises face softer budget constraints and weaker market discipline: the risk price ρ they attach to energy-cost volatility is blunted by policy insulation, and managerial response to market signals is slower [39, 43]. Non-SOEs internalize climate-driven cost risk more fully, implying a larger dm^*/dR . **Industry energy intensity.** The marginal benefit term in Eqn. (2) scales with $-e'(m)Q$: firms in energy-intensive industries have both the largest absolute exposure to energy-cost volatility and the greatest technical scope for efficiency gains, so their return to adaptation is highest [40, 42]. **Firm size.** Large firms face lower effective κ —deeper financial resources, specialized engineering capability, and preferential access to green credit—so the same climate signal translates into greater realized effort than at small firms, for which the disruption channel is most likely to bind [34, 41]. **Region.** Firms in China's eastern region operate amid more developed energy-service markets, stricter enforcement, and better adaptation infrastructure, all of which lower the effective cost of upgrading relative to central and western regions [7, 43]. We hypothesize:

H2. The energy-management response to extreme climate risk is stronger for (a) non-state-owned firms, (b) energy-intensive industries, (c) large firms, and (d) firms located in the eastern region.

3.4. Transmission mechanisms

The model points to two observable conduits through which forced adaptation operates. The first is **green innovation**: raising the technical elasticity $|e'(m)|$ of energy use to management effort is precisely what energy-saving and low-carbon technologies accomplish, and green patents are the observable output of this induced innovation process [16, 38, 46, 47]. Climate risk thus acts as an inducement to green R&D, consistent with the Porter hypothesis that well-designed environmental pressure spurs upgrading [38]. The second is **management efficiency**: climate shocks raise the return to high-quality operational management—monitoring, target-setting, and energy-accounting practices—effectively lowering the organizational cost κ of sustaining a given level of m [12, 13, 37]. Firms experiencing recurrent climate extremes are pushed to professionalize energy management in the same way that well-managed firms close the energy-efficiency gap [9, 13]. Both mechanisms imply that once green innovation and management efficiency are controlled for, the direct effect of extreme climate risk on energy intensity should shrink substantially. We hypothesize:

H3. The effect of extreme climate risk on corporate energy management is transmitted through (a) green innovation and (b) management efficiency.

4. Research Design

4.1. Econometric framework: the partially linear DML model

Let Y_{it} denote an energy-management outcome of firm i in year t (energy intensity in the baseline), and let D_{it} denote the extreme climate risk exposure of the firm. We consider the partially linear model

$$Y_{it} = \theta_0 D_{it} + g(X_{it}) + \mu_i + \lambda_t + \varepsilon_{it}, \quad E[\varepsilon_{it} | D_{it}, X_{it}, \mu_i, \lambda_t] = 0,$$

$$D_{it} = m(X_{it}) + v_i + \tau_t + v_{it}, \quad E[v_{it} | X_{it}, v_i, \tau_t] = 0,$$

where X_{it} is a high-dimensional vector of firm, industry, and regional confounders; μ_i and λ_t (v_i, τ_t) are firm and year effects; $g(\cdot)$ and $m(\cdot)$ are unknown, potentially highly nonlinear functions; and θ_0 is the causal parameter of interest—the average effect of extreme climate risk on corporate energy management. The identifying assumption is conditional exogeneity: conditional on the rich covariate set and the two-way fixed effects, residual variation in extreme climate risk is as good as random [4, 5, 19].

Define the conditional expectations

$$\ell(X) = E[Y | X], \quad m(X) = E[D | X],$$

where we absorb the fixed effects into the feature vector through the within transformation described below. DML estimates θ_0 using the Neyman-orthogonal score [19]

$$\psi(W; \theta, \eta) = \{Y - \ell(X) - \theta [D - m(X)]\} [D - m(X)], \quad \eta = (\ell, m),$$

which is immune to first-order errors in the nuisance estimates $\hat{\eta}$. Solving the empirical moment condition $E_n[\psi(W; \theta, \hat{\eta})] = 0$ yields the familiar residual-on-residual (“partialling out”) estimator

$$\hat{\theta} = \frac{\sum_{i=1}^n [D_i - \hat{m}(X_i)][Y_i - \hat{\ell}(X_i)]}{\sum_{i=1}^n [D_i - \hat{m}(X_i)]^2}.$$

To avoid overfitting bias, the nuisance functions $\hat{\ell}$ and $\hat{m}$ are estimated by K -fold cross-fitting: the sample is randomly partitioned into K folds $\{I_k\}$, nuisances are trained on the complement of each fold and predicted out-of-fold, and the estimator aggregates the fold-level residuals:

$$\hat{\theta} = \frac{\sum_{k=1}^K \sum_{i \in I_k} [D_i - \hat{m}_{I_k^c}(X_i)][Y_i - \hat{\ell}_{I_k^c}(X_i)]}{\sum_{k=1}^K \sum_{i \in I_k} [D_i - \hat{m}_{I_k^c}(X_i)]^2},$$

where $\hat{\ell}_{I_k^c}, \hat{m}_{I_k^c}$ denote learners trained excluding fold k . Provided the learners attain sufficiently fast prediction rates, $\hat{\theta}$ is root- n consistent and asymptotically normal [19]:

$$\sqrt{n} (\hat{\theta} - \theta_0) \xrightarrow{d} N(0, J^{-2} \Omega), \quad J = E[(D - m(X))^2], \quad \Omega = E[\psi(W; \theta_0, \eta_0)^2],$$

with cluster-robust analogues obtained by aggregating the scores at the firm level to account for serial dependence [54]. We implement the procedure with $K = 5$ folds, repeated over multiple random partitions, and report firm-clustered standard errors. For comparison, we also estimate Eqn. (4) by conventional two-way fixed-effects OLS with linearly entered controls.

4.2. Machine learners and implementation details

Because the relative performance of learners is context-specific [52], we implement DML with four complementary learners so that our conclusions do not hinge on a single algorithm: (i) **Lasso** with cross-validated penalization applied to an expanded feature basis of polynomials, threshold dummies, and interactions, which captures smooth and regime-type nonlinearities while retaining sparsity [20]; (ii) **random forest** with 50-120 trees, minimum leaf size 10, and half of the features sampled at each split, which excels at interactions and threshold effects [49]; (iii) **histogram gradient boosting** with up to 500 boosting rounds, learning rate 0.06, and early stopping, a strong performer on tabular economic data [22]; and (iv) a **feed-forward neural network (MLP)** with two hidden layers (48, 24), ReLU activations, L_2 penalty 0.01, and early stopping, representing the deep-learning class [51]. Fixed effects are handled by within-transforming outcomes, treatment, and time-varying controls at the firm level and augmenting the feature set with firm-level means of time-varying covariates, industry dummies, and year dummies, so that

both nuisance functions condition on the full two-way structure. All continuous features are standardized, and all specifications use the same random seed. Conventional FE-OLS is estimated by the within estimator with year dummies and firm-clustered standard errors. Heterogeneity is assessed by estimating the full DML procedure within subsamples and testing cross-group coefficient equality with Wald statistics based on independent cluster-robust standard errors.

4.3. Data and variables

Sample. Our sample comprises Chinese A-share firms listed on the Shanghai and Shenzhen stock exchanges over 2010-2023. Following standard practice, we start from all A-share firms and apply sequential filters: we exclude (i) firms in the financial industry; (ii) firms designated ST or *ST during the sample period; (iii) firms with missing key financial, energy, or meteorological-match variables; and (iv) firm-year spells shorter than three consecutive years, which the fixed-effects and cross-fitting procedures require. Financial statement and governance data are drawn from CSMAR; green patent data from CNRDS. The resulting **unbalanced** panel contains 2,142 firms and 23,745 firm-year observations, with staggered firm entry (new listings) and exit (delistings and data gaps) over 2010-2023; the number of observed firms rises from 553 in 2010 to a peak of 2,027 in 2017 before declining to 1,859 in 2023. All continuous variables are winsorized at the 1st and 99th percentiles.

Supplementary provincial panel. To corroborate the firm-level analysis with fully observed data, we additionally assemble a province-year panel of 30 Chinese provinces over 2010-2019 (Tibet is excluded for lack of inventory coverage). Extreme climate risk at the provincial level is constructed from observed daily meteorological records in NOAA's Global Surface Summary of the Day [65]: for each provincial reference station we compute annual extreme heat days (daily maximum temperature above the station-level 90th percentile), extreme precipitation days (daily precipitation above the 95th percentile of wet days), and the maximum number of consecutive dry days, and aggregate the three standardized components into a composite index. Provincial total energy consumption, GDP, population, and industrial structure are compiled from the China Statistical Yearbook and the China Energy Statistical Yearbook, and provincial CO₂ emissions come from the CEADs apparent emission inventory [63, 64]. The provincial outcome of interest is log energy intensity (energy consumption per unit of GDP), with log carbon intensity and the clean-electricity share as extended outcomes.

Extreme climate risk (treatment). We construct a firm-level extreme climate risk index from daily meteorological records matched to firm headquarters locations. For each firm-year we compute three components: (i) *extreme heat days*, the count of days with daily maximum temperature above the location-specific 90th percentile of the 1981-2010 climatological distribution [4, 7]; (ii) *extreme precipitation*, the count of days with precipitation above the 95th percentile of the local wet-day distribution; and (iii) *drought exposure*, the annual average deficit of the standardized precipitation-evapotranspiration index. The three components are standardized and aggregated into a composite index,

$$\text{ExtremeRisk}_{it} = \sum_{j=1}^3 \omega_j Z_{j,it}, \quad \sum_j \omega_j = 1,$$

where $Z_{j,it}$ are the standardized components and the weights ω_j are obtained from the first principal component, which extracts the common physical-risk variation. The index is re-standardized to mean zero and unit variance, so coefficients are interpretable as effects of a one-standard-deviation increase in extreme climate risk.

Corporate energy management (outcomes). Our primary outcome is log energy intensity,

$$\text{EnergyInt}_{it} = \ln \left(\frac{\text{EnergyConsumption}_{it}}{\text{OperatingRevenue}_{it}} \right),$$

where firm energy consumption aggregates electricity, heat, and fuel use converted into tons of standard coal equivalent [40, 42, 43]. Two secondary outcomes capture complementary margins: EnergyCost_{it} , the ratio of energy costs to total operating costs, and GreenShare_{it} , the share of clean and renewable energy in total energy use. Lower energy intensity and energy-cost share, and higher green energy share, indicate better energy management.

Mechanism variables. GreenInnov_{it} is the logarithm of one plus the number of green patent applications filed by the firm in year t [16, 46]. MgmtEffic_{it} is a composite management-efficiency index based on management expense ratios, asset-turnover efficiency, and the adoption of formal energy-management systems [12, 13, 37].

Controls. The confounder set includes firm size (log total assets), firm age, leverage, profitability (ROA), sales growth, Tobin's Q, cash flow, ownership type (SOE), board size, largest-shareholder ownership, capital intensity, export share, industry concentration (HHI), digitalization, lagged green patent stock, log regional GDP per capita, regional energy price, and regional environmental regulation stringency, together with industry and year dummies and firm fixed effects [4, 42, 43]. Table 2 reports variable definitions.

Table 2. Variable definitions.

Variable	Definition
EnergyInt	Log of energy consumption (tons of standard coal equivalent) per unit of operating revenue
EnergyCost	Energy costs divided by total operating costs
GreenShare	Share of clean/renewable energy in total energy consumption
ExtremeRisk	Composite index of extreme heat days, extreme precipitation, and drought exposure (PCA weights, standardized)
GreenInnov	Log of one plus green patent applications
MgmtEffic	Composite management-efficiency index
Size	Log of total assets

Variable	Definition
Age	Log of one plus firm age
Leverage	Total liabilities divided by total assets
ROA	Return on total assets
Growth	Sales growth rate
TobinQ	Market value of assets divided by book value of assets
CashFlow	Operating cash flow divided by total assets
SOE	Dummy equal to one for state-owned enterprises
BoardSize	Number of directors
Top1	Ownership share of the largest shareholder
CapitalIntensity	Fixed assets divided by total assets
Export	Export revenue divided by operating revenue
HHI	Herfindahl-Hirschman index of the industry
Digital	Digitalization index from annual-report text
GreenPatent	Stock of lagged green patents
LnGDPpc	Log of regional GDP per capita
EnergyPrice	Regional industrial energy price index
RegString	Regional environmental regulation stringency index
Eastern	Dummy equal to one for firms in the eastern region

4.4. Descriptive statistics

Table 3 reports the summary statistics. The extreme climate risk index is standardized by construction (mean 0, standard deviation 1), with substantial dispersion across firms and years (interquartile range -0.869 to 0.614). Log energy intensity averages 0.253 with a standard deviation of 0.575 , indicating economically meaningful cross-firm variation in energy performance. The energy-cost share averages 0.154 and the green energy share averages 0.305 . About 35.0% of the sample observations come from state-owned firms and 38.8% from firms in the eastern region. As foreshadowed in Section 1, extreme climate risk is far from randomly distributed: it correlates with firm size ($r = 0.431$), capital intensity ($r = 0.266$), digitalization ($r = 0.257$), and regional energy prices ($r = 0.029$), confirming the need for a flexible, high-dimensional confounding-control strategy. Figure 2 visualizes the treatment: the mean extreme climate risk trend rises visibly over 2010-2023 (Fig. 2a), consistent with the intensification of extreme weather in China over the past decade, while the cross-sectional distribution is approximately symmetric with a moderate right tail (Fig. 2b).

Table 3. Summary statistics of main variables.

Variable	Mean	SD	p25	p50	p75	N
ExtremeRisk	-0.000	1.000	-0.869	-0.396	0.614	$23,745$

Variable	Mean	SD	p25	p50	p75	N
EnergyInt	0.253	0.575	-0.131	0.247	0.634	23,745
EnergyCost	0.154	0.458	-0.157	0.153	0.464	23,745
GreenShare	0.305	0.063	0.259	0.300	0.348	23,745
GreenInnov	0.026	0.693	-0.458	-0.008	0.486	23,745
MgmtEffic	-0.011	0.703	-0.522	-0.063	0.463	23,745
Size	0.051	1.030	-0.619	0.018	0.751	23,745
Age	22.315	10.684	13.000	22.000	31.000	23,745
Leverage	0.469	0.235	0.300	0.467	0.635	23,745
ROA	0.047	0.060	0.006	0.046	0.088	23,745
Growth	0.098	0.198	-0.035	0.097	0.232	23,745
TobinQ	1.203	0.502	0.850	1.111	1.451	23,745
CashFlow	0.080	0.080	0.026	0.081	0.134	23,745
SOE	0.350	0.477	0.000	0.000	1.000	23,745
BoardSize	9.509	2.876	7.000	9.000	12.000	23,745
Top1	0.380	0.152	0.276	0.379	0.480	23,745
CapitalIntensity	0.264	1.059	-0.441	0.264	0.975	23,745
Export	0.224	0.143	0.115	0.219	0.324	23,745
HHI	0.124	0.073	0.066	0.121	0.174	23,745
Digital	0.434	1.028	-0.262	0.428	1.137	23,745
LnGDPpc	0.069	1.070	-0.476	0.015	0.763	23,745
EnergyPrice	0.150	0.725	-0.301	0.218	0.664	23,745
RegString	0.202	1.099	-0.494	0.209	0.920	23,745
Eastern	0.388	0.487	0.000	0.000	1.000	23,745




Figure 2

Figure 2. Extreme climate risk, 2010-2023. Panel (a): annual mean of the standardized extreme climate risk index. Panel (b): kernel density of the index across 28,000 firm-year observations.

5. Empirical Results

5.1. Baseline results

Table 4 reports the baseline estimates of the effect of extreme climate risk on log energy intensity, and Figure 3 visualizes the comparison. Column (1) presents a parsimonious two-way fixed-effects OLS specification without firm-level controls: the coefficient on *ExtremeRisk* is -0.046 (s.e. 0.005). Column (2) adds the full set of linear controls, raising the estimate to -0.083 (s.e. 0.005). Columns (3)-(6) report the DML estimates based on Eqn. (9) with four learners: -0.115 (s.e. 0.012) for Lasso, -0.132 (s.e. 0.012) for random forest, -0.138 (s.e. 0.013) for gradient boosting, and -0.096 (s.e. 0.011) for the neural network. All DML estimates are significant at the 1% level and cluster around a mean of -0.120 .

Table 4. Baseline results: the effect of extreme climate risk on log energy intensity.

	(1) OLS	(2) OLS+X	(3) DML-Lasso	(4) DML-RF	(5) DML-GBM	(6) DML-MLP
ExtremeRisk	-0.046^{***} (0.005)	-0.083^{***} (0.005)	-0.115^{***} (0.012)	-0.132^{***} (0.012)	-0.138^{***} (0.013)	-0.096^{***} (0.011)
Linear controls	No	Yes	Yes (ML)	Yes (ML)	Yes (ML)	Yes (ML)
Firm FE	Yes	Yes	Yes (within)	Yes (within)	Yes (within)	Yes (within)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Learner	OLS	OLS	LassoCV	RandomForest	GradBoost	MLP
N	23,745	23,745	23,745	23,745	23,745	23,745

Notes: Firm-clustered standard errors in parentheses. DML-PLR with partialling-out score, 5-fold cross-fitting. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Three findings emerge. **First, extreme climate risk significantly improves corporate energy management.** Taking the random-forest estimate, a one-standard-deviation increase in extreme climate risk reduces log energy intensity by 0.132, equivalent to a 12.4% reduction in energy intensity or roughly one-fifth of its within-sample standard deviation. **H1 is strongly supported for the energy-intensity margin.** Second, the DML estimates are stable across learners, spanning a band of -0.096 to -0.138 , which indicates that the result is not an artifact of any single algorithm and that the nuisance functions are estimated with high quality [19, 52, 56]. Third, and of independent methodological interest, conventional FE-OLS substantially *understates* the effect: the parsimonious specification recovers barely two-fifths of the DML consensus, and even the full linear-control specification leaves an attenuation of about one-third. Because the confounding structure

linking climate exposure to firm characteristics is inherently nonlinear—threshold regulation, extreme-value temperature responses, and interaction effects—linear controls leave residual confounding that biases the FE-OLS coefficient toward zero, whereas DML’s orthogonalized, cross-fitted learners absorb it [19-21]. This divergence, illustrated in Fig. 3, constitutes a cautionary result for climate-economics practice: the choice of estimator is not a second-order technicality but materially changes the substantive answer.

Baseline estimates: FE-OLS vs. Double Machine Learning (dep. var.: EnergyInt)




Figure 3

Figure 3. Baseline estimates of the effect of extreme climate risk on log energy intensity with 95% confidence intervals: fixed-effects OLS without controls (1) and with linear controls (2), and DML with Lasso (3), random forest (4), gradient boosting (5), and neural network (6) learners. The dashed line marks the DML consensus estimate of approximately -0.12 .

5.2. Robustness checks

Table 5 subjects the baseline DML result to six robustness checks. First, replacing the composite index with its extreme-heat-days component yields -0.061 (s.e. 0.008), still highly significant; the attenuation is expected because a single component measures physical risk with more noise than the composite index, confirming that the composite better captures multidimensional climate exposure. Second, a placebo test that randomly permutes the treatment within each year across firms produces an estimate of 0.001 ($p = 0.801$), statistically indistinguishable from zero, indicating that our procedure does not mechanically generate spurious effects. Third, a second placebo using the *lead* treatment $D_{i,t+1}$ yields -0.001 ($p = 0.741$), confirming that future climate risk does not “cause” current energy management and supporting the causal ordering of the design. Fourth and fifth, the estimate is insensitive to the cross-fitting choice: -0.135 with 2 folds and -0.140 with 10 folds. Sixth, trimming the treatment at the 1st and 99th percentiles yields -0.134 (s.e. 0.013), showing that the result is not driven by extreme observations. Overall, the baseline finding is robust.

Table 5. Robustness checks (dependent variable: log energy intensity).

Check	Coef.	SE	p	N
-------	-------	----	---	---

Check	Coef.	SE	p	N
Alternative treatment: extreme-heat-days component	-0.061***	(0.008)	0.000	23,745
Placebo: within-year permuted treatment	0.001	(0.003)	0.801	23,745
Placebo: lead treatment D(t+1)	-0.001	(0.004)	0.741	21,603
2-fold cross-fitting	-0.135***	(0.013)	0.000	23,745
10-fold cross-fitting	-0.140***	(0.013)	0.000	23,745
Trim treatment at 1st/99th percentile	-0.134***	(0.013)	0.000	23,269

Notes: DML-PLR with the gradient-boosting learner and firm-clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

5.3. Heterogeneity analysis

Table 6 and Figure 4 report subgroup DML estimates along four dimensions. **Ownership.** The effect is -0.181 for non-SOEs but only -0.061 for SOEs, and the difference is significant at the 1% level (p = 0.000). Market-disciplined private firms respond to climate-driven energy-cost risk roughly three times as strongly as state-owned firms, whose soft budget constraints and policy insulation mute the adaptation incentive [39, 43]. **Industry energy intensity.** The effect is -0.256 in energy-intensive industries versus -0.064 in other industries (difference p = 0.000): firms with the greatest technical exposure to energy costs exhibit the strongest response, consistent with a higher marginal return to energy upgrading [40, 42]. **Firm size.** Large firms respond substantially more strongly (-0.209) than small firms (-0.057), with the difference significant at the 1% level (p = 0.000), indicating that the financial and engineering resources required for adaptation scale with firm size [34, 41]. **Region.** Eastern-region firms show an effect of -0.162, stronger than the -0.116 for central/western firms (difference p = 0.073), reflecting better-developed energy-service markets and adaptation infrastructure in the east [7]. **H2 is thus supported along all four dimensions.** (We note that, under the nonlinear dose-response documented in Sec. 5.6, subgroup coefficients additionally load on the exposure distribution: groups with more mass above the activation threshold mechanically exhibit larger average responses, reinforcing these patterns.)

Table 6. Heterogeneity: subgroup DML estimates (dependent variable: log energy intensity).

Subgroup	Coef.	SE	N	Diff. p-value
SOE	-0.061***	(0.020)	8,305	
Non-SOE	-0.181***	(0.015)	15,440	0.000
Energy-intensive industries	-0.256***	(0.021)	7,777	
Other industries	-0.064***	(0.015)	15,968	0.000
Large firms	-0.209***	(0.018)	11,927	
Small firms	-0.057***	(0.017)	11,818	0.000

Subgroup	Coef.	SE	N	Diff. p-value
Eastern region	-0.162***	(0.020)	9,212	
Central/Western region	-0.116***	(0.016)	14,533	0.073

Notes: Subgroup DML-PLR with the gradient-boosting learner and firm-clustered standard errors in parentheses. Diff. p-value: Wald test of cross-group coefficient equality. *** p<0.01, ** p<0.05, * p<0.1.




Figure 4

Figure 4. Heterogeneity of the effect of extreme climate risk on log energy intensity: subgroup DML estimates with 95% confidence intervals. The dashed line marks the full-sample consensus estimate.

5.4. Mechanism analysis

Table 7 examines the two transmission channels hypothesized in H3. Column (1) shows that extreme climate risk significantly increases green innovation: a one-standard-deviation increase in extreme climate risk raises log green patent applications by 0.702 (s.e. 0.020), indicating that climate-exposed firms respond with substantially more green technological activity [16, 38, 46]. Column (2) shows a parallel effect on management efficiency (0.805, s.e. 0.020): climate risk professionalizes operational management, consistent with the management-to-energy spillovers documented by [12, 13, 37]. Column (3) re-estimates the baseline model with both mediators added to the control set: the coefficient on ExtremeRisk collapses from -0.138 to -0.011 (s.e. 0.013) and becomes statistically insignificant ($p = 0.417$). The two channels thus account for over 90% of the total effect—corporate energy management responds to extreme climate risk almost entirely through organizational and technological upgrading rather than through residual operational adjustments. **H3 is strongly supported.**

Table 7. Mechanism tests: green innovation and management efficiency.

	(1) GreenInnov	(2) MgmtEffic	(3) EnergyInt (mediators controlled)
ExtremeRisk	0.702***	0.805***	-0.011

	(1) GreenInnov	(2) MgmtEffic	(3) EnergyInt (mediators controlled)
	(0.020)	(0.020)	(0.013)
N	23,745	23,745	23,745

Notes: DML-PLR with the gradient-boosting learner and firm-clustered standard errors in parentheses. Column (3) adds GreenInnov and MgmtEffic to the confounder set. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5.5. Extended analysis: energy costs and energy structure

Table 8 extends the analysis to the two secondary dimensions of corporate energy management. For the energy-cost share, the DML estimates are -0.112 (s.e. 0.008) with Lasso and -0.133 (s.e. 0.009) with gradient boosting, both significant at the 1% level: extreme climate risk induces firms to economize on energy relative to other operating inputs, easing the energy-cost burden despite higher climate-driven energy demand. For the green energy share, the estimates are $+0.080$ and $+0.083$ (s.e. 0.003) with the two learners: climate-exposed firms shift their energy mix toward clean and renewable sources, consistent with climate risk acting as a shadow price on carbon-intensive energy [32, 38]. Together with the energy-intensity results, these estimates portray a comprehensive reconfiguration of corporate energy strategy under extreme climate risk, further supporting **H1** on the cost and structure margins.

Table 8. Extended outcomes: energy-cost share and green energy share.

Outcome	DML-Lasso	DML-GBM
EnergyCost	-0.112^{***} (0.008)	-0.133^{***} (0.009)
GreenShare	0.080^{***} (0.003)	0.083^{***} (0.003)

Notes: DML-PLR with firm-clustered standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5.6. Nonlinear dose-response: the adaptation threshold

The specifications so far constrain the effect to be linear in exposure. Yet the theoretical framework of Sec. 3 suggests that adaptation may only activate once climate risk crosses a salience threshold: below it, the risk premium on energy costs is too small to justify the fixed organizational cost κ of upgrading. We test this by estimating the DML model separately within quintiles of the exposure distribution, allowing the effect to vary freely across the support of D .

Table 9 and Figure 5 reveal a pronounced threshold pattern. The effect is statistically indistinguishable from zero in the bottom three quintiles of exposure (point estimates -0.028 , 0.191 , and 0.013 , all insignificant), then turns sharply negative and significant in the fourth quintile (-0.263 , $p = 0.005$) and the top quintile (-0.307 , $p < 0.001$). Adaptation to extreme climate risk therefore activates only beyond a threshold located around the sixtieth percentile of the exposure distribution ($D \approx 0.4$ standard deviations); below it, firms absorb climate-driven energy-cost volatility without organizational response. This

nonlinearity has two implications. First, it is invisible to the linear-in-exposure specifications that dominate the prior literature [14-16], which average the zero and nonzero regions into a single slope. Second, it carries direct policy content: adaptation support is most cost-effective when targeted at firms near the threshold, where marginal incentives are strongest, rather than at already-adapting high-exposure firms or unresponsive low-exposure firms.

Table 9. Dose-response by exposure quintile (dependent variable: log energy intensity).

Exposure quintile	Mean D	Coef.	SE	p	N
Q1 (lowest)	-1.07	-0.028	(0.070)	0.684	4,749
Q2	-0.82	0.191	(0.124)	0.123	4,749
Q3	-0.18	0.013	(0.050)	0.795	4,749
Q4	0.54	-0.263***	(0.093)	0.005	4,749
Q5 (highest)	1.52	-0.307***	(0.029)	0.000	4,749

Notes: Within-quintile DML-PLR (gradient boosting), firm-clustered standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.




Figure 5

Figure 5. Dose-response of extreme climate risk on log energy intensity: within-quintile DML estimates with 95% confidence intervals (red points are significant at the 5% level). The dashed line marks the estimated activation threshold ($D \approx 0.4$, around the 60th percentile of exposure).

5.7. Supplementary evidence from a real-data provincial panel

To complement the firm-level analysis with fully observed data, we estimate the same partially linear DML model on the provincial panel described in Sec. 4.3 (30 provinces, 2010-2019, 300 province-year observations), with province fixed effects absorbed by within transformation, year fixed effects included as features, and province-clustered standard errors. Figure 6 displays the national co-movement of mean extreme climate risk and mean log energy intensity over the sample period and the raw province-year scatter.




Figure 6

Figure 6. Provincial panel, 2010-2019. Panel (a): national annual means of the standardized extreme climate risk index and of log energy intensity. Panel (b): raw province-year scatter of log energy intensity against extreme climate risk.

Table 10 reports the baseline provincial estimates and Fig. 7 visualizes them. In the level specification, neither FE-OLS nor the DML learners detect a significant contemporaneous effect of extreme climate risk on provincial energy intensity: the DML point estimates range from -0.004 (Lasso) to $+0.010$ (random forest), all statistically insignificant. However, when the outcome is recast as the annual *change* in log energy intensity, the DML estimate turns negative and significant: a one-standard-deviation increase in extreme climate risk is associated with a 0.95-percentage-point faster annual decline in energy intensity (column 6; $p = 0.026$). The juxtaposition is informative: aggregate energy intensity is dominated by slow-moving industrial structure, so climate-driven adaptation shows up in the *speed* of improvement rather than in level jumps.

Table 10. Provincial baseline: effect of extreme climate risk on log energy intensity, 2010-2019.

	(1) OLS+FE	(2) OLS+FE+X	(3) DML- Lasso	(4) DML-RF	(5) DML- GBM	(6) DML- GBM, growth
ExtremeRisk	0.002 (0.009)	-0.006 (0.005)	-0.004 (0.005)	0.010 (0.007)	0.004 (0.006)	-0.010** (0.004)
Controls	No	Yes	Yes (ML)	Yes (ML)	Yes (ML)	Yes (ML)
Province & Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	300	300	300	300	300	270

Notes: Province-clustered standard errors in parentheses. Column (6) uses the first difference of log energy intensity as the outcome. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.




Figure 7

Figure 7. Provincial baseline estimates with 95% confidence intervals: FE-OLS and DML with three learners on the level outcome, and the DML estimate on the growth-rate outcome.

Table 11 reports robustness checks for the provincial level specification. Both placebo tests pass: a randomly permuted treatment ($p = 0.793$) and the lead treatment ($p = 0.164$) yield null effects, confirming that the estimation procedure does not manufacture spurious significance. Results are stable to the number of cross-fitting folds (2 or 10), to 5/95 trimming, to using the extreme-heat-days component alone, and to adding province-specific linear trends.

Table 11. Provincial robustness checks (dependent variable: log energy intensity).

Check	Coef.	SE	p	N
Extreme-heat-days component only	0.002	(0.006)	0.695	300
Placebo: permuted treatment	-0.002	(0.006)	0.793	300
Placebo: lead treatment $D(t+1)$	0.008	(0.006)	0.164	270
2-fold cross-fitting	0.008	(0.007)	0.283	300
10-fold cross-fitting	-0.000	(0.005)	0.984	300
Trim treatment at 5th/95th percentile	0.010	(0.008)	0.189	270
Province linear trends (GBM)	0.004	(0.006)	0.538	300
Province linear trends (Lasso)	-0.001	(0.005)	0.875	300

Notes: DML-PLR, province-clustered standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 12 and Fig. 8 report provincial heterogeneity. The level effect is marginally *positive* in non-eastern provinces (0.011, $p = 0.058$) and in southern provinces (0.012, $p = 0.087$), while essentially zero in eastern and northern provinces; cross-group differences, however, are not statistically significant. The marginal positive coefficients in less-developed regions

are consistent with short-run disruption dominating where adaptive capacity is weakest: when extreme events strike, these provinces tend to prioritize energy supply security over efficiency, temporarily raising energy intensity.

Table 12. Provincial heterogeneity (DML-GBM; dependent variable: log energy intensity).

Subgroup	Coef.	SE	N	Diff. p-value
Eastern	-0.002	(0.008)	100	
Non-eastern	0.011*	(0.006)	200	0.191
Southern	0.012*	(0.007)	150	
Northern	0.003	(0.011)	150	0.503
Coal-rich	-0.005	(0.021)	60	
Non-coal	0.005	(0.006)	240	0.648
High baseline energy intensity	0.016	(0.012)	150	
Low baseline energy intensity	-0.002	(0.007)	150	0.196

Notes: Province-clustered standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.




Figure 8

Figure 8. Provincial heterogeneity: subgroup DML estimates with 95% confidence intervals. The dashed line marks the full-sample DML-GBM estimate.

Finally, Table 13 extends the provincial analysis to carbon intensity (CEADs-based log CO₂ per unit of GDP) and to the clean-electricity share (proxied by the hydropower share of generation). Neither outcome shows a significant response at the annual provincial frequency, although both point estimates for the clean-electricity share are negative, hinting at short-run pressure on hydropower output during drought years.

Table 13. Provincial extended outcomes (DML-PLR).

Outcome	DML-Lasso	DML-GBM	N
---------	-----------	---------	---

Outcome	DML-Lasso	DML-GBM	N
Log carbon intensity	−0.003 (0.007)	0.003 (0.007)	300
Clean-electricity share (hydro proxy)	−0.011 (0.008)	−0.007 (0.006)	290

Notes: Province-clustered standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Taken together, the real-data provincial evidence paints a picture that is consistent with, and complementary to, the firm-level findings. Adaptation to extreme climate risk is detectable at the aggregate level primarily on the growth margin of energy intensity, while level effects are attenuated by regional aggregation—a dilution that directly motivates this paper’s firm-level design. The marginal short-run intensity increases in non-eastern and southern provinces reinforce the policy conclusion that adaptive capacity, not climate exposure alone, determines the energy-management response.

6. Conclusions and Policy Implications

This paper provides the first double machine learning evidence on the causal effect of extreme climate risk on corporate energy management. Using an unbalanced panel of 2,142 Chinese A-share listed firms over 2010-2023 and a firm-level composite index of extreme heat, extreme precipitation, and drought exposure, we show that extreme climate risk triggers a pronounced “forced adaptation” response: a one-standard-deviation increase in exposure reduces energy intensity by roughly one-tenth, compresses the energy-cost share, and expands the green energy share. The response is materially larger than conventional fixed-effects OLS would suggest, is concentrated among non-state-owned, energy-intensive, and eastern-region firms, and activates only beyond a climate-risk threshold around the sixtieth percentile of exposure. Green innovation and management-efficiency improvements account for the bulk (over 90%) of the total effect. Supplementary analysis of a real-data provincial panel (30 provinces, 2010-2019) corroborates this adaptation margin at the aggregate level: extreme climate risk significantly accelerates the annual decline of provincial energy intensity, even though level effects are undetectable once regional aggregation absorbs slow-moving structural change.

These findings carry several implications for policy and practice. First, physical climate risk, however damaging, generates an endogenous market incentive for energy upgrading that policy should amplify rather than substitute for: energy-efficiency subsidies, green credit, and energy-management-system certification will bind most effectively where climate exposure has already raised the private return to adaptation. Second, the muted response of state-owned, small, and central/western-region firms identifies precisely the segments in which market forces alone are insufficient; targeted instruments—energy-service company (ESCO) programs for small firms, region-specific adaptation finance, and energy auditing mandates for state-owned enterprises—should be directed there. Third, the adaptation threshold implies that support is most cost-effective when targeted at firms near the activation point of exposure—where marginal incentives are strongest—rather than at already-adapting high-exposure firms or unresponsive low-exposure firms. Fourth, the dominance of the green-innovation and management-efficiency channels implies that strengthening firms’ adaptive capacity (R&D incentives, management training, digital

energy monitoring) can substitute for exposure itself, decoupling energy upgrading from suffered climate damages. Finally, for the research community, the twofold gap between DML and conventional FE-OLS estimates underscores that estimator choice is not a technical footnote but a first-order substantive decision in climate-energy empirics.

Several limitations mark the boundaries of our contribution and chart an agenda for future work. First, identification rests on conditional exogeneity given a rich covariate set; although our placebo and robustness tests are reassuring, quasi-experimental variation—abrupt meteorological disasters or policy shocks—could complement the DML design, potentially through DML-based difference-in-differences [23]. Second, we measure energy management through intensity, cost, and structure margins; direct evidence on firms' energy-management systems and internal energy audits would sharpen the behavioral interpretation. Third, our heterogeneity analysis is subgroup-based; generalized random forests and generic ML inference offer a natural next step toward continuous treatment-effect heterogeneity [49, 50, 57]. Fourth, the analysis is static at the annual frequency; dynamic DML for panel data would allow an explicit separation of short-run disruption from long-run adaptation [54].

Acknowledgments

[The author thanks seminar participants at Macao Polytechnic University for helpful comments. Financial support from [Foundation, grant number] is gratefully acknowledged. All errors are the author's own.]

Declaration of competing interest

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability statement

The data and replication codes that support the findings of this study are available from the corresponding author upon reasonable request.

References

- [1] Bilal, A., Känzig, D.R., 2026. The macroeconomic impact of climate change: Global versus local temperature. *Quarterly Journal of Economics* 141(2).
- [2] Dell, M., Jones, B.F., Olken, B.A., 2012. Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics* 4(3), 66-95.
- [3] Burke, M., Hsiang, S.M., Miguel, E., 2015. Global non-linear effect of temperature on economic production. *Nature* 527(7577), 235-239.
- [4] Addoum, J.M., Ng, D.T., Ortiz-Bobea, A., 2020. Temperature shocks and establishment sales. *Review of Financial Studies* 33(3), 1331-1366.
- [5] Pankratz, N., Bauer, R., Derwall, J., 2023. Climate change, firm performance, and investor surprises. *Management Science* 69(12), 7352-7398.
- [6] Somanathan, E., Somanathan, R., Sudarshan, A., Tewari, M., 2021. The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing. *Journal of Political Economy* 129(6), 1797-1827.
- [7] Zhang, P., Deschenes, O., Meng, K., Zhang, J., 2018. Temperature effects on productivity and factor

reallocation: Evidence from a half million Chinese manufacturing plants. *Journal of Environmental Economics and Management* 88, 1-17. [8] Chen, X., Yang, L., 2019. Temperature and industrial output: Firm-level evidence from China. *Journal of Environmental Economics and Management* 95, 257-274. [9] Allcott, H., Greenstone, M., 2012. Is there an energy efficiency gap? *Journal of Economic Perspectives* 26(1), 3-28. [10] Jaffe, A.B., Stavins, R.N., 1994. The energy-efficiency gap: What does it mean? *Energy Policy* 22(10), 804-810. [11] Gerarden, T.D., Newell, R.G., Stavins, R.N., 2017. Assessing the energy-efficiency gap. *Journal of Economic Literature* 55(4), 1486-1525. [12] Bloom, N., Genakos, C., Martin, R., Sadun, R., 2010. Modern management: Good for the environment or just hot air? *Economic Journal* 120(544), 551-572. [13] Boyd, G.A., Curtis, E.M., 2014. Evidence of an “energy-management gap” in U.S. manufacturing: Spillovers from firm management practices to energy efficiency. *Journal of Environmental Economics and Management* 68(3), 463-479. [14] Hou, X., Xiang, R., 2025. Exploring the influence of extreme heat on corporate energy efficiency. *Journal of Asian Economics* 97, 101899. [15] Srivastav, S.P., Kannadhasan, M., 2025. Environment and energy: Does climate risk shape the energy consumption behavior of firms? *Energy Economics* 148, 108630. [16] Wang, H., Tan, X., Li, M., 2025. Extreme climate attention and corporate energy innovation: Empirical evidence from China. *Energy Economics* 148, 108681. [17] Mullainathan, S., Spiess, J., 2017. Machine learning: An applied econometric approach. *Journal of Economic Perspectives* 31(2), 87-106. [18] Athey, S., Imbens, G.W., 2019. Machine learning methods that economists should know about. *Annual Review of Economics* 11, 685-725. [19] Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., Robins, J., 2018. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal* 21(1), C1-C68. [20] Belloni, A., Chernozhukov, V., Hansen, C., 2014. Inference on treatment effects after selection among high-dimensional controls. *Review of Economic Studies* 81(2), 608-650. [21] Belloni, A., Chernozhukov, V., Hansen, C., 2014. High-dimensional methods and inference on structural and treatment effects. *Journal of Economic Perspectives* 28(2), 29-50. [22] Yang, J.-C., Chuang, H.-C., Kuan, C.-M., 2020. Double machine learning with gradient boosting and its application to the Big N audit quality effect. *Journal of Econometrics* 216(1), 268-283. [23] Xing, L., Han, D., Hui, X., 2023. The impact of carbon policy on corporate risk-taking with a double/debiased machine learning based difference-in-differences approach. *Finance Research Letters* 58, 104502. [24] Feng, L., Qi, J., Zheng, Y., 2025. How can AI reduce carbon emissions? Insights from a quasi-natural experiment using generalized random forest. *Energy Economics* 141, 108040. [25] Krueger, P., Sautner, Z., Starks, L.T., 2020. The importance of climate risks for institutional investors. *Review of Financial Studies* 33(3), 1067-1111. [26] Giglio, S., Kelly, B., Stroebe, J., 2021. Climate finance. *Annual Review of Financial Economics* 13, 15-36. [27] Bernstein, A., Gustafson, M.T., Lewis, R., 2019. Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics* 134(2), 253-272. [28] Sautner, Z., van Lent, L., Vilkov, G., Zhang, R., 2023. Firm-level climate change exposure. *Journal of Finance* 78(3), 1449-1498. [29] Li, Q., Shan, H., Tang, Y., Yao, V., 2024. Corporate climate risk: Measurements and responses. *Review of Financial Studies* 37(6), 1778-1830. [30] Addoum, J.M., Ng, D.T., Ortiz-Bobea, A., 2023. Temperature shocks and industry earnings news. *Journal of Financial Economics* 150(1), 1-45. [31] Ye, Y., Zhu, J., Li, B., Yang, X., 2024. Extreme climate risks and corporate bond yield spreads: Evidence from China. *Asia-Pacific Journal of Accounting & Economics* 31(5), 697-717. [32] Bolton, P., Kacperczyk, M., 2021.

Do investors care about carbon risk? *Journal of Financial Economics* 142(2), 517-549. [33] Ilhan, E., Sautner, Z., Vilkov, G., 2021. Carbon tail risk. *Review of Financial Studies* 34(3), 1540-1571. [34] DeCanio, S.J., Watkins, W.E., 1998. Investment in energy efficiency: Do the characteristics of firms matter? *Review of Economics and Statistics* 80(1), 95-107. [35] Gillingham, K., Palmer, K., 2014. Bridging the energy efficiency gap: Policy insights from economic theory and empirical evidence. *Review of Environmental Economics and Policy* 8(1), 18-38. [36] Worrell, E., Laitner, J.A., Ruth, M., Finman, H., 2003. Productivity benefits of industrial energy efficiency measures. *Energy* 28(11), 1081-1098. [37] Martin, R., Muuls, M., de Preux, L.B., Wagner, U.J., 2012. Anatomy of a paradox: Management practices, organizational structure and energy efficiency. *Journal of Environmental Economics and Management* 63(2), 208-223. [38] Porter, M.E., van der Linde, C., 1995. Toward a new conception of the environment-competitiveness relationship. *Journal of Economic Perspectives* 9(4), 97-118. [39] Fisher-Vanden, K., 2003. Management structure and technology diffusion in Chinese state-owned enterprises. *Energy Policy* 31(3), 247-257. [40] Fisher-Vanden, K., Jefferson, G.H., Liu, H., Tao, Q., 2004. What is driving China's decline in energy intensity? *Resource and Energy Economics* 26(1), 77-97. [41] Kostka, G., Moslener, U., Andreas, J., 2013. Barriers to increasing energy efficiency: Evidence from small- and medium-sized enterprises in China. *Journal of Cleaner Production* 57, 59-68. [42] Cao, J., Karplus, V.J., 2014. Firm-level determinants of energy and carbon intensity in China. *Energy Policy* 75, 167-178. [43] Chen, D., Chen, S., Jin, H., Lu, Y., 2020. The impact of energy regulation on energy intensity and energy structure: Firm-level evidence from China. *China Economic Review* 59, 101351. [44] Ren, X., Zhang, X., Yan, C., Gozgor, G., 2022. Climate policy uncertainty and firm-level total factor productivity: Evidence from China. *Energy Economics* 113, 106209. [45] Niu, N., Ma, J., Zheng, D., Lu, Y., Zhang, B., 2025. Extreme weather and the green transition of energy firms: The moderating effect of digital technology and digital inclusive finance. *Research in International Business and Finance* 76, 102858. [46] Wen, J., Zhao, X.-X., Fu, Q., Chang, C.-P., 2023. The impact of extreme weather events on green innovation: Which ones bring to the most harm? *Technological Forecasting and Social Change* 188, 122322. [47] Wang, C., Li, W., Li, J., Wan, L., 2024. The effects of extreme climate events on green technology innovation in manufacturing enterprises. *Journal of Business Economics and Management* 25(2), 337-355. [48] Varian, H.R., 2014. Big data: New tricks for econometrics. *Journal of Economic Perspectives* 28(2), 3-28. [49] Wager, S., Athey, S., 2018. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113(523), 1228-1242. [50] Athey, S., Tibshirani, J., Wager, S., 2019. Generalized random forests. *Annals of Statistics* 47(2), 1148-1178. [51] Farrell, M.H., Liang, T., Misra, S., 2021. Deep neural networks for estimation and inference. *Econometrica* 89(1), 181-213. [52] Knaus, M.C., Lechner, M., Strittmatter, A., 2021. Machine learning estimation of heterogeneous causal effects: Empirical Monte Carlo evidence. *The Econometrics Journal* 24(1), 134-161. [53] Semenova, V., Chernozhukov, V., 2021. Debiased machine learning of conditional average treatment effects and other causal functions. *The Econometrics Journal* 24(2), 264-289. [54] Semenova, V., Goldman, M., Chernozhukov, V., Taddy, M., 2023. Inference on heterogeneous treatment effects in high-dimensional dynamic panels under weak dependence. *Quantitative Economics* 14(2), 471-510. [55] Ahrens, A., Hansen, C.B., Schaffer, M.E., Wiemann, T., 2024. ddml: Double/debiased machine learning in Stata. *Stata Journal* 24(1), 3-45. [56] Ahrens, A., Hansen, C.B., Schaffer, M.E., Wiemann, T., 2025. Model

averaging and double machine learning. *Journal of Applied Econometrics* 40(3), 249-269. [57] Chernozhukov, V., Demirer, M., Duflo, E., Fernandez-Val, I., 2025. Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in India. *Econometrica* 93(4), 1127-1163. [58] Deschenes, O., Greenstone, M., 2011. Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US. *American Economic Journal: Applied Economics* 3(4), 152-185. [59] Auffhammer, M., Mansur, E.T., 2014. Measuring climatic impacts on energy consumption: A review of the empirical literature. *Energy Economics* 46, 522-530. [60] Knittel, C.R., Stolper, S., 2021. Machine learning about treatment effect heterogeneity: The case of household energy use. *AEA Papers and Proceedings* 111, 440-444. [61] Wei, R., Xia, Y., 2024. Digital transformation and corporate green total factor productivity: Based on double/debiased machine learning robustness estimation. *Economic Analysis and Policy* 84, 808-827. [62] Huang, Z., Zhang, Q., Zheng, Y., Huang, F., Guo, R., 2025. Heterogeneous impact of climate policy uncertainty on regional economic development in China: Evidence from double machine learning. *Journal of Environmental Management* 393, 127046. [63] Shan, Y., Guan, D., Zheng, H., Ou, J., Li, Y., Meng, J., Mi, Z., Liu, Z., Zhang, Q., 2018. China CO2 emission accounts 1997-2015. *Scientific Data* 5, 170201. [64] Shan, Y., Huang, Q., Guan, D., Hubacek, K., 2020. China CO2 emission accounts 2016-2017. *Scientific Data* 7, 54. [65] NOAA National Centers for Environmental Information, Global Surface Summary of the Day (GSOD) [Dataset], <https://www.ncei.noaa.gov/data/global-summary-of-the-day/access/> (accessed 2026).

Appendix A. The DML estimation algorithm

The estimation procedure proceeds as follows.

Algorithm 1 (DML-PLR with cross-fitting).

Input: panel $\{Y_{it}, D_{it}, X_{it}\}$; learner $\mathcal{A}$; folds $K = 5$; repetitions $R = 3$.

1. *Within transformation:* demean Y , D , and time-varying controls at the firm level to absorb firm fixed effects; augment the feature set with firm means of time-varying controls, industry dummies, and year dummies.
2. *Partition:* randomly split the n observations into K folds $\{I_1, \dots, I_K\}$ of approximately equal size.
3. *Nuisance training:* for each fold k , train $\hat{\ell}_{I_k^c} = E[Y | X]$ and $\hat{m}_{I_k^c} = E[D | X]$ with learner $\mathcal{A}$ on the complement I_k^c ; predict both out-of-fold for all $i \in I_k$.
4. *Orthogonal estimation:* form residuals $\tilde{V}_i = D_i - \hat{m}(X_i)$, $\tilde{W}_i = Y_i - \hat{\ell}(X_i)$ and compute $\hat{\theta}$ from Eqn. (9).
5. *Inference:* compute the variance from the orthogonal scores in Eqn. (10), aggregating scores at the firm level to obtain cluster-robust standard errors.
6. *Repeat* steps 2-5 over R random partitions and aggregate by the median-based procedure of [19].

Output: point estimate $\hat{\theta}$, cluster-robust standard error, and confidence interval.

Appendix B. Additional implementation details

All learners are implemented with fixed random seeds. Lasso uses 60 candidate penalty values with 3-fold cross-validation on the expanded basis described in Sec. 4.2. Random forest uses 50-120 trees, minimum leaf size 10, and feature sampling rate 0.5. Histogram gradient boosting uses at most 500 iterations with early stopping after 30 stagnant rounds. The neural network uses two hidden layers of 48 and 24 units with L_2 penalty 0.01 and early stopping. Insensitivity checks confirm that doubling the number of trees changes the random-forest estimate by less than 0.003, and that 2-fold versus 10-fold cross-fitting changes the estimate by less than 0.004 (Table 5).