

Extreme Heat and Directed Innovation

Enjie (Jack) Ma
Cornell University*
em686@cornell.edu

October 14, 2025

Abstract

Can economies innovate their way out of climate trouble? I provide the first systematic evidence outside agriculture that firms adapt to extreme heat through directed technological change. Linking firm-level data with patent records for nine EU countries (2000–2020), I establish three findings. First, extreme heat operates as a labor-biased productivity shock: labor-intensive firms suffer disproportionate losses and cede market share to capital-intensive competitors. Second, firms respond by shifting production toward capital and redirecting innovation toward labor-saving technologies, with the strongest responses in labor-intensive industries facing the greatest heat exposure. Third, this endogenous innovation response has quantitatively meaningful consequences: labor-saving patents filed in response to heat attenuate aggregate productivity losses by 26 percent over the study period. These findings demonstrate that innovation is not merely a driver of growth but an active margin of climate adaptation.

Keywords: Climate change, Climate adaptation, directed technical change.

JEL Classification Codes: Q54, O33, Q55

*I am grateful to my advisors Shanjun Li, Ivan Rudik, Ezra Oberfield, and Julieta Caunedo for their guidance and support. I also thank Jacob Moscona, Zebang Xu, Todd Gerarden, Levon Barseghyan, Adam Harris, and seminar participants at Cornell for their helpful comments. All errors are my own.

1 Introduction

Climate change is reshaping the global economy, with extreme heat—one of its clearest and most precisely measured manifestations—already depressing productivity and growth. Innovation is the engine of long-run prosperity and a critical margin of economic resilience, yet a fundamental question remains unresolved: can economies innovate their way out of climate trouble? Existing research on climate adaptation has focused on quantifying the degree of adaptation and evaluating policy responses, particularly in agriculture and public health (Burke and Emerick, 2016; Burke et al., 2024; Barreca et al., 2016). Far less is known about whether and how innovation serves as an adaptation margin in the industrial sector—whether climate risk redirects inventive activity toward resilience-enhancing technologies, and if so, whether that redirection materially cushions economic damages. These questions carry high stakes: if innovation adapts effectively, historical damage estimates may overstate future costs.

This paper provides systematic evidence outside agriculture on how extreme heat shapes the direction of innovation and quantifies the economic payoff from induced technological change. Using firm-level data from nine EU countries and regional patent records spanning 2000–2020, I exploit variation in extreme heat exposure—which depresses labor productivity across industries (Graff Zivin and Neidell, 2014; Somanathan et al., 2021)—to answer two questions: Does extreme heat redirect innovation toward labor-saving technologies? And to what extent does induced innovation attenuate heat-related productivity damage?

I begin by establishing that extreme heat operates as a *labor-biased* productivity shock. Comparing firms with heterogeneous factor intensities, I document that labor-intensive firms experience substantially larger productivity losses under extreme heat, whereas capital-intensive firms exhibit only mild effects. Market shares reallocate accordingly: labor-intensive firms lose ground to capital-intensive competitors under heat.

How do firms respond to this labor-biased climate shock? I examine two underexplored adaptation margins: *technology change*—shifts in production methods that reduce labor intensity—and *innovation*—redirection of inventive activity toward labor-saving technologies. Both margins reveal clear movement away from the heat-vulnerable labor input.

Measuring technology change through output elasticities estimated from time-varying translog production functions (Akerberg et al., 2015), I find that capital elasticity rises while labor elasticity falls under heat. Ten additional days above 90°F raise capital elasticity by 0.46 percent and reduce labor elasticity by 0.07 percent, relative to the [60, 70) F reference bin. Exploiting the structural relationship between output elasticities and production input under the translog specification, I demonstrate that these shifts in output elasticity reflect input substitution: firms reallocate from labor toward capital in response to heat. At the industry level, ex ante labor-intensive sectors increase aggregate capital intensity in hotter years, driven by both within-firm technological adjustment and between-firm reallocation.

I then turn to the innovation margin and test whether labor-biased heat exposure reorients inventive effort toward labor-saving (LS) technologies. LS patents are identified through a unified classification: a patent qualifies as LS if any CPC/IPC code appears in the automation set of Hémous et al. (2025) or if its title contains canonical automation keywords. Using a region–industry–year panel constructed from OECD REGPAT and interacting regional cooling degree-days above 85°F with pre-sample industry labor intensity, $LI_{i,0}$, I estimate Poisson (PPML) models with high-dimensional fixed effects. Hotter conditions systematically shift innovation toward LS technologies, and the response is markedly stronger in labor-intensive industries. Crucially, the result persists across multiple lag structures: the $CDD \times LI$ interaction

remains positive and statistically significant when temperature exposure is lagged two to five years, indicating that innovation responds to sustained heat exposure rather than transitory fluctuations. This persistence aligns with models of directed technical change under R&D gestation lags and adjustment costs—firms require a persistent signal before committing resources to new technologies.

To organize these empirical findings conceptually, I embed the labor-biased heat shock in a directed technical change (DTC) framework (Acemoglu, 2002). Innovators choose between capital-augmenting and labor-augmenting R&D; a heat shock that reduces effective labor supply raises the relative return to labor-saving innovation. When the elasticity of substitution between capital and labor is sufficiently high, equilibrium innovation endogenously tilts toward the less heat-sensitive capital sector. Importantly, because such innovation, when successful, raises factor efficiency and expands the production possibility frontier, the model also provides a mechanism through which induced technological change can offset the productivity losses caused by extreme heat.

Motivated by this DTC logic, I quantify the role of labor-saving innovation in mitigating heat-induced productivity damage. At the firm level, I compare heat–TFP responses for firms that did versus did not file an LS patent in the prior two years. The heterogeneity is striking: under an additional 100 Cooling Degree Days above 85°F , non-innovating firms experience a productivity decline of 0.05 percent, whereas firms with recent LS patents see productivity *increase* by 0.56 percent. Firms that successfully redirect inventive effort toward labor-saving technologies are not merely insulated from climate damages—they realize net productivity gains under extreme heat.

Building on these two empirical relationships—that heat systematically redirects innovation toward LS technologies, and that LS innovation in turn flattens the heat–productivity gradient—I construct a back-of-the-envelope (BOTE) calculation that compares predicted aggregate productivity under two scenarios: the observed world, where innovation responds endogenously to heat, and a counterfactual "no-innovation" world with the same temperature trajectory but where LS innovation does not respond to heat. The exercise reveals that directed LS innovation offset roughly **26 percent** of aggregate heat-related productivity losses over 2000–2020 in my EU sample. Absent the redirection of inventive activity documented in the patent data, productivity damages from extreme heat would have been approximately one-third larger. I interpret this 26 percent mitigation as a lower bound on innovation’s adaptive contribution: the calculation captures only the intensive margin of in-house LS innovation by incumbent producers, omitting technology diffusion to non-patenting firms, spatial reallocation within the integrated EU market, general equilibrium price adjustments, and entry-exit dynamics. The mitigation share varies substantially across countries (roughly 0–40 percent), reflecting heterogeneity in both warming severity and the scope for LS redirection given local industrial composition.

These findings establish that innovation is not a passive accompaniment to economic growth but an active, endogenous response to environmental stress—one that materially shapes the economic consequences of climate change. The direction of technical change adapts to the structure of damages, and this adaptation recovers a quantitatively meaningful share of lost output. Yet the results also underscore the limits of spontaneous, market-driven adaptation: even in a high-capacity institutional environment, roughly three-quarters of heat-induced productivity losses persist. Directed innovation is a powerful margin of climate adaptation, but far from sufficient on its own. The policy implication is clear: fostering innovation capacity, accelerating technology diffusion, and correcting the market failures that constrain adaptive investment should be central pillars of climate resilience strategies—alongside emissions mitigation and social insurance

against climate damages.

Related Literature This paper speaks to three literatures. First, it contributes to research on climate impacts on firms and productivity. A growing body of evidence documents that rising temperatures depress economic performance through reduced labor supply (Graff Zivin and Neidell, 2014), diminished on-the-job productivity (Somanathan et al., 2021), and lower plant-level output (Zhang et al., 2018; Chen and Yang, 2019; Addoum et al., 2020). Recent work explores heterogeneous impacts across firms (Xie, 2024; Ponticelli et al., 2023). These studies typically estimate climate–productivity relationships holding technology fixed. I show that this assumption can be misleading: production technology and innovation direction respond endogenously to climate exposure, and this adjustment materially flattens the damage gradient. Ignoring endogenous technical change risks overstating net climate impacts. The closest paper is Long and Wang (2025), who document increased climate-related patenting in China. My analysis differs in two ways: I exploit cross-firm and cross-industry variation in labor intensity to identify a *factor-biased* mechanism, and I quantify whether induced innovation mitigates realized productivity damages, finding that labor-saving innovation offsets roughly 26 percent of aggregate heat-related losses.

Second, the paper extends the climate adaptation literature. Existing work has focused on whether aggregate damage functions have flattened over time (Burke et al., 2024; Dell et al., 2014) or on specific behavioral margins such as agricultural practices (Burke and Emerick, 2016), mortality responses (Barreca et al., 2016), and energy use (Heutel et al., 2021). Moscona and Sastry (2023) is the first paper that examines directed innovation as adaptation to climate shocks, showing that agricultural technology reorients toward climate-stressed crops and mitigates 19 percent of heat-induced yield losses in U.S. agriculture. I extend this framework from cross-crop to cross-factor heterogeneity. Where Moscona and Sastry (2023) exploit variation in crop exposure, I leverage variation in industry labor intensity interacted with spatial heat exposure, linking innovation direction to the structure of climate damages at the production input factor level. This reveals that industries facing labor-biased heat shocks systematically redirect inventive effort toward labor-saving technologies and suggests that adaptation through innovation is not confined to agriculture sector but operates as a first-order margin in industrial sector as well.

Third, the paper also speaks to the directed technical change (DTC) literature. Acemoglu (2002) establishes that innovation responds endogenously to market size and relative factor prices, with empirical applications spanning energy (Newell et al., 1999; Popp, 2002), environmental regulation (Acemoglu et al., 2012; Calel and Dechezleprêtre, 2016; Aghion et al., 2016), and labor costs (Hémous et al., 2025; Acemoglu and Restrepo, 2022). I provide the first evidence that climate shocks can induce factor-biased technological change. The mechanism I document—a labor-biased climate shock that reduces effective labor productivity and shifts innovation toward capital-augmenting technologies—extends the directed technical change framework beyond traditional factor price shocks to exogenous productivity shocks driven by environmental conditions. Furthermore, while the current DTC-environment literature focuses on technologies related to climate mitigation (low-carbon technologies that reduce future emissions), my paper focuses on directed technologies related to climate adaptation (technologies that cope with existing impacts).

The paper also touches on several adjacent literatures. The evidence of labor-biased heat damage and the consequent capital-deepening adaptation contributes to the literature on falling labor share (Grossman and Oberfield, 2022) and the broader debate over whether this decline is driven by technological change, globalization, or shifts in market power. Relatedly, the finding that capital-intensive firms are less vulnerable

to heat connects to work on automation’s distributional consequences (Acemoglu and Restrepo, 2022; Autor et al., 2020), suggesting climate change may accelerate labor-displacing technological change. The heterogeneity in adaptive capacity across firms relates to misallocation and productivity dispersion (Baqee and Farhi, 2020), as climate shocks may widen gaps between innovating and non-innovating firms. Finally, the role of patents connects to the broader innovation literature on knowledge spillovers and technology diffusion (Bloom et al., 2019, 2013).

The rest of the paper proceeds as follows. Section 2 describes the data construction and measurement. Section 3 establishes that extreme heat operates as a labor-biased productivity shock, documenting differential effects across firms and the resulting within-market reallocation from labor-intensive to capital-intensive producers. Section 4 examines how firms adapt along two margins: first, short-run adjustments in production technology and factor usage, and second, medium-run redirection of inventive effort toward labor-saving innovation. Section 5 quantifies the aggregate economic consequences of induced innovation, showing that directed technological change materially attenuates heat-related productivity losses. Section 6 concludes and discusses implications for climate policy and future research.

2 Data

My analysis examines how temperature shocks affect firm-level outcomes and innovation responses in heat-exposed industries across Europe. This requires combining three distinct data sources: (i) granular climate data to measure local temperature exposure, (ii) comprehensive firm-level panel data to track economic outcomes, and (iii) patent records to identify innovation responses.

I construct firm-specific climate exposure measures from daily temperature and precipitation data, matched to each establishment’s geographic coordinates. I link these to firm balance-sheet data from ORBIS, focusing on heat-exposed sectors. To measure innovation responses, I combine patent applications with a classification scheme that identifies labor-saving technologies, and I match patents to firms using harmonized applicant names. The final data is a matched panel in which I observe temperature exposure, economic outcomes, and patenting activity.

The remainder of this section proceeds as follows. Section 2.1 describes the construction of temperature and precipitation variables. Section 2.2 details the firm sample and variable definitions. Sections 2.3–2.5 explain the patent data sources, the classification of labor-saving innovations, and the matching procedure.

2.1 Climate Data

This subsection describes the construction of my climate exposure variables, provides summary statistics on their magnitudes and cross-country variation, and introduces the figures that document these patterns.

Temperature. I obtain daily temperature from the *NOAA Physical Sciences Laboratory* (PSL), provided on a $0.5^\circ \times 0.5^\circ$ global grid. I assign the *daily maximum* temperature to each firm using its recorded latitude–longitude.¹ I focus on daily maxima because they capture the daytime conditions workers face and therefore better mirror heat exposure than daily averages (Graff Zivin and Neidell, 2014; Somanathan et al., 2021).

From this daily *TMAX* series, I construct two complementary annual measures of heat exposure:

¹When coordinates are missing, I average over the firm’s postal code as a proxy for local exposure.

(i) *Daily-maximum temperature bins (TMAX bins)*. I group daily maxima into eight 10°F intervals: <30°F, 30–40°F, 40–50°F, 50–60°F, 60–70°F, 70–80°F, 80–90°F, and ≥90°F. Each bin records the number of firm-days in that range within a year; the eight counts sum to 365.

(ii) *Cooling degree-days (CDD)*. I also compute a single continuous index that aggregates the intensity of hot days above a threshold:

$$\text{CDD}_{rt} = \sum_{d=1}^{365} \max\{0, \text{TMAX}_{rt,d} - 85^\circ\text{F}\},$$

reported in °F-days.²

The two measures serve different purposes. TMAX bins offer a flexible way to locate *where* in the temperature distribution responses arise and are well suited to describing contemporaneous, within-year variation. CDD, by contrast, collapses the high-temperature tail into a parsimonious intensity index above a damage threshold, which is convenient in specifications with interactions, lags, or high-dimensional fixed effects (as is standard in patent regressions). Because both are built from the same raw series of daily maximum data and are informative of exposure to extreme heat, they deliver similar qualitative patterns throughout the empirical results. I use bins when flexibility is paramount and CDD when parsimony helps identification.

Across firm-year observations in the EU sample (2000–2020), the annual CDD distribution has a median of about 94 °F-days (mean ≈ 180), with the 90th percentile near 478 and a long upper tail (max ≈ 1,887). Year-to-year within-firm changes are sizable: the median absolute change is about 43 °F-days (90th percentile ≈ 136). Cross-country differences are pronounced, with higher exposures in Mediterranean economies and lower exposures in the Nordics.³

Figure 1 summarizes these patterns. Panel (a) shows that increases in cooling degree-days between 2000 and 2020 are concentrated in southern and western Europe, particularly Spain and France, while northern and central countries exhibit modest changes or mild declines. Panel (b) illustrates the corresponding shift in the daily maximum temperature distribution: compared to 2000, the year 2020 exhibits fewer days in cooler bins (30–60°F) and noticeably more days in the hottest ≥90°F bin.⁴ Together these patterns highlight both the spatial heterogeneity and the compositional shifts in heat exposure that underlie the firm-level variation exploited in later sections. Appendix A1 reports additional summary tables, smoothed maps, and country-level distributions.


Precipitation. Precipitation comes from the same NOAA PSL gridded product (0.5° resolution). I extract daily precipitation at each firm’s location and average to the calendar year.

Aggregation. Firm-level TMAX bins, CDD, and precipitation enter the firm-year analysis directly. For market-level statistics, I aggregate to country×NACE-4 using time-varying market-share weights.


²The 85°F (29.4°C) threshold aligns with evidence of sharp nonlinear losses in labor supply and productivity as temperatures exceed this point. Park et al. (2021) find workplace injuries increase 5-7% in the 85-90°F range relative to days in the 60s, while Graff Zivin and Neidell (2014) document that workers in high-exposure industries work nearly an hour less when daily maximum temperatures exceed 85°F. Somanathan et al. (2021) show productivity declines of 2% per degree Celsius in Indian manufacturing. This range reflects physiological heat stress impacts on labor rather than building comfort thresholds.

³Country-level averages and additional distributional figures are summarized in Appendix A1.

⁴The average increase for Bin [≥90F] is around 10 more days.



(a) Δ CDD: 2020–2000 (EU)



(b) Shift in daily max temperature bins, 2000 vs. 2020

Figure 1: Changes in climate exposure over the sample period

Notes: Panel (a): change in cooling degree-days, $\Delta \text{CDD} = \text{CDD}_{2020} - \text{CDD}_{2000}$ ($\text{TMAX} > 29.4^\circ\text{C} / 85^\circ\text{F}$). Panel (b): average daily maximum temperature bin counts in 2000 and 2020 across firms/countries; bins are eight 10°F intervals summing to 365 days.

2.2 Firm Data

To measure firm-level economic outcomes, I use the ORBIS Database provided by Bureau van Dijk, which compiles firm-level financial statements and balance sheets from national information providers across Eu-

rope. I extract data on annual revenue, labor costs, capital costs, and the number of employees. I also use firm-level geographical information (address, longitude, and latitude) to match ORBIS records to the weather data described above.

I estimate firm-level production functions using the ACF control function approach (Akerberg et al., 2015). To allow production technology to evolve over time, the estimation is performed on five-year rolling windows, following Hubmer and Restrepo (2021). This procedure recovers firm-level revenue Total Factor Productivity (TFP) and output elasticities with respect to capital and labor. Appendix A3 provides details on the production function estimation.

Following Graff Zivin and Neidell (2014), I define "heat-exposed" industries as those in which workers are particularly vulnerable to high temperatures: mining, manufacturing, utilities, and construction (NACE Rev. 2 sections B, C, D, F). While agriculture is also heat-exposed in principle, ORBIS does not capture farm-level production establishments but rather headquarters and logistical firms. I therefore exclude NACE section A to avoid misclassifying non-production activities as directly climate-exposed. Appendix A2 provides additional details on data cleaning and imputation procedures.

2.3 OECD REGPAT and TPF Data

To measure innovation activity at the regional and industry level, I use patent application data from the OECD REGPAT database (January 2024 release), which links European Patent Office (EPO) and Patent Cooperation Treaty (PCT) applications to geographic regions. I extract each patent's NUTS3 region—the finest European regional classification—from applicant address data (EPO_APP_REG table) and retrieve the associated technology codes (CPC_CLASS table) based on the Cooperative Patent Classification (CPC) system. To assign each patent to a 4-digit NACE industry, I merge these records with the Goldschlag et al. (2016) probabilistic crosswalk that maps CPC codes to ISIC industries, followed by a standard ISIC-to-NACE conversion. When a patent lists multiple CPC codes, I apply fractional weighting so that each patent's total weight sums to one across industries, avoiding double-counting. I also perform standard data-cleaning steps—deduplication of application identifiers, consistency checks on region codes, and standardization of CPC strings—to ensure internally consistent observations.

I supplement these data with patent titles from the OECD Triadic Patent Families (TPF) dataset (TPF_EPO file).⁵ Each triadic family member links to REGPAT via a common application identifier (`Appln_id`), allowing me to add English-language titles to the patent records. The combined dataset contains, for each application, the NUTS3 region of the applicant, a set of CPC technology codes (fractionally weighted by industry), and a descriptive title. After filtering out records with incomplete identifiers, I obtain a final sample of EPO applications ready for classification and aggregation to the region–industry–year level. The final data includes 89068 region–industry–year observations from 2000 to 2020.

2.4 Classification of Labor-Saving Patents

I construct a single indicator for labor-saving (LS) innovation that combines technology-code information with patent text. A patent is classified as LS if *either* (i) any of its CPC/IPC technology codes belongs to

⁵Triadic patents are those filed at the EPO, USPTO, and JPO, representing high-value inventions with international market potential.

the automation set defined by [Hémous et al. \(2025\)](#) (DHOZ) *or* (ii) its title contains automation keywords drawn from the canonical terminology used in DHOZ and the broader automation literature.

For the code component, I adopt the DHOZ automation list, which is obtained by computing the prevalence of automation terms at the technology-class level and flagging classes above a high threshold (their *auto95* cutoff; examples include classes associated with robots, numerical control/CNC, and CAD/CAM). I match these classes to REGPAT using the CPC classification table (CPC_CLASS) and set the code-based flag to one if any listed code is present. For the patent title text component, I standardize titles (lowercasing, accent stripping) and search for a curated lexicon of automation terms—e.g., *robot*, *automat** (with stemming), *numerical control/CNC*, *computer-aided/CAD/CAM*, and closely related synonyms—using word-boundary matching to avoid spurious hits. The final LS indicator equals one if either the code-based flag or the text-based flag equals one. Table 1 shows a snapshot of the patent application that is classified as labor-saving.

Table 1: Examples of labor-saving patents

PATENT APPLICATION TITLE	CPC
Robot vacuum cleaner equipped with at least one fixed lateral cleaning element	A47L11
Apparatus for filling food containers and corresponding method	A23G9
Method for producing a relief-like printed image on containers	B29C64
Feed device for automatically feeding a sheet of material into a stretching unit, installation and method	B29C2037
Device for attaching a sprayer to a robot arm	B05B13

Finally, I aggregate to the region–industry–year level by computing the count (and share) of labor-saving patent applications in each NUTS-3 region $\times$ NACE 4-digit industry $\times$ year cell. These counts serve as my measure of local labor-saving innovation intensity over time. Table 2 shows the summary of the region-industry panel.

Table 2: Region-Industry Panel — Summary Statistics

Variable	N	Mean	SD	Median	Max
Total Patents	89,068	6	22.91	2	1,724
Labor-saving Patents	89,068	0.21	0.90	0	61
Labor-saving Share	89,068	0.06	0.20	0	1
CDD	89,068	24.07	47.37	9.48	913.45
Labor Intensity	80,701	0.44	0.06	0.43	0.72

Figure 2 shows the spatial distribution of LS patent share across NUTS 3 regions in my EU sample. Spain, France, and Germany have relatively higher concentration of labor-saving innovations. Figure 3 shows the trend of labor-saving patent share over time from 2000 to 2020. The stark rise around 2010s is consistent

Figure 2: Share of Labor-saving Patents





Figure 3: Labor-saving patents application share

with the increased automation activity documented in the literature ([Hémous et al., 2025](#)).

2.5 Linking REGPAT Patents to ORBIS Firms

To link firm characteristics from ORBIS to patent applications in REGPAT, I use the OECD’s *Harmonised Applicant Names* (HAN) dataset as a bridge. HAN standardizes applicant strings across patent records and consolidates name variants into persistent group identifiers (HAN_id), explicitly designed to facilitate linkages to business registers such as ORBIS (OECD Directorate for Science, Technology and Innovation, 2024a; Dernis and Squicciarini, 2021). Each REGPAT application (appln_id) is first associated with a HAN_id, which I then match to ORBIS firm identifiers (BvDID).

I construct this HAN_id–BvDID crosswalk following a precision-first approach that prioritizes exact string matches within country blocks and admits fuzzy matches only under conservative thresholds. The procedure enforces a unique one-to-one mapping and yields approximately 240,000 applicant–firm links across European countries. Roughly half arise from strict exact matches, with the remainder from high-confidence fuzzy pairs after legal-form normalization. Complete details on preprocessing rules, matching phases, thresholds, and validation diagnostics are provided in Appendix A2.2.

Final sample coverage. The analysis sample spans 2000–2020 and covers nine EU countries: Austria, Belgium, Denmark, Finland, France, Germany, Italy, Luxembourg, and Spain.⁶ The firm-level analysis focuses on heat-exposed sectors (NACE B, C, D, F), yielding approximately 8 million firm-year observations. The region–industry innovation analysis contains 89,068 NUTS3 $\times$ NACE-4 $\times$ year cells.

3 Descriptive Evidence of Labor-Biased Shock

This section documents the factor-biased nature of temperature shocks through reduced-form evidence on productivity and market-share responses. Section 3.1 shows that extreme heat disproportionately harms labor-intensive firms’ productivity, while Section 3.2 demonstrates that these firms consequently lose market share to capital-intensive competitors within product markets.

Empirical Framework. All specifications in Section 3 adopt a common "bin regression" template. I partition the distribution of daily *maximum* temperatures into eight 10-degree-Fahrenheit bins:

$$B = \{<30, 30\text{--}40, 40\text{--}50, 50\text{--}60, 60\text{--}70, 70\text{--}80, 80\text{--}90, \geq 90\}.$$

The interval $[60, 70)^\circ F$ is omitted as the reference category, chosen to represent mild, productivity-optimal conditions. For firm i in year t , let Bin_{it}^b denote the count of days with maximum temperature in bin b . The baseline estimating equation is

$$y_{it} = \sum_{\substack{b \in B \\ b \neq [60-70)^\circ F}} \beta_b \text{Bin}_{it}^b + \gamma \text{Prec}_{it} + \delta_i + \xi_t + \varepsilon_{it}, \quad (1)$$

where y_{it} is the outcome of interest (log productivity, market share, etc.), Prec_{it} is annual precipitation, δ_i are firm fixed effects, and ξ_t are year fixed effects. Standard errors are two-way clustered by firm and

⁶Country coverage is determined by data availability after merging firm-level (ORBIS), climate (NOAA), and patent (REGPAT-HAN) data sources. The nine countries represent all those with sufficient matched observations; this is the maximal feasible sample given data constraints, not an ad hoc selection.

country-year to allow for arbitrary within-firm serial correlation and spatial correlation in weather shocks across all firms within a country in a given year. Where appropriate, I include country-year and sector-year effects; see the details in each subsection.

The coefficients β_b measure semi-elasticities: the percentage-point change in y_{it} from one additional day in bin b relative to the reference bin $[60, 70)^\circ F$. To contextualize exposure levels: in the full sample, firms experience an average of 14.7 days above $90^\circ F$ annually, with the interquartile range spanning 1 to 20 days (see Appendix Table A1 for the complete distribution across all temperature bins). Year-to-year variation within firms is also considerable, with a standard deviation of 19.5 days in the $> 90^\circ F$ bin. The fixed effects structure isolates within-firm deviations from typical exposure patterns, so a “10 extra days” coefficient corresponds to roughly half a standard deviation increase in extreme heat relative to a firm’s historical average.

Each subsection specifies the outcome variable y_{it} and any deviations from (1); otherwise the framework is unchanged. This standardization enables transparent comparison of temperature effects across margins.

3.1 Heterogeneous Productivity Effects

I first examine whether the temperature-productivity relationship varies systematically with firms’ factor intensity. The outcome is log total factor productivity (TFP), estimated from production function estimation. Firms are ranked within country-industry-year by labor intensity⁷ and split into quartiles. For clarity of exposition, I compare the *High* labor intensity group (top quartile) to the *Low* labor intensity group (bottom three quartiles) and estimate (1) separately for each group.

The specification augments the baseline with higher-dimensional fixed effects:

$$\ln(\text{TFP})_{it} = \sum_{\substack{b \in B \\ b \neq [60-70)^\circ F}} \beta_{bg} \text{Bin}_{it}^b + \gamma_g \text{Prec}_{it} + \delta_i + \eta_{c(i),t} + \zeta_{s(i),t} + \varepsilon_{it}, \quad (2)$$

where $g \in \{\text{High}, \text{Low}\}$ indexes the labor-intensity group, $\eta_{c(i),t}$ are country-year fixed effects, and $\zeta_{s(i),t}$ are sector-year fixed effects. These high-dimensional controls absorb common macroeconomic shocks and sector-specific trends, isolating within-firm temperature variation. Standard errors remain clustered by firm and country-year.

Figure 4 reveals striking heterogeneity in heat sensitivity across firms. Panel (a) shows that the full-sample productivity effect is negative but modest in the hottest bin. Panel (b) demonstrates that this aggregate pattern masks substantial variation by factor intensity. Labor-intensive firms (High group) experience sharp productivity losses under extreme heat, while capital-intensive firms (Low group) exhibit only mild effects. The divergence is most pronounced in the hottest bins: ten additional days above $90^\circ F$ reduce log TFP by **0.051%** (s.e. = 0.013%) for the High labor-intensity group, compared to **0.022%** (s.e. = 0.012%) for the Low group—a difference of 0.029 percentage points that is both economically and statistically significant. The $[80, 90)^\circ F$ bin shows a similar pattern: High firms lose 0.041% while Low firms lose only 0.012%. Table 6 reports the full set of coefficients, confirming that heat systematically penalizes labor-intensive production.

Interpretation. The differential productivity response establishes that extreme heat operates as a *labor-biased productivity shock*. Firms whose production relies more heavily on labor—where the heat-vulnerable

⁷measured by output elasticity with respect to labor




Figure 4: Effect of daily maximum temperature bins on log TFP

Notes: Coefficients β_{bg} (scaled $\times 100$) measure the effect of one additional day in temperature bin b relative to $[60, 70]^{\circ}F$; ribbons show 95% CIs. Specification (2) includes firm, country-year, and sector-year fixed effects. Groups are defined by within-country-industry-year quartile ranks of labor intensity: *High* = most labor-intensive quartile, *Low* = pooled bottom three quartiles. Standard errors clustered by firm and country-year.

human input constitutes a larger share of the production function—suffer disproportionately larger TFP losses when temperatures rise. This heterogeneity is precisely what one would expect if heat primarily impairs labor effectiveness (through reduced worker productivity, increased fatigue, or absenteeism) while leaving capital largely unaffected. The pattern provides the foundation for understanding subsequent adaptation responses: if heat damages are concentrated on the labor margin, profit-maximizing firms have clear incentives to substitute toward capital-intensive production methods and to direct innovation toward labor-saving technologies.

Robustness. The heterogeneous productivity response is robust to alternative specifications and temperature measures. First, I estimate a pooled specification interacting temperature bins with continuous labor intensity:

$$\ln(\text{TFP})_{it} = \sum_{\substack{b \in B \\ b \neq [60-70]^{\circ}F}} [\beta_b \text{Bin}_{it}^b + \theta_b (\text{Bin}_{it}^b \times \text{LI}_{it})] + \phi \text{LI}_{it} + \delta_i + \eta_{c(i),t} + \zeta_{s(i),t} + \varepsilon_{it}. \quad (3)$$

The interaction coefficients θ_b are negative and statistically significant in the hot bins: for $[80, 90]^{\circ}F$, $\theta_b = -0.220$ (s.e. = 0.065), and for $> 90^{\circ}F$, $\theta_b = -0.111$ (s.e. = 0.067), confirming that productivity losses increase monotonically with labor intensity (Table A9). Second, replacing temperature bins with cooling degree-days (CDD) yields qualitatively identical results: the High labor-intensity group experiences a CDD coefficient of -0.0004 (s.e. = 0.0002) compared to -0.0001 (s.e. = 0.00005) for the Low group (Table A10), confirming that the labor-biased damage pattern does not depend on the specific functional form used to measure heat exposure.

3.2 Market Share Reallocation

The productivity heterogeneity documented in Section 3.1 suggests that extreme heat should reallocate output within product markets, penalizing labor-intensive producers. I test this prediction by examining

market-share responses. Markets are defined as country $\times$ four-digit NACE industry (country $\times$ NACE4). For firm i in market m and year t , log market share is

$$y_{it} = \ln(\text{MS}_{it}) = \ln\left(\frac{\text{Sales}_{it}}{\sum_{k \in (m,t)} \text{Sales}_{kt}}\right).$$


⁸ For exposition, I isolate the most labor-intensive quartile as *High* labor intensity group, and pool the remaining three quartiles as *Low*.

The specification follows (1), estimated separately by group:

$$\ln(\text{MS})_{it} = \sum_{\substack{b \in B \\ b \neq [60-70]^\circ F}} \beta_{bg} \text{Bin}_{it}^b + \gamma_g \text{Prec}_{it} + \delta_i + \eta_{c(i),t} + \zeta_{s(i),t} + \varepsilon_{it}, \quad (4)$$

where $g \in \{\text{Low}, \text{High}\}$, and δ_i , $\eta_{c(i),t}$, $\zeta_{s(i),t}$ are firm, country-year, and sector-year fixed effects. Standard errors are two-way clustered by market and country-year to account for within-market correlation and common regional shocks.

Figure 5: Effect of daily maximum temperature bins on log market share



Notes: Coefficients β_{bg} (scaled $\times 100$) measure the effect of one additional day in temperature bin b relative to $[60, 70]^\circ F$; ribbons show 95% CIs. Specification (4) includes firm, country-year, and sector-year fixed effects. Groups defined by within-market-year quartile ranks: *High* = most labor-intensive quartile, *Low* = pooled bottom three quartiles. Standard errors two-way clustered by market and country-year.

Results. Figure 5 reveals a clear split mirroring the productivity findings. Log market shares decline in hot bins for more labor-intensive firms (*High*), while the less labor-intensive firm (*Low*) firms show no detectable response. The divergence is sharpest in the hottest bin: ten additional days above $90^\circ F$ reduce the *High* group's log market share by **0.4%** (s.e. $\approx 0.12\%$), whereas the *Low* group exhibits a statistically insignificant

⁸Labor-intensity ranks are recomputed annually to allow for time-varying technology.

0.0% change. The $[80, 90)^{\circ}F$ bin also shows meaningful losses for the *High* group. Full coefficient estimates appear in Table 7.

Interpretation. Combined with the productivity evidence in Section 3.1, the pattern supports a *labor-biased* damage channel. Extreme heat depresses the labor-augmenting component of productivity, imposing larger cost shocks on firms whose technologies rely more heavily on labor. These firms face tighter margins and lose competitiveness, ceding market share to capital-intensive rivals within the same product market. The reallocation operates through standard competitive forces: less-affected (capital-intensive) producers expand at the expense of more-affected (labor-intensive) ones when heat shocks strike.

Robustness. As a complementary test, I estimate a pooled specification interacting temperature bins with continuous labor intensity:

$$\ln(\text{MS})_{it} = \sum_{\substack{b \in B \\ b \neq [60-70)^{\circ}F}} [\beta_b \text{Bin}_{it}^b + \theta_b (\text{Bin}_{it}^b \times \text{LI}_{it})] + \phi \text{LI}_{it} + \delta_i + \eta_{c(i),t} + \zeta_{s(i),t} + \varepsilon_{it}. \quad (5)$$

The main effect β_b for the hottest bin is positive and statistically significant, indicating market-share gains for very capital-intensive firms; the interaction θ_b is negative and precisely estimated, confirming that losses increase monotonically with labor intensity (Table A11).

Summary. The evidence above establishes that extreme heat imposes a *labor-biased* productivity shock. Labor-intensive firms experience larger TFP losses (Section 3.1) and consequently lose market share to capital-intensive competitors (Section 3.2). This raises a natural question: **do firms adapt to this factor-biased shock?** Section 4 examines two adjustment margins that remain understudied in the climate-economy literature: (1) technology choice, and (2) innovation.

4 Adaptation to Labor-Biased Shock

Having documented the factor-biased nature of temperature shocks, I now examine whether and how firms adjust their production technology in response. I focus on two margins of adjustment that operate at different time horizons. Section 4.1 documents *short-run factor intensity responses*—immediate reoptimization of the capital-labor mix when heat makes labor relatively less effective. Section 4.2 examines *long-run directed innovation*—whether extreme heat induces persistent R&D investment in labor-saving technologies, revealing technological change that extends beyond contemporaneous input substitution.

4.1 Short-Run Adjustment: Factor Intensity


The productivity and market-share patterns in Section 3 suggest that extreme heat impairs the labor margin more than the capital margin. If firms recognize this asymmetry, they should respond by adjusting factor intensity—substituting toward the relatively less-affected input. I test this prediction by estimating the temperature sensitivity of firm-level output elasticities with respect to capital and labor.

Measurement and Specification. Let OE_{it}^m denote firm i 's *log* output elasticity with respect to factor $m \in \{K, L\}$ in year t , estimated via production-function estimation (see Appendix A3 for details). These elasticities capture the share of output attributable to each factor and are the key technology parameters governing how firms combine inputs. I estimate the baseline bin specification (1) separately for capital and labor elasticities:

$$OE_{it}^m = \sum_{\substack{b \in B \\ b \neq [60-70]^\circ F}} \beta_b^m \text{Bin}_{it}^b + \gamma^m \text{Prec}_{it} + \delta_i^m + \eta_{c(i),t}^m + \zeta_{s(i),t}^m + \varepsilon_{it}^m, \quad (6)$$

where δ_i^m , $\eta_{c(i),t}^m$, and $\zeta_{s(i),t}^m$ are firm, country $\times$ year, and sector $\times$ year fixed effects. Standard errors are clustered by firm and country-year. Coefficients β_b^m (scaled $\times 100$) measure semi-elasticities: the percentage-point change in log output elasticity from one additional day in bin b relative to the $[60, 70)^\circ F$ reference.

Figure 6: Short-run factor intensity adjustment: effect of temperature on output elasticities



Notes: Coefficients β_b^m (scaled $\times 100$) from specification (6), estimated separately for capital (red) and labor (blue) output elasticities. Ribbons show 95% confidence intervals. Specification includes firm, country $\times$ year, and sector $\times$ year fixed effects; standard errors clustered by firm and country-year. The omitted reference bin is $[60, 70)^\circ F$.

Results. Figure 6 reveals clear factor-specific responses. Relative to the $[60, 70)^\circ F$ reference, additional very hot days ($\geq 90^\circ F$) are associated with a *rise* in capital elasticity and a *decline* in labor elasticity. Quantitatively, ten extra days above $90^\circ F$ raise the log capital output elasticity by **0.46%** (s.e. = 0.13%) and reduce the log labor elasticity by **0.07%** (s.e. = 0.03%).⁹ Given that firms experience an average of 14.7 days $\geq 90^\circ F$ annually (SD = 19.5 days; see Table A1), a one-standard-deviation increase in extreme heat implies a capital elasticity increase of 0.90% and labor elasticity decline of 0.14%. The divergence is sharpest in the high-temperature tail, consistent with nonlinear heat damage to the labor margin. The pattern

⁹Full coefficient estimates for all temperature bins appear in Table 8.

indicates *capital-using, labor-saving* technological adjustment: firms shift the input mix toward the factor that is relatively less heat-sensitive, rather than uniformly scaling down production or holding technology fixed.

Interpretation and Mechanisms. The elasticity responses reflect *input reallocation* along the production function. Under the translog specification, output elasticities are:¹⁰

$$\theta_{it}^L = \beta_l + 2\beta_{ll}l_{it} + \beta_{lk}k_{it}, \quad (7)$$

$$\theta_{it}^K = \beta_k + 2\beta_{kk}k_{it} + \beta_{lk}l_{it}. \quad (8)$$

The structural parameters $(\beta_l, \beta_k, \beta_{ll}, \beta_{kk}, \beta_{lk})$ are estimated separately for each country-industry pair over five-year windows. Within each window, these parameters remain constant, so year-to-year elasticity variation arises primarily from changes in factor inputs (l_{it}, k_{it}) rather than shifts in underlying technology. Averaging across all country-industry estimations yields $\beta_{ll} = 0.068$, $\beta_{kk} = 0.018$, and $\beta_{lk} = -0.037$.¹¹ Given these parameter values, the observed elasticity pattern follows mechanically from input adjustment: when extreme heat reduces l_{it} and raises k_{it} , labor elasticity falls (via both the $2\beta_{ll}l$ and $\beta_{lk}k$ terms) while capital elasticity rises (via $2\beta_{kk}k$ and $\beta_{lk}l$).

Direct evidence confirms this input reallocation. Figure 7 shows the temperature responses of *each* input separately. Relative to the $[60, 70)^\circ F$ reference bin, hot days—especially $\geq 90^\circ F$ —are associated with an *increase* in capital input and a *decrease* in labor input. The full coefficients are documented in Table 10. In the data, K is proxied by the log book value of fixed assets and L by the log wage bill. These opposite movements mechanically imply a higher capital–labor ratio on hot days. Taken together with the elasticity results, the input patterns point to *short-run capital deepening*: firms substitute toward capital and away from labor when heat reduces labor effectiveness, rather than proportionally scaling both inputs.

The factor intensity responses I document are consistent with growing evidence on temperature effects in manufacturing. Somanathan et al. (2021) show that extreme heat reduces labor productivity and increases absenteeism in Indian plants, with annual plant output falling by about 2% per degree Celsius—a response "driven by a reduction in the output elasticity of labor." Adhvaryu et al. (2020) find a negative, nonlinear productivity-temperature gradient in Indian garment factories, with productivity losses concentrated on hot days. Colmer (2021) documents labor reallocation away from heat-exposed manufacturing tasks during hot months in India.


My elasticity estimates provide a distinct perspective on these adjustments. Output elasticities θ^L and θ^K measure the *marginal contribution* of each factor to production—specifically, the percentage increase in output from a one-percent increase in labor or capital. This differs fundamentally from measuring input quantities alone: a firm might maintain constant employment (L unchanged) yet see labor’s productivity contribution decline sharply if heat impairs worker effectiveness, causing θ^L to fall.

When heat reduces labor productivity—as documented in the studies above—firms respond by adjusting their input mix. Equations (7)–(8) reveal that this reallocation mechanically alters the marginal product of each factor: reducing l_{it} and raising k_{it} lowers θ^L and raises θ^K given the estimated parameter values. The elasticity framework thus captures *technological adjustment* in the sense that it measures changes in how productively firms deploy each input, not merely changes in input levels. A firm experiencing heat shocks

¹⁰See Appendix A3 for derivation.

¹¹See Appendix Table A5 for the full distribution of estimated structural parameters across country-industry pairs.

Figure 7: Temperature effects on capital and labor inputs



Notes: Points plot coefficients from bin regressions of log inputs on daily maximum temperature bins, omitting $[60, 70)^{\circ}F$. Capital (K) is measured as the log book value of fixed assets; labor (L) as the log wage bill (cost). Shaded areas show 95% confidence intervals. Specifications include firm, country $\times$ year, and industry $\times$ year fixed effects; standard errors clustered at the firm and country-year level. Coefficients are scaled $\times 100$ and can be read as percentage-point changes in log inputs per additional day in each bin relative to the reference.


might use similar quantities of capital and labor yet operate at a different point on its production function where capital's marginal contribution is higher and labor's is lower. Zhang et al. (2018) establish labor-biased productivity losses in Chinese manufacturing; my results show that this bias manifests in measurable shifts in factor productivity contributions.

Industry-Level Amplification. To assess whether market-share reallocation amplifies the firm-level responses, I aggregate output elasticities to the industry level using sales-weighted averages (see Appendix A5.3.2 for construction details). By construction, the industry average $\overline{OE}_{Ict}^m$ embeds both within-firm technology adjustment and between-firm reallocation via changing market shares. Figure 8 shows that industry-level capital elasticity responses substantially exceed firm-level magnitudes: the coefficient on the hottest bin approaches 1.0 at the industry level, compared to 0.05 at the firm level¹². This amplification is consistent with market-share reallocation toward more capital-intensive producers when heat shocks strike, as documented in Section 3.2. The pattern suggests that both margins—*within-firm substitution* and *between-firm reallocation*—operate in the same labor-saving direction, with competitive forces magnifying the aggregate technological response.¹³

¹²Table 9 shows the exact coefficients.

¹³An Olley–Pakes decomposition would formally quantify the within versus between components; I relegate detailed industry heterogeneity and decomposition exercises to Appendix A5.3.2.

Figure 8: Industry-level amplification: temperature effects on average capital elasticity



Notes: Coefficients from bin regression of log industry-average capital elasticity (sales-weighted) on temperature bins, omitting $[60, 70]^{\circ}\text{F}$. Specification includes country $\times$ year and sector $\times$ year fixed effects; standard errors clustered by country-industry and year. Industry aggregation follows (A20). The larger industry-level coefficients relative to firm-level estimates (Figure 6) suggest that between-firm reallocation amplifies within-firm adjustment. See Appendix A5.3.2 for additional heterogeneity by industry labor-intensity.

Robustness and Dynamics. The main findings are robust to alternative temperature measures. Replacing daily maximum temperature bins with cooling degree-days (base 29.4°C / 85°F) yields the same directional pattern: positive coefficients for capital elasticity, negative for labor (Appendix Table A12). To assess persistence, I include two annual lags of degree-days; the cumulative response (sum of contemporaneous and lagged coefficients) remains positive for capital and negative for labor (Appendix Table ??), indicating that factor intensity adjustments are not purely transitory.

4.2 Long-Run Response: Directed Innovation

Section 4.1 documented short-run factor intensity adjustment—firms substitute toward capital when heat makes labor relatively less effective. But do temperature shocks also induce *persistent* technological change through directed R&D? Patents offer a window into long-run innovation responses that extend beyond contemporaneous input reoptimization. I examine whether extreme heat tilts the direction of innovation toward labor-saving technologies, particularly in labor-intensive industries where the incentive to automate should be strongest.

Empirical Design. The analysis exploits two-dimensional variation in heat exposure: spatial heterogeneity across NUTS-3 regions (temperature shocks vary geographically) interacted with cross-industry differences in labor intensity. The identification strategy asks whether labor-intensive industries located in regions experiencing extreme heat develop more labor-saving innovations, capturing directed technological change where innovation responds to the economic environment by targeting the factor whose productivity has deteriorated.

The outcome is the annual count of labor-saving (LS) patent applications at the region–industry–year level, identified using a single indicator that flags a patent as LS if either (i) any CPC/IPC code lies in the

DHOZ automation set or (ii) its title contains canonical automation keywords (Section 2.4). For temperature exposure, I use cooling degree-days above 85 °F,

$$\text{CDD}_{rt} = \sum_{d=1}^{365} \max\{0, \text{TMAX}_{rt,d} - 85^\circ\text{F}\},$$

a parsimonious scalar suited to interactions and lags.¹⁴ Industry labor intensity, $\text{LI}_{i,0}$, is the pre-sample NACE-2 industry average output elasticity w.r.t labor, a measure of ex-ante industry exposure to labor-biased heat-shock. I interact $\text{CDD}_{r,t-2}$ with $\text{LI}_{i,0}$ and use a two-year lag as the baseline.¹⁵

I estimate Poisson (PPML) count models with high-dimensional fixed effects, the standard approach for patent data.¹⁶ The baseline specification is:

$$\mathbb{E}[P_{rit}] = \exp\left(\beta_1 \text{CDD}_{r,t-2} + \beta_2 \text{LI}_{i,0} + \beta_3 \text{CDD}_{r,t-2} \times \text{LI}_{i,0} + \theta_{r,s(i)} + \nu_t\right), \quad (9)$$

where P_{rit} is the LS patent count, $\theta_{r,s(i)}$ are region-sector fixed effects, ν_t are priority-year fixed effects, and standard errors are two-way clustered by region–year and industry to allow for spatial-temporal correlation in temperature shocks and common innovation trends across regions within the same industry. The coefficient of interest is β_3 , capturing whether heat exposure induces more LS innovation in labor-intensive industries.

Main Results. Table 3 reveals directed innovation responses to heat exposure. The $\text{CDD} \times \text{LI}$ interaction coefficient is $\hat{\beta}_3 = 0.0144$ (s.e. = 0.0034), statistically significant at the 1% level. The main effect of CDD is negative ($\hat{\beta}_1 = -0.0069$, s.e. = 0.0025), but the interaction dominates: for industries with labor intensity above approximately 0.48, higher heat exposure increases labor-saving patenting.¹⁷ The pattern indicates that innovation responds to the economic environment—regions experiencing more extreme heat shift R&D toward labor-saving technologies, with the response concentrated in labor-intensive industries where the returns to automation are highest.

To interpret the effect size, consider Spain’s textile industry (NACE 13), which has high labor intensity ($\text{LI} \approx 0.72$) and experienced substantial warming between 2000 and 2020. During this period, some Spanish regions saw CDD increases exceeding 500 degree-days—roughly the 95th percentile of observed changes from 2000 to 2020 in the sample. Under this observed scenario, the model predicts that such exposure would raise labor-saving patent counts by a factor of approximately 5.7—that is, LS patents in Spanish textiles would be 470% higher than in the absence of this heat increase.

Robustness. The directed-innovation result does not hinge on the two-year lag. Figure 9 plots the interaction coefficient $\hat{\beta}_3(\text{CDD}_{t-k} \times \text{LI}_{i,0})$ estimated separately for $k = 2, \dots, 5$. Estimates are positive and statistically significant at all lags, with similar magnitudes (≈ 0.012 – 0.016). Thus, LS patenting at time t responds to a *history* of heat exposure rather than a one-off shock, consistent with R&D gestation and investment thresholds that require a persistent signal. Appendix Table A15 reports the underlying regressions and alternative specifications.

¹⁴Results are similar with temperature bins.

¹⁵As is standard in the innovation literature, patent applications typically reflect R&D initiated earlier (Hall et al., 1986; Popp, 2002; Aghion et al., 2016). The full lag profile and alternatives $k = 2$ – 5 are reported in Appendix Table A15.

¹⁶Suitable for counts with many zeros; the multiplicative mean specification is robust to overdispersion and accommodates rich fixed effects.


¹⁷This threshold represents the first quartile across NACE-2 industries in the sample. See Appendix Table ?? for the summary labor intensity across industries.

Table 3: Heat exposure and labor-saving innovation: region–industry Poisson estimates

	Labor-saving patents (1)
CDD	-0.0069*** (0.0016)
Labor Intensity	-2.797*** (0.1440)
CDD $\times$ Labor Intensity	0.0144*** (0.0034)
Observations	76,400
Region $\times$ Sector FE	✓
Year FE	✓

Notes: Poisson estimates at NUTS-3 region–industry–year level. Temperature exposure measured as cooling degree-days above 85 °F, lagged two years. Specification (9) includes region-sector and priority-year fixed effects. Standard errors clustered by region-year and industry. Coefficients interpretable as semi-elasticities. *** p<0.01, ** p<0.05, * p<0.1.

Figure 9: Dynamic innovation response across lags




Notes: Points show PPML estimates of $\hat{\beta}_3$ on $CDD_{t-k} \times LI_{i,0}$ for $k = 2, \dots, 5$; bars are 95% confidence intervals. The baseline $k = 2$ estimate is highlighted for reference.

Long-difference specifications cumulating effects over 2000-2020 yield similar magnitudes: $\Delta CDD \times LI = 0.0268$ (s.e. = 0.0051), consistent with persistent innovation responses to sustained heat exposure (Appendix Table A16). Using patent shares (LS patents as a fraction of total patents) rather than counts yields comparable estimates, confirming the effect operates through increased labor-saving innovation rather than

general patenting activity (Appendix Table A14).

Implied timeline—from shock to adaptation The evidence from Section 3 and 4 delivers a simple timeline from shock to adaptation (Figure 10). Heat shocks at time t act as labor-biased productivity shocks: they bite harder in labor-intensive activities and induce within-period substitution away from the exposed factor. I observe contemporaneous shifts in output elasticities and higher capital intensity, reflecting input substitution (OE shifts, $K/L \uparrow$). The innovation margin operates in the medium run. In the dynamic Poisson estimates, the interaction coefficient $\beta_3(\text{CDD}_{t-k} \times \text{LI})$ is positive and statistically significant for $k = 2, 3, 4, 5$ (Figure 9). This pattern is consistent with directed technical change under adjustment costs and R&D gestation: the shock raises the expected return to labor-saving technologies, but investment is triggered once exposure proves persistent enough to clear fixed costs. Repeated hot years update beliefs about future exposure, making LS innovation privately optimal. In short: labor-biased shock $\rightarrow$ within-period substitution $\rightarrow$ medium-run directed innovation.

Figure 10: From heat shock to adaptation.



4.2.1 Theoretical Framework: Why Heat Redirects Innovation

The finding that heat-exposed industries patent more in labor-saving domains fits the logic of Directed Technical Change (DTC): when a factor of production becomes scarce or costly, the direction of innovation responds to the incentives created by that scarcity (Hicks, 1932; Acemoglu, 2002, 2010). In our setting, heat shocks reduce *effective* labor supply, raising the effective cost of labor (i.e., the shadow price of $D(T)A_L L$ increases when $D(T) \downarrow$).

Two forces shape the direction of innovation in DTC. *First (price effect)*, a higher price of a factor raises the payoff to *augmenting that same factor*; when labor becomes more expensive in efficiency units, the price effect points toward *labor-augmenting* innovation. *Second (market-size effect)*, innovation profits are larger in the market where the quantity used is larger; because heat depresses effective labor, the market for labor-augmenting technologies shrinks relative to the market for capital-augmenting (labor-saving) technologies, tilting incentives toward *labor-saving* innovation (Acemoglu, 2002).¹⁸ Theoretically, the strength and net sign of these forces depend on the elasticity of substitution between capital and labor: with $\sigma > 1$ the market-size force tends to dominate (favoring labor-saving), whereas with $\sigma < 1$ the price effect can dominate (favoring labor-augmenting). Empirical estimates of σ vary—from about 0.3–0.8 in U.S./European manufacturing to 3–4 in Chinese manufacturing (Raval, 2019; Oberfield and Raval, 2021; Liu et al., 2020). Using my translog production function estimates across country-industry pairs, I find suggestive evidence consistent with increasing capital-labor substitutability over the 2000–2019 period: the interaction parameter governing

¹⁸In Appendix A4, I develop a stylized DTC model following Acemoglu (2003) and Acemoglu et al. (2012) and show conditions under which a labor-biased shock (a fall in $D(T)$) leads to labor-saving innovation. The net direction reflects the interaction of the price and market-size forces and depends on the elasticity of substitution, σ .

substitutability shows a clear upward trend (See Appendix Section A3 for more details).¹⁹

Recent empirical work in DTC documents analogous links between factor scarcity/abundance and the direction of innovation: Acemoglu and Restrepo (2022) show demographic aging increases robot adoption across countries, and Moscona and Sastry (2023) find drought exposure induces directed innovation toward drought-resistant crops. My setting—heat shocks that shrink effective labor supply—provides novel variation in effective labor costs, and the innovation response fits this broader pattern.

But so what? Does induced innovation matter for productivity? Through the DTC lens, heat affects productivity through two channels: (i) a *direct* impact, as heat lowers production efficiency via the labor damage function; and (ii) an *indirect* impact, as induced innovation shifts technology toward labor-saving (or factor-augmenting) directions. Formally, with a factor-augmenting representation $Y = F(A_K K, D(T) A_L L)$, a first-order log approximation yields that changes in $D(T)$ and factor-specific technologies (A_K, A_L) contribute to $\Delta \ln Y$ with weights related to factor shares and σ . For exposition, I summarize the net effect as

$$\Delta \ln \text{Productivity} \approx \underbrace{\Delta \ln D(T)}_{\text{Direct damage} < 0} + \underbrace{\Delta \ln(A)}_{\text{Innovation response}},$$

where the “ A ” term captures the composition of factor-augmenting improvements induced by heat (notably, capital-augmenting/labor-saving in environments with higher substitutability). Whether the innovation response quantitatively mitigates the direct damage is an empirical question that I address in Section 5: does induced innovation attenuate the heat-to-productivity gradient?

5 Quantifying Mitigation: Does Innovation Attenuate Heat Damages?

Building on the evidence that extreme heat is a labor-biased shock (Section 3) and systematically redirects inventive effort toward Labor-Saving (LS) technologies (Section 4), this section quantifies whether that redirected innovation functions as a tool for *adaptation*. Specifically, I test whether firms that recently innovated along the LS margin experience a weaker heat–productivity gradient. To answer this, I link firm–patent matches to a firm-year productivity panel and estimate interaction regressions in which exposure to cooling degree–days (CDD) is interacted with lagged LS measures. The parameter of interest is the CDD×LS interaction coefficient: a positive estimate indicates that LS innovation attenuates the *marginal* productivity damage of heat by shifting firms onto a technology frontier less sensitive to heat-induced labor inefficiency.

Design and identification. I take total factor productivity (TFP) in logs at the firm-year level as the outcome. Temperature exposure is measured using cooling degree–days above 85°F, denoted as CDD, following the definition in Section 2.1. Labor-saving (LS) innovation is proxied by a firm-year indicator `ls_any`, which equals one if the firm has filed any LS patent in the preceding period.²⁰ I estimate the following firm-level

¹⁹This rising substitutability is qualitatively consistent with the recent wave of automation and robot adoption—capital that substitutes for routine and some non-routine tasks—suggesting increasing substitutability at the margin in economies undergoing automation (e.g., Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020; Karabarbounis and Neiman, 2014).

²⁰Patent outcomes are linked to ORBIS firms via the HAN–ORBIS name matching procedure described in Section 2.5.

specification:

$$\ln \text{TFP}_{iy} = \beta_1 \text{CDD}_{iy} + \beta_2 \text{LS}_{i,y-k} + \beta_3 \text{CDD}_{iy} \times \text{LS}_{i,y-k} + \Gamma \text{Prec}_{iy} + \alpha_i + \alpha_{\text{iso}(i) \times y} + \alpha_{\text{sector}(i) \times y} + \varepsilon_{iy}, \quad (10)$$

where α_i are firm fixed effects, $\alpha_{\text{iso}(i) \times y}$ are country–year effects, and $\alpha_{\text{sector}(i) \times y}$ are sector–year effects. Standard errors are clustered by firm and country–year.

Identification comes from within-firm variation in CDD exposure combined with cross-sectional heterogeneity in firms’ innovative activity. The key parameter of interest is the interaction coefficient β_3 , which tests whether the marginal productivity loss from extreme heat is mitigated for firms with recent LS innovation. Because the adoption and diffusion of new technologies take time to translate into measurable productivity gains (Brynjolfsson and Hitt, 2003; Comin and Hobijn, 2010), I introduce a lag between LS innovation and productivity. For the baseline specification, I focus on a two-year lag ($k = 2$) as the benchmark.²¹

Firm-level evidence. Table 4 reports the benchmark specification that interacts cumulative degree days (CDD) with an indicator for recent labor-saving (LS) innovation, measured as having filed at least one LS patent in the previous two years ($k = 2$). For readability, all coefficients are rescaled so that they can be interpreted directly in percentage terms. The baseline effect of temperature (column 1, β_1) is negative: additional CDD lowers firm-level productivity, consistent with the expectation that extreme heat depresses output. By contrast, the interaction coefficient $\hat{\beta}_3$ is positive and statistically significant, indicating that firms with recent labor-saving innovation experience attenuated—and potentially reversed—productivity losses under heat shocks.

To gauge economic magnitudes, I benchmark the estimates against a realistic climate scenario of an additional 100 CDD.²² A shock of this size lowers productivity by about 0.05 percent for firms without recent LS innovation, based on the baseline coefficient β_1 in Table 4. In sharp contrast, for firms with recent LS patents, the positive interaction term β_3 more than offsets this loss: the net effect of 100 additional CDD is an *increase* in productivity of about 0.56 percent ($0.61 - 0.05$).

This striking reversal—from productivity loss to productivity gain—warrants careful interpretation. The result does *not* imply that heat makes innovating firms more productive in an absolute sense. Rather, it reflects three complementary mechanisms. First, LS innovation allows firms to substitute away from heat-sensitive labor toward more climate-resilient capital and automation, directly reducing exposure to the heat shock. Second, the timing is important: firms that successfully patent LS technologies two years prior are likely those that also invested in complementary organizational changes and worker training, amplifying the measured TFP effect. Third, there may be selection: firms that innovate in response to heat may be better-managed or have higher baseline adaptive capacity, though the firm fixed effects α_i absorb time-invariant differences in management quality.

Taken together, the evidence shows that firms innovating along the labor-saving margin not only offset climate-induced damage but experience net productivity gains under extreme heat. This provides direct

²¹Alternative lags $k \in \{1, 3\}$ yield qualitatively similar patterns, whereas contemporaneous innovation ($k = 0$) shows no significant effect on productivity, consistent with the view that technology adoption requires time. Figure 11 presents the full lag profile.

²²Appendix Table A3 documents that between 2000 and 2020, the 75th percentile of country-level increases in CDD is about 100 degree-days.

Table 4: Firm-level mitigation: TFP on CDD $\times$ LS indicator (benchmark)

	log TFP (1)
CDD	-0.0005** (0.0002)
Any LS (t-k)	-0.1638 (0.5309)
CDD $\times$ Any LS (t-k)	0.0061*** (0.0021)
Observations	6,096,992
Firm FE	✓
Country $\times$ Year FE	✓
Sector $\times$ Year FE	✓
Lag Revenue Control	✓

Notes: Outcome: log TFP . Temperature exposure: CDD defined as cumulative degree days above 85 °F . LS indicator equals one if the firm filed at least one labor-saving patent in the preceding two years. All coefficients expressed in percentage terms. Fixed effects: firm, country $\times$ year, and sector $\times$ year. Standard errors clustered by firm and country-year. *** p<0.01, ** p<0.05, * p<0.1.

micro-level evidence that directed innovation can mitigate—and even locally reverse—the adverse productivity impacts of climate change, at least for firms with the capacity and resources to innovate.

Robustness. Whether LS innovation is measured by a binary indicator or by the continuous count of patents produces the same conclusion: firms with labor-saving patents are significantly less exposed to climate damages, with some evidence of a dose-response relationship (higher LS stocks yield stronger mitigation).²³


Timing and dynamics of mitigation. To examine the temporal profile of the mitigation effect and rule out contemporaneous confounding, I estimate equation (10) across multiple lag specifications ($k \in \{0, 1, 2, 3\}$), holding the sample constant. Figure 11 plots the estimated interaction coefficient $\hat{\beta}_3$ for each lag k . The pattern reveals three key features that illuminate the shock $\rightarrow$ innovation $\rightarrow$ mitigation sequence.

First, the contemporaneous effect ($k = 0$) is statistically indistinguishable from zero. This rules out reverse causality where heat shocks simultaneously depress productivity *and* trigger LS patenting in the same year, which would spuriously generate a positive interaction even absent any true mitigation. The null contemporaneous effect confirms that patenting itself is not an immediate response to within-year productivity fluctuations.

Second, the mitigation effect *emerges* at a one-year lag ($k = 1$) and *peaks* at the two-year lag ($k = 2$, the baseline specification). This timing is consistent with the well-documented adoption and diffusion lags in technology implementation (Brynjolfsson and Hitt, 2003; Comin and Hobijn, 2010): patents filed in

²³See Appendix Table A17 for patent count specifications.

Figure 11: Mitigation effect across innovation lags: β_3 profile



Notes: Each point represents the estimated interaction coefficient β_3 from equation (10) using a different lag k for the LS indicator. The baseline specification ($k = 2$, in blue) corresponds to Table 4. Bars show 95% confidence intervals. Standard errors clustered by firm and country-year. Coefficients expressed in percentage points per additional CDD.

year $t - 2$ represent inventions that subsequently require time for prototyping, process integration, worker retraining, and organizational adjustment before measurable productivity gains materialize. The magnitude at $k = 2$ ($\hat{\beta}_3 \approx 0.006$, or 0.6 percentage points per 100 CDD) captures the full mitigation potential once these complementary investments are in place.

Third, the effect *persists* at longer lags ($k = 3$), though with slightly wider confidence intervals due to reduced sample overlap. This persistence indicates that LS innovation does not deliver a transitory productivity spike but rather shifts firms onto a durably flatter heat-damage trajectory, consistent with the interpretation that these technologies fundamentally alter the firm's production frontier.

Taken together, the lag profile provides strong evidence that the mitigation channel operates through *ex post* technology adoption, not *ex ante* firm selection or contemporaneous shocks. The sequence—heat exposure (year $t - 2$) $\rightarrow$ LS innovation (year $t - 2$) $\rightarrow$ adoption/diffusion (years $t - 1$ to t) $\rightarrow$ mitigation (year t)—is precisely what one would expect if directed innovation serves as an adaptive response to climate risk. Importantly, this timing also differentiates the mitigation mechanism from confounds where high-TFP firms simply innovate more: such selection would generate positive interactions at *all* lags, including $k = 0$, which I do not observe.

Region–industry counterpart. To examine if the mitigation result applies to broader level, I estimate the region–industry specification using the *LS patent count* as the regressor. Within each NUTS3 $\times$ industry cell and year, I form sales-weighted means of log TFP and CDD using each firm's *sample-average* sales as a fixed weight (so composition is held constant over time). Let $C_{r,j,t}$ denote the number of LS patent applications filed by firms in industry-region cell (r, j) in year t . I interact contemporaneous heat with the

lagged LS count ($k=2$):

$$\log \text{TFP}_{rjt} = \beta_1 \text{CDD}_{rt} + \beta_3 (\text{CDD}_{rt} \times C_{rj,t-2}) + \theta_{rj} + \nu_t + \varepsilon_{rjt}, \quad (11)$$

estimated with region $\times$ industry and year fixed effects and two-way clustering by region and year.²⁴ Intuitively, the regression asks whether region-industries that produced *more* LS patents two years earlier experience a *smaller* decline in TFP for the same heat shock, holding cell and year effects fixed. The estimates replicate the micro pattern: the baseline temperature effect is negative, while the interaction term is positive and statistically significant, implying that cells with more recent LS innovation suffer attenuated productivity losses under the same heat shock. This count specification is the one used in the BOTE, ensuring that the object predicted in the innovation stage maps directly into the mitigation stage.

Table 5: Region-industry mitigation: TFP on CDD $\times$ LS Count

	log TFP (1)
CDD	-0.0004*** (0.0001)
LS count (t-k)	-0.5731 (1.093)
CDD $\times$ LS count (t-k)	0.0147*** (0.0011)
Observations	71,800
Region $\times$ Industry FE	✓
Year FE	✓

Notes: Unit: NUTS3 region $\times$ industry (cell). Outcome: log TFP. Temperature: cooling degree days above 85°F. LS regressor: the cell's *lagged two-year* count of labor-saving patent applications (single-year flow, aggregated across firms). Aggregation/weights: cell-year means are computed using *sample-average* firm sales as fixed weights within the cell (composition held constant over time); weighted specification shown. Fixed effects: region $\times$ industry and year. Standard errors: two-way clustered by region and year. Coefficients are reported in percent per +1 CDD (scaling matches the firm table). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Summary. At the firm level, recent LS patenting attenuates the heat-TFP gradient. Importantly, the *same qualitative mitigation* appears at the region-industry level, indicating that region-industry with more recently accumulated LS innovation suffer smaller heat-related productivity losses. Conceptually, the empirical pattern aligns with the mechanism implied by the DTC framework in Section 4.2.1: innovations act as a buffer to the labor-biased heat shock.

²⁴A two-year lag reflects the typical gestation from patenting to productivity and ensures the LS regressor is predetermined relative to current heat; using a single-year count keeps the object identical to Stage 1 and avoids additional assumptions about depreciation or stock windows.

5.1 Aggregate Mitigation: Back-of-Envelope Quantification

This back-of-the-envelope (BOTE) calculation quantifies how much directed labor-saving (LS) innovation offsets aggregate productivity losses from rising heat. The approach leverages two empirical relationships established earlier in the paper. From Section 4.2, I have estimates of how heat exposure redirects innovation toward LS technologies—the *heat* $\rightarrow$ *innovation* link. From Section 5, I have estimates of how LS patents flatten the heat-TFP damage curve—the *innovation* $\rightarrow$ *mitigation* link. Combining these two relationships allows me to trace the full causal chain: heat induces LS innovation (Stage 1), and that induced innovation attenuates productivity damage (Stage 2). The calculation compares the realized path—where LS innovation responds endogenously to heat—against a counterfactual with the *same* climate trajectory but *no induced innovation response* (cf. Moscona and Sastry, 2023). This counterfactual isolates the protective value of the observed technological adjustment.

Let $\Delta\text{TFP}^{\text{I}}$ denote the TFP change in the world with induced innovation operating as observed, and $\Delta\text{TFP}^{\text{NI}}$ the change when induced LS is shut down but all else (including the climate shock) remains the same. The mitigation share—the fraction of heat-related productivity losses offset by innovation—is then

$$\text{Mitigation} := \frac{\Delta\text{TFP}^{\text{NI}} - \Delta\text{TFP}^{\text{I}}}{\Delta\text{TFP}^{\text{NI}}}. \quad (12)$$

The numerator captures the productivity gain attributable to innovation: it is the difference between what would have been lost without innovation and what was actually lost with it. The denominator is the baseline damage absent any innovation response.

To operationalize this calculation, I return to the two regressions estimated in earlier sections, now recast at the region-industry level to align the units and timing across both stages. The construction uses estimates from the *same spatial unit* (NUTS3 region $\times$ 2-digit industry), with the *same lag structure* ($k = 2$ years), and the *same LS measure* (single-year patent counts at $t - 2$). This consistency ensures that the induced innovation in Stage 1 is precisely the innovation that mitigates damages in Stage 2, making the counterfactual comparison internally coherent.

Stage 1: Heat induces LS innovation. The first stage re-estimates the heat-to-innovation relationship from Section 4.2 at the region-industry-year level. Index NUTS3 regions by r , 2-digit industries by j , and years by t . Let C_{rjt} denote the *count* of LS patent applications with priority year t in region r and industry j . I estimate a Poisson pseudo-maximum likelihood (PPML) model with region $\times$ sector fixed effects (ϕ_{rs}) and year fixed effects (τ_t), interacting lagged cooling degree-days with pre-sample labor intensity:

$$\log \mathbb{E}[C_{rjt} \mid \cdot] = \phi_{rs} + \tau_t + \alpha_1 \text{CDD}_{r,t-2} + \alpha_2 \text{LI}_{j,0} + \alpha_3 (\text{CDD}_{r,t-2} \times \text{LI}_{j,0}),$$

where $\text{LI}_{j,0}$ is industry j 's labor intensity measured in the pre-sample period.²⁵ The two **blue coefficients**— α_1 and α_3 —are the key parameters from Stage 1. They tell us *how many* additional LS patents are induced by heat: α_1 captures the average innovation response, and α_3 captures the differential response in labor-intensive industries, where heat's impact on labor productivity is most severe.

To construct the counterfactual, define the realized heat shock as

$$\Delta\text{CDD}_r := \text{CDD}_{r,2020} - \text{CDD}_{r,2000}.$$

²⁵The two-year lag reflects typical R&D gestation lags and ensures that the LS patent stock in Stage 2 is predetermined relative to contemporaneous heat exposure.

Using the estimated coefficients, I predict LS patent counts under two scenarios. Under the *realized* climate (setting $CDD_{r,t-2}$ to its 2020 value and $LI_{j,0}$ to its observed level), I obtain the induced count $\widehat{C}_{rj,t-2}^I$. Under the *baseline* climate (holding $CDD_{r,t-2}$ at its 2000 level while keeping $LI_{j,0}$ constant), I obtain the no-innovation counterfactual count $\widehat{C}_{rj,t-2}^{NI}$. The heat-induced change in LS patent counts is then

$$\widehat{\Delta C}_{rj,t-2} := \widehat{C}_{rj,t-2}^I - \widehat{C}_{rj,t-2}^{NI} \quad (\text{approximately proportional to } [\alpha_1 + \alpha_3 LI_{j,0}] \Delta CDD_r).$$

This object quantifies the additional LS innovation attributable to the observed warming. Regions that experienced larger temperature increases (ΔCDD_r higher) and industries that are more labor-intensive ($LI_{j,0}$ higher) see the largest innovation responses, reflecting the targeted nature of the adaptive technological adjustment.

Stage 2: LS innovation flattens the damage curve. The second stage re-estimates the mitigation relationship from Section 5 at the same region-industry-year level. Using cell-year means of log TFP and CDD constructed with *sample-average* sales weights, I estimate

$$\log \text{TFP}_{rjt} = \beta_1 CDD_{rt} + \beta_2 C_{rj,t-2} + \beta_3 (CDD_{rt} \times C_{rj,t-2}) + \theta_{rj} + \nu_t + \varepsilon_{rjt},$$

where θ_{rj} are region $\times$ industry fixed effects and ν_t are year fixed effects. The two **red coefficients**— β_1 and β_3 —are the key parameters from Stage 2. They tell us *how much* LS innovation mitigates heat damage. The coefficient β_1 captures the baseline heat-TFP slope: it is the marginal effect of an additional CDD on log TFP when there are zero LS patents in the region-industry. The coefficient $\beta_3 > 0$ captures the mitigation effect: each additional LS patent filed two years earlier *flattens* the damage slope, making productivity less sensitive to heat.²⁶

The interpretation is straightforward. In a region-industry with C LS patents filed at $t-2$, the marginal effect of heat on productivity is $\beta_1 + \beta_3 \cdot C$. Since $\beta_3 > 0$, larger LS stocks make this marginal effect less negative (or even positive, if mitigation is sufficiently strong). This is the mechanism through which innovation attenuates climate damages: it does not eliminate the physical climate shock, but it reduces the economic sensitivity to that shock.

Combining the stages: From innovation to mitigation. With estimates of the four key coefficients— α_1, α_3 from Stage 1 and β_1, β_3 from Stage 2—I can now compute the mitigation gain for each region-industry cell. Let $C_{rj,t-2}^{NI}$ be the counterfactual LS count predicted from Stage 1 under baseline climate (no warming), and $C_{rj,t-2}^I$ be the realized count under observed warming. Over the climate window ΔCDD_r , the predicted TFP changes are

$$\Delta \log \text{TFP}_{rj}^{NI} \approx (\beta_1 + \beta_3 C_{rj,t-2}^{NI}) \Delta CDD_r, \quad \Delta \log \text{TFP}_{rj}^I \approx (\beta_1 + \beta_3 C_{rj,t-2}^I) \Delta CDD_r.$$

The first expression is the TFP loss that would occur if innovation remained at its baseline level (no induced response); the second is the actual TFP change with innovation responding to heat. The difference between them is the *mitigation gain*:

$$\Delta \log \text{TFP}_{rj}^I - \Delta \log \text{TFP}_{rj}^{NI} \approx \beta_3 \Delta CDD_r \underbrace{(C_{rj,t-2}^I - C_{rj,t-2}^{NI})}_{\widehat{\Delta C}_{rj,t-2} \text{ from Stage 1}}.$$


²⁶The level effect β_2 does not enter the BOTE calculation because it reflects the direct productivity contribution of LS innovation independent of heat, which does not affect the *slope* of the heat-TFP relationship.

This expression has a transparent economic interpretation. The term $\beta_3 \Delta \text{CDD}_r$ is the marginal mitigation per additional LS patent, scaled by the size of the climate shock. Multiplying by $\widehat{\Delta C}_{rj,t-2}$ —the heat-induced increase in LS patents from Stage 1—yields the total productivity protection for cell (r, j) . Intuitively, the **blue coefficients** (α_1, α_3) tell us *how many* extra LS patents heat induces in each region-industry, and the **red coefficients** (β_1, β_3) tell us *how much* those patents reduce the marginal damage of heat. The product of "how many" and "how much," scaled by the climate shock, gives the cell-level mitigation.

To obtain the aggregate mitigation share, I average the cell-level gains $\Delta \log \text{TFP}_{rj}^I - \Delta \log \text{TFP}_{rj}^{\text{NI}}$ and losses $\Delta \log \text{TFP}_{rj}^{\text{NI}}$ across region-industry cells within each country, then compute a GDP-weighted mean across the nine EU countries. This yields the headline mitigation share in equation (12): the fraction of aggregate heat-induced productivity losses offset by directed LS innovation over 2000–2020.²⁷

Results. The headline finding is that directed LS innovation mitigates approximately **26 percent** of aggregate heat-induced productivity damages over 2000–2020 (GDP-weighted across nine EU countries). Relative to a counterfactual world with the same warming trajectory but no induced LS innovation response, the observed redirection of inventive effort recovers nearly one-third of the output that would otherwise be lost to extreme heat.

Figure 12: Heat shock and share of damage mitigated by induced LS innovation



Notes: Panel (a) shows the change in cooling degree days (CDD, base 85 °F) between 2000 and 2020 across European regions. Panel (b) shows the country-level mitigation share (in percent) computed from the BOTE exercise, defined as the fraction of heat-induced TFP losses offset by induced LS innovation. Darker shading indicates higher mitigation. See text for details on construction.

This aggregate figure masks considerable heterogeneity across countries and regions. Figure 12(a) visualizes the spatial pattern of the climate shock between 2000 and 2020, while Figure 12(b) shows the corresponding country-level mitigation shares. The dispersion is substantial—mitigation ranges from roughly 0 to

²⁷The $k = 2$ timing ensures temporal ordering—LS patent counts at $t - 2$ are predetermined relative to contemporaneous heat at t —and aligns with typical gestation lags between patent filings and productivity realizations. Using patent counts rather than stocks avoids assumptions about knowledge depreciation rates. Results are robust to alternative lag structures ($k = 1$ to $k = 5$) and to alternative climate windows (e.g., 2000–2015, 2000–2023).

40 percent—reflecting differences in both warming severity and adaptive capacity. Countries with larger temperature increases and labor-intensive industrial structures experienced greater absolute heat exposure, but also stronger innovation responses where baseline R&D capacity permitted. Countries with limited warming or capital-intensive production to begin with saw smaller damages and correspondingly smaller scope for LS mitigation. The mitigation is therefore *not* a uniform cushion but a *targeted response*: innovation accrues disproportionately where the returns to adaptation are highest, yet that targeting depends critically on local innovation capacity and industrial composition.

Three forces interact to produce the aggregate 26 percent. First, spatial heterogeneity in realized warming creates differential exposure: southern European regions face larger heat shocks than northern ones. Second, the endogenous innovation response is concentrated in labor-intensive industries within hot regions—precisely where the interaction of heat exposure and labor dependence creates the strongest incentives for LS R&D. Third, the damage-flattening effect captured by β_3 translates this targeted innovation into productivity protection. The aggregate mitigation reflects the product of exposure, response, and effectiveness across all region-industry cells, weighted by economic importance. That roughly one-quarter of damages are offset through this single channel—in-house innovation by incumbent firms—suggests innovation is a quantitatively meaningful margin of adaptation, though clearly not sufficient to eliminate climate impacts.²⁸

Discussion The headline finding—that directed labor-saving innovation mitigates approximately 26 percent of aggregate heat-induced productivity damages over 2000–2020—should be interpreted as a conservative lower bound on innovation’s adaptive contribution. This figure captures only one measurable channel (in-house LS patents by incumbent firms) and omits several amplification mechanisms.

The 26 percent figure likely understates the true effect for several reasons. First, the calculation excludes technology diffusion beyond patenting firms. Many small and medium enterprises adopt automation equipment, climate-resilient processes, or labor-saving software developed elsewhere, yet leave no patent trail. Knowledge spillovers through labor mobility, supplier relationships, and imitation further spread LS technologies. Second, conditioning on firm survival ignores selection: if heat accelerates exit of vulnerable firms while enabling expansion of resilient competitors, the aggregate economy becomes more climate-resistant than incumbent-based estimates suggest. Third, the analysis abstracts from general equilibrium adjustments—LS adoption may lower output prices, shift factor returns, and reallocate market shares in ways that amplify welfare gains beyond measured TFP. Fourth, empirical scope is restricted to nine EU manufacturing economies, omitting adaptation in agriculture, services, construction, and other heat-exposed sectors. Fifth, many firms adapt through informal process improvements and purchased capital that generate no patent record.

Together, these omissions imply the true aggregate contribution could exceed 26 percent. While some LS innovation would have occurred absent climate stress, the robust empirical relationship between heat exposure and LS patenting—particularly the positive CDD×LI interaction and lagged responses—indicates climate shock is one of the important drivers. I therefore interpret 26 percent as a lower bound, with actual mitigation likely considerably higher once diffusion, spillovers, selection, and unmeasured sectors are accounted for.

The magnitude is also setting-specific. EU manufacturing operates in a high-innovation environment with strong IP protections, integrated markets, abundant capital, and institutional R&D support, facilitating rapid technological adjustment. The 26 percent observed here likely represents an upper bound for

²⁸A one-row summary table of the GDP-weighted headline is provided below.

developing economies with weaker ecosystems, limited automation capital, credit constraints, and slower diffusion. Conversely, policies that actively facilitate R&D and adoption—through targeted subsidies, public procurement, or complementary infrastructure—could yield mitigation exceeding these estimates. This heterogeneity underscores that innovation-led adaptation is not uniformly available: climate damages will be distributed not only by exposure but also by adaptive capacity.

The observed innovations emerged endogenously from profit incentives, without explicit climate mandates, revealing substantial private-sector capacity for autonomous adjustment. Yet market failures—knowledge spillovers preventing firms from capturing full social value, unpriced climate externalities—likely lead to underinvestment relative to the social optimum. The incomplete mitigation (26 percent as lower bound, yet leaving most damages unaddressed) underscores this gap. Policy interventions addressing these distortions—R&D subsidies, technology transfer programs, mechanisms that internalize climate costs—could amplify adaptive innovation beyond what markets achieve alone.

6 Conclusion

This paper establishes that directed innovation is an important margin of adaptation to climate change. Using firm-level and patent data across nine EU countries from 2000–2020, I show that extreme heat—operating as a labor-biased productivity shock—systematically redirects innovation toward labor-saving technologies, and that this endogenous response mitigates approximately 26 percent of aggregate heat-related productivity losses.

The findings carry implications for how we estimate and interpret climate damages. The climate-economy literature typically measures economic losses from temperature shocks holding technology fixed, either explicitly through short panels or implicitly by omitting innovation responses. This simplifies estimation but may introduce systematic bias if technology adapts meaningfully over policy-relevant timescales. The evidence here—that firms reshape production technologies in response to sustained environmental stress, flattening damage curves through factor reallocation and directed R&D—suggests that projections abstracting from endogenous technical change may overstate long-run costs, particularly for economies with strong innovation capacity. This implies that fostering innovation capacity through R&D subsidies, IP protection, and technology diffusion may be as central to climate resilience as emissions mitigation itself.

The mechanism generates testable predictions about distributional consequences across multiple dimensions. Labor-biased heat shocks induce capital-labor substitution and redirect R&D toward automation, implying that adaptation will be stronger in sectors with high factor substitutability and robust innovation ecosystems. This creates three forms of uneven adjustment. First, across factors: labor-saving technical change imposes costs on workers in automatable tasks while capital owners capture productivity gains. Second, across firms: the innovation response documented here concentrates among larger, resource-rich incumbents with patent capacity, while smaller firms lacking R&D budgets or access to capital cannot adapt as readily. This may widen productivity dispersion and concentrate market power as climate-resilient innovators capture share from vulnerable non-innovators, amplifying inequality within industries. Third, across regions: areas with weak innovation ecosystems lag in adaptive capacity even as climate exposure intensifies. Future work tracing these channels—particularly linking patent data to matched employer-employee records and firm dynamics—would illuminate who adapts, who bears costs, and where policy can most effectively

intervene.

As climate change is a global problem, external validity is critical for assessing global adaptation potential. When it comes to the role of innovation as an adaptation margin, two dimensions vary systematically between developed and developing economies, but in opposing directions: innovation capacity (IP strength, capital availability, R&D infrastructure, diffusion speed) and climate exposure severity. My setting—EU manufacturing with strong institutions and abundant resources—represents a best case for innovation supply but experiences relatively mild heat. Developing countries face the reverse: weaker innovation ecosystems constrain adaptive technology supply, yet far more extreme heat creates stronger demand and potentially larger returns. Whether demand-side forces can overcome supply-side constraints in determining net adaptive responses remains an open empirical question with profound implications. If innovation capacity binds more tightly than climate incentives, directed technical change will remain concentrated in high-capacity economies despite their milder exposure, and climate vulnerability will be determined as much by adaptive capacity as by physical exposure. This could reshape comparative advantage and widen international inequality, as regions unable to innovate experience persistent productivity losses while climate-resilient economies capture market share in heat-exposed sectors. Conversely, if strong climate pressure stimulates innovation in resource-constrained settings—through informal adaptation, South-South technology transfer, or policy interventions easing adoption barriers—developing countries may exhibit comparable adaptive responses. A natural question then follows: could easing cross-border technology transfer help bridge the supply-demand gap, connecting high-capacity innovators in temperate regions with high-demand adopters in severely exposed economies? Understanding which scenario prevails, and whether policy can facilitate these transfer mechanisms, will shape international climate finance, technology transfer agreements, and development assistance strategies. Extending this analysis to low-income countries is therefore paramount.

Figures and Table

Productivity (TFP)

Table 6: Effect of Temperature Bins on Productivity

Labor Intensity Group	log(TFP)	
	Low (1)	High (2)
<30	0.0020 (0.0018)	-0.0032 (0.0032)
30-40	-0.0009 (0.0007)	-0.0028** (0.0011)
40-50	0.0007 (0.0011)	0.0003 (0.0006)
50-60	0.0008** (0.0004)	-0.0002 (0.0004)
70-80	-0.0015** (0.0007)	-0.0035*** (0.0011)
80-90	-0.0012 (0.0008)	-0.0041*** (0.0014)
>90	-0.0022* (0.0012)	-0.0051*** (0.0013)
Precipitation	0.0001 (0.0001)	-2.4×10^{-5} (0.0001)
Observations	6,915,846	2,148,213
Firm FE	✓	✓
Country×Year FE	✓	✓
Sector×Year FE	✓	✓

Notes: Estimating equation uses firm FE, country-year FE, and sector-year FE. Standard errors clustered by firm and country-year. Coefficients are relative to $[60, 70)^{\circ}\text{F}$; temperature-bin effects shown in percent per day. “High” is the most labor-intensive quartile (within country-year); “Low” pools the remaining quartiles. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Market Share

Table 7: Market Share and Temperature Bins: Heterogeneity by Labor Intensity

Labor Intensity Group	Market Share (log)	
	Low (1)	High (2)
<30	0.0238 (0.0375)	-0.0427** (0.0169)
30-40	0.0141 (0.0118)	0.0041 (0.0072)
40-50	0.0159 (0.0125)	0.0173*** (0.0054)
50-60	0.0107 (0.0115)	0.0078** (0.0036)
70-80	0.0058 (0.0113)	-0.0094* (0.0049)
80-90	-0.0002 (0.0140)	-0.0253*** (0.0067)
>90	0.0025 (0.0224)	-0.0377*** (0.0123)
Precipitation	-0.0002 (0.0032)	0.0026** (0.0012)
Observations	6,915,846	2,148,213
Firm FE	✓	✓
Country×Year FE	✓	✓
Sector×Year FE	✓	✓

Notes: Dependent variable is $\ln(\text{Market Share})$. Markets are defined at the country by NACE-4 industry level. Coefficients are semi-elasticities relative to the $[60, 70)^{\circ}F$ reference bin; precipitation is included as a control. Specification includes firm fixed effects, country-year fixed effects, and sector-year fixed effects. Labor-intensity ranks are computed annually within country-year and split into quartiles (Q4). “High” is the most labor-intensive quartile; “Low” pools the remaining quartiles. Standard errors are two-way clustered by market and country-year. Temperature-bin coefficients are reported in percent. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Elasticity and Factor Ratio

Elasticity

Table 8: Output Elasticities (Capital and Labor) vs *TMAX* Bins — Firm Level

	log(Capital Output Elasticity)	log(Labor Output Elasticity)
	(1)	(2)
<30	0.0041 (0.0205)	0.0026 (0.0018)
30-40	0.0084*** (0.0023)	0.0004 (0.0004)
40-50	0.0008 (0.0058)	0.0013*** (0.0004)
50-60	-0.0054 (0.0050)	0.0015* (0.0008)
70-80	0.0059 (0.0101)	3.89×10^{-5} (0.0015)
80-90	0.0082 (0.0052)	-0.0002 (0.0012)
>90	0.0463*** (0.0161)	-0.0071*** (0.0020)
Precipitation	-0.0034* (0.0020)	0.0006*** (9.52×10^{-5})
Observations	9,064,059	9,064,059
Firm FE	✓	✓
Sector×Year FE	✓	✓
Country×Year FE	✓	✓

Notes: Coefficients are semi-elasticities (scaled $\times 100$) relative to the $[60, 70)^\circ\text{F}$ bin. Specification includes firm fixed effects, country×year fixed effects, and sector×year fixed effects; standard errors are clustered by firm and country–year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 9: Capital and Labor Output Elasticities vs. *TMAX* Bins—Industry Level

Series	log(Capital Output Elasticity) (1)	log(Labor Output Elasticity) (2)
<30	0.5320 (0.4006)	-0.2835 (0.3814)
30-40	-0.1156 (0.3590)	0.2406 (0.3970)
40-50	-0.0930 (0.3409)	-0.0679 (0.3866)
50-60	-0.2138 (0.2264)	-0.0886 (0.2020)
70-80	-0.3635 (0.2825)	0.1392 (0.4086)
80-90	-0.1324 (0.4406)	0.2513 (0.4435)
>90	1.055*** (0.2902)	-0.4884 (0.3820)
Observations	7,311	7,311
Country×Year FE	✓	✓
Sector×Year FE	✓	✓

Notes: Entries report coefficients (semi-elasticities, $\times 100$) for bins of daily maximum temperature *TMAX*, relative to $[60, 70)^\circ\text{F}$. Dependent variables are industry-level output elasticities for capital (col. 1) and labor (col. 2), computed as sales-share-weighted means of firm-level elasticities within country–industry–year cells. Specifications include country×year and sector×year fixed effects; standard errors are clustered by country–industry and year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Factor Input

Table 10: Capital and Labor Cost vs $TMAX$ Bins — Firm Level

	Capital (1)	Labor Cost (2)
<30	-0.0449** (0.0202)	-0.0314* (0.0182)
30-40	0.0192 (0.0130)	0.0016 (0.0144)
40-50	0.0359*** (0.0076)	0.0234*** (0.0070)
50-60	0.0094* (0.0055)	0.0108 (0.0068)
70-80	-0.0137* (0.0079)	-0.0113 (0.0089)
80-90	0.0061 (0.0174)	-0.0267*** (0.0095)
>90	0.0527 (0.0364)	-0.0330** (0.0149)
Precipitation	0.0006 (0.0033)	0.0022 (0.0017)
Observations	9,064,059	9,064,059
Firm FE	✓	✓
Country×Year FE	✓	✓
Sector×Year FE	✓	✓

Notes: Coefficients are semi-elasticities (scaled $\times 100$) relative to the $[60, 70)^\circ\text{F}$ bin. Specification includes firm fixed effects, country×year fixed effects, and sector×year fixed effects; standard errors are clustered by firm and country–year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Mitigation

Table 11: Summary of aggregate mitigation

	Mitigation (%)
GDP-weighted Mean	26.1
Unweighted Mean	18.1
Range	[0.0, 40.0]

Notes: This table shows the aggregate summary of the country level mitigation measure calculated using equation (12). I use country-level GDP at 2020 as weights for aggregating across country-level mitigation values.

References

- Acemoglu, Daron**, “Directed technical change,” *The review of economic studies*, 2002, *69* (4), 781–809.
- , “Labor-and capital-augmenting technical change,” *Journal of the European Economic Association*, 2003, *1* (1), 1–37.
- , “When Does Labor Scarcity Encourage Innovation?,” *Journal of Political Economy*, 2010, *118* (6), 1037–1078.
- and **Pascual Restrepo**, “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 2020, *128* (6), 2188–2244.
- and —, “Demographics and automation,” *The Review of Economic Studies*, 2022, *89* (1), 1–44.
- , **Philippe Aghion**, **Leonardo Bursztyn**, and **David Hemous**, “The environment and directed technical change,” *American economic review*, 2012, *102* (1), 131–166.
- Akerberg, Daniel A**, **Kevin Caves**, and **Garth Frazer**, “Identification properties of recent production function estimators,” *Econometrica*, 2015, *83* (6), 2411–2451.
- Addoum, Jawad M**, **David T Ng**, and **Ariel Ortiz-Bobea**, “Temperature shocks and establishment sales,” *The Review of Financial Studies*, 2020, *33* (3), 1331–1366.
- Adhvaryu, Achyuta**, **Namrata Kala**, and **Anant Nyshadham**, “The light and the heat: Productivity co-benefits of energy-saving technology,” *Review of Economics and Statistics*, 2020, *102* (4), 779–792.
- Aghion, Philippe**, **Antoine Dechezleprêtre**, **David Hémous**, **Ralf Martin**, and **John Van Reenen**, “Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry,” *Journal of Political Economy*, 2016, *124* (1), 1–51.
- Autor, David**, **David Dorn**, **Lawrence F Katz**, **Christina Patterson**, and **John Van Reenen**, “The Fall of the Labor Share and the Rise of Superstar Firms*,” *The Quarterly Journal of Economics*, 02 2020, *135* (2), 645–709.
- Bajgar, Matej**, **Giuseppe Berlingieri**, **Sara Calligaris**, **Chiara Criscuolo**, and **Jonathan Timmis**, “Industry Concentration in Europe and North America,” 2019, (18).
- , —, —, —, —, and —, “Coverage and representativeness of Orbis data,” 2020.
- Baqae, David Rezza** and **Emmanuel Farhi**, “Productivity and misallocation in general equilibrium,” *The Quarterly Journal of Economics*, 2020, *135* (1), 105–163.
- Barreca, Alan**, **Karen Clay**, **Olivier Deschenes**, **Michael Greenstone**, and **Joseph S Shapiro**, “Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the twentieth century,” *Journal of Political Economy*, 2016, *124* (1), 105–159.
- Berndt, Ernst R** and **David O Wood**, “Technology, prices, and the derived demand for energy,” *Review of Economics and Statistics*, 1975, *57* (3), 259–268.

- Bloom, Nicholas, John Van Reenen, and Heidi Williams**, “A toolkit of policies to promote innovation,” *Journal of Economic Perspectives*, 2019, *33* (3), 163–184.
- , **Mark Schankerman, and John Van Reenen**, “Identifying Technology Spillovers and Product Market Rivalry,” *Econometrica*, 2013, *81* (4), 1347–1393.
- Brynjolfsson, Erik and Lorin M. Hitt**, “Computing Productivity: Firm-Level Evidence,” *Review of Economics and Statistics*, 2003, *85* (4), 793–808.
- Burke, Marshall and Kyle Emerick**, “Adaptation to climate change: Evidence from US agriculture,” *American Economic Journal: Economic Policy*, 2016, *8* (3), 106–140.
- , **Mustafa Zahid, Mariana CM Martins, Christopher W Callahan, Richard Lee, Tumenkhusel Avirmed, Sam Heft-Neal, Mathew Kiang, Solomon M Hsiang, and David Lobell**, “Are we adapting to climate change?,” Technical Report, National Bureau of Economic Research 2024.
- Calel, Raphael and Antoine Dechezleprêtre**, “Environmental policy and directed technological change: evidence from the European carbon market,” *Review of Economics and Statistics*, 2016, *98* (1), 173–191.
- Chen, Xiaoguang and Lu Yang**, “Temperature and Industrial Output: Firm-level Evidence from China,” *Journal of Environmental Economics and Management*, 2019, *95*, 257–274.
- Christensen, Laurits R, Dale W Jorgenson, and Lawrence J Lau**, “Transcendental logarithmic production frontiers,” *Review of Economics and Statistics*, 1973, *55* (1), 28–45.
- Colmer, Jonathan**, “Weather, labor reallocation, and industrial production: Evidence from India,” *American Economic Journal: Applied Economics*, 2021, *13* (4), 101–124.
- Comin, Diego and Bart Hobijn**, “An Exploration of Technology Diffusion,” *American Economic Review*, 2010, *100* (5), 2031–2059.
- Dell, Melissa, Benjamin F Jones, and Benjamin A Olken**, “What do we learn from the weather? The new climate-economy literature,” *Journal of Economic literature*, 2014, *52* (3), 740–798.
- Dernis, H  l  ne and Mariagrazia Squicciarini**, “Harmonisation of IP Applicant Names: The OECD Experience,” 2021. Presentation, OECD STI Micro-data Lab.
- Gal, Peter N**, “Measuring total factor productivity at the firm level using OECD-ORBIS,” 2013.
- Ganglmair, Bernhard, Nadine Hahn, Michael Hellwig, Alexander Kann, Bettina Peters, and Ilona Tsanko**, “Price markups, innovation, and productivity: Evidence from Germany,” Technical Report, Produktivit  t f  r Inklusives Wachstum 2020.
- Goldschlag, Nathan, Travis J Lybbert, and Nikolas Jason Zolas**, “An ‘algorithmic links with probabilities’ crosswalk for uspc and cpc patent classifications with an application towards industrial technology composition,” *US Census Bureau Center for Economic Studies paper no. CES-WP-16-15*, 2016.
- Graetz, Georg and Guy Michaels**, “Robots at Work,” *Review of Economics and Statistics*, 2018, *100* (5), 753–768.

- Grossman, Gene M and Ezra Oberfield**, “The elusive explanation for the declining labor share,” *Annual Review of Economics*, 2022, 14 (1), 93–124.
- Hall, Bronwyn H., Zvi Griliches, and Jerry A. Hausman**, “Patents and R&D: Is There a Lag?,” *International Economic Review*, 1986, 27 (2), 265–283.
- Hémous, David, Morten Olsen, Carlo Zanella, and Antoine Dechezleprêtre**, “Induced Automation Innovation: Evidence from Firm-Level Patent Data,” *Journal of Political Economy*, 2025, 133 (6), 1975–2028.
- Heutel, Garth, Nolan H. Miller, and David Molitor**, “Adaptation and the mortality effects of temperature across US climate regions,” *Review of Economics and Statistics*, 2021, 103 (4), 740–753.
- Hicks, John R.**, *The Theory of Wages*, London: Macmillan, 1932.
- Hubmer, Joachim and Pascual Restrepo**, “Not a typical firm: The joint dynamics of firms, labor shares, and capital–labor substitution,” Technical Report, National Bureau of Economic Research 2021.
- Kalemli-Ozcan, Sebnem, Bent Sorensen, Carolina Villegas-Sanchez, Vadym Volosovych, and Sevcen Yesiltas**, “How to construct nationally representative firm level data from the Orbis global database: New facts and aggregate implications,” Technical Report, National Bureau of Economic Research 2015.
- Karabarbounis, Loukas and Brent Neiman**, “The Global Decline of the Labor Share,” *Quarterly Journal of Economics*, 2014, 129 (1), 61–103.
- Liu, Yang, Fang Zhang, and Shiyi Chen**, “Estimating the Capital–Labor Substitution in China: Evidence from the VAT Reform,” *China Economic Review*, 2020, 63, 101536.
- Loecker, Jan De and Chad Syverson**, “An industrial organization perspective on productivity,” in “Handbook of industrial organization,” Vol. 4, Elsevier, 2021, pp. 141–223.
- **and Frederic Warzynski**, “Markups and Firm-Level Export Status,” *American Economic Review*, May 2012, 102 (6), 2437–71.
- Long, Xianling and Zhiqiang Wang**, “From heat to high-tech: How innovation responds to climate change,” *Journal of Development Economics*, 2025, 176, 103525.
- Moscona, Jacob and Karthik A. Sastry**, “Does Directed Innovation Mitigate Climate Damage? Evidence from U.S. Agriculture,” *The Quarterly Journal of Economics*, 2023, 138 (2).
- Newell, Richard G., Adam B. Jaffe, and Robert N. Stavins**, “The Induced Innovation Hypothesis and Energy-Saving Technological Change,” *The Quarterly Journal of Economics*, 1999, 114 (3), 941–975.
- Oberfield, Ezra and Devesh Raval**, “Micro Data and Macro Technology,” *Econometrica*, 2021, 89 (2), 703–732.
- OECD Directorate for Science, Technology and Innovation**, “OECD HAN Database: Harmonised Applicant Names, January 2024 edition,” 2024. Database documentation and methodology.

- , “OECD REGPAT Database: EPO & PCT Patent Applications at Regional Level, January 2024 edition,” 2024. Database documentation and methodology.
- Park, R. Jisung, Nora Pankratz, and A. Patrick Behrer**, “Temperature, Workplace Safety, and Labor Market Inequality,” IZA Discussion Paper 14560, Institute of Labor Economics (IZA) July 2021.
- Ponticelli, Jacopo, Qiping Xu, and Stefan Zeume**, *Temperature and local industry concentration*, National Bureau of Economic Research, 2023.
- Popp, David**, “Induced Innovation and Energy Prices,” *American Economic Review*, 2002, 92 (1), 160–180.
- Raval, Devesh**, “The Micro Elasticity of Substitution and Non-Neutral Technology,” *American Economic Journal: Applied Economics*, 2019, 11 (3), 190–217.
- Somanathan, Eswaran, Rohini Somanathan, Anant Sudarshan, and Meenu Tewari**, “The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing,” *Journal of Political Economy*, 2021, 129 (6), 1797–1827.
- Weche, John P. and Achim Wambach**, “The Fall and Rise of Market Power in Europe,” *Jahrbucher fur Nationalokonomie und Statistik*, 11 2021, 241, 555–575.
- Xie, Victoria Wenxin**, “Heterogeneous firms under regional temperature shocks: exit and reallocation, with evidence from Indonesia,” *Economic Development and Cultural Change*, 2024, 72 (2), 659–690.
- Zhang, Peng, Olivier Deschenes, Kyle Meng, and Junjie Zhang**, “Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants,” *Journal of Environmental Economics and Management*, 2018, 88, 1–17.
- Zivin, Joshua Graff and Matthew Neidell**, “Temperature and the allocation of time: Implications for climate change,” *Journal of Labor Economics*, 2014, 32 (1), 1–26.

Appendix

Extreme Heat and Directed Innovation

Enjie (Jack) Ma

Table of Contents

A1	Climate data: additional tables and figures	44
A2	ORBIS Data	46
A2.1	Data Cleaning and Imputation	46
A2.2	HAN-ORBIS Matching	48
A3	Production Function Estimation	50
A4	Model of Directed Technical Change	56
A4.1	Setup	56
A4.2	Direction of Innovation	57
A4.3	Innovation as a Buffer Against Direct Damage	58
A5	Additional Figures and Tables	60
A5.1	Temperature and Firm Productivity	60
A5.2	Heat-Induced Market Share Dynamics	62
A5.3	Technology Adjustment: Output Elasticities	63
A5.3.1	Factor Ratio	64
A5.3.2	Industry-Level Technology Response	65
A5.4	Directed Innovation: Patent Evidence	67
A5.4.1	Alternative Patent Measure: Count vs Share	67
A5.4.2	Lag Robustness	68
A5.4.3	Long-Difference Specifications	69
A5.5	Damage Mitigation from Innovation	70

A1 Climate data: additional tables and figures

Table A1: Summary Statistics: Annual Days in Daily Maximum Temperature Bins

TMAX Bin	Mean	SD	p25	p50	p75	min	max
<30 °F	3.8	12.9	0	0	1	0	157
30–40°F	20.1	24.5	1	11	28	0	161
40–50°F	46.9	26.3	26	51	67	0	135
50–60°F	79.3	21.7	67	78	93	0	195
60–70°F	81.6	19.7	68	80	94	1	194
70–80°F	70.8	19.0	59	71	82	0	285
80–90°F	48.1	26.4	26	49	68	0	135
>90°F	14.7	19.5	1	7	20	0	225

Notes: Summary statistics for the annual distribution of daily maximum temperature bins across all firm-year observations. Each row shows the number of days per year falling in the specified temperature range.

Table A2: Country averages of CDD (2000–2020)

iso	N_years	mean_CDD	sd_CDD	p50_CDD
ES	21	384.7	92.1	360.5
IT	21	157.9	71.7	138.2
FR	21	87.5	55.7	77.9
AT	21	47.6	33.3	39.2
LU	21	42.7	38.3	26.2
DE	21	41.6	30.7	24.2
BE	21	31.7	24.4	27.1
FI	21	2.9	5.1	0.4
DK	19	1.8	3.7	0.6


Notes: Average and dispersion of country-year mean CDD computed from firm-level exposures within each country; units are °F-days.

Table A3: Country-level CDD in 2000 and 2020, and 20-year changes

iso	CDD_2000	CDD_2020	delta	pct
ES	310.5	492.3	181.8	58.6
FR	33.6	160.7	127.1	377.6
LU	8.5	95.0	86.5	1017.1
BE	10.6	86.3	75.7	715.1
DE	18.8	55.1	36.3	192.7
IT	114.6	134.1	19.6	17.1
DK	3.1	16.4	13.3	431.1
FI	0.4	3.2	2.8	696.5
AT	35.7	26.9	-8.8	-24.6
Overall (mean across countries)	59.5	118.9	59.4	99.7

Notes: Country averages of Cooling Degree Days (CDD, measured as extreme heat degree-days above 85 °F, or 29.4°C). Reported values are computed as averages of firm-level exposures aggregated to the country-year level. The final column reports the 20-year change between 2000 and 2020, both in absolute and percentage terms.

Figure A1: Average cooling degree-days by country (EU; 2000–2020)



Notes: Country polygons colored by mean CDD over the sample window (°F-days).

A2 ORBIS Data

A2.1 Data Cleaning and Imputation

Data Cleaning There are several types of firm accounts available in our ORBIS dataset: Consolidated (consolidation code C1, C2), Unconsolidated (U1,U2), and Limited Financial(LF). I use only Unconsolidated accounts data from ORBIS, as I want to match the economic variables of a plant to its specific location and the local temperature shock it experiences, as well as to avoid double counting. Restricting the sample to be only Unconsolidated accounts also better approximates a specific product market, which further validates the production function estimation approach. On the other hand, Consolidated accounts observations might not reflect the output of the local plant, as it can be balance sheet data of a headquarter, instead of a specific plant.

Using only Unconsolidated accounts information can understate the level of concentration measure, since it can ignore firm-to-firm linkages within different business groups (Kalemli-Ozcan et al., 2015). The literature has adopted different approaches when it comes to measuring concentration: whether it is at the individual firm-level or the business group level. Each of these approach has its advantages and limitations, depending on the data availability and research questions²⁹. However, in our context, I believe focusing on individual plants' performances is more appropriate for analyzing the effect of local temperature shock. In addition, across most countries and industries, Unconsolidated firms account for more than the majority of the market share (Bajgar et al., 2020). The broad trend of concentration change in the EU carries through, regardless of different unit of measures, though the trajectory is a bit different(Bajgar et al., 2019).³⁰

From our production function estimation, our data cleaning procedure closely follows Weche and Wambach (2021) and Ganglmair et al. (2020), albeit a bit less restrictive. Observations that have missing or insensible data in these variables are dropped : *OPERATING_REVENUE_TURNOVER*, *NUMBER_OF_EMPLOYEES*. I also keep only Unconsolidated accounts, as in the context of production function estimation Unconsolidated accounts are closer to a product market compared to Consolidated.

Many researchers using ORBIS tend to impose a fixed cutoff threshold for number of employees to address the problem of ORBIS under-representing small firms. Doing so can ensure a more stable coverage over time and a better defined distribution. However, it would also diminish the ability to approximate the true distribution of firms, especially for the micro firms, which can be more vulnerable to climate change. Imposing a fixed cutoff threshold across different country-industries can also be misleading, as it can cutoff different parts of the firm distribution in different country-industries. Therefore, in the context of our research questions, I think it's also crucial to analyze how smaller firms are affected by climate change. Thus, I adopt a less restrictive data cleaning approach than the literature and do not impose any size restrictions at the moment.

The climate data extraction procedure requires data on geographical locations. My preferred geographical variables are longitude and latitude of a plant. When those are not available, I use street address or postcode to geocode the location. Observations that do not have any of these geographical variables are dropped.

²⁹For a more complete discussion on the accounts types in ORBIS, see Bajgar et al. (2020)

³⁰For more discussions on different approaches to calculate concentration in EU, see Bajgar et al. (2019)

Deflator For output measures, I use Producer Price Index data from Eurostat, which is available at the country by NACE 2-digit industries level. For labor cost, I use Labor Cost Index at the same country by NACE 2-digit industries level.

Imputation Following [Gal \(2013\)](#), I perform internal imputation for the Value-added and Material Costs variable before the production function estimation. As suggested by [Bajgar et al. \(2020\)](#), imputation of value added using information on wage bill and earnings can partially improve the representativeness of ORBIS data, which has been a well-documented problem with ORBIS. In particular, the mean characteristics and the representativeness of the bottom half of firm distribution is closer to that of the true population after imputation. My imputation to date is internal, using information from within the ORBIS data. External imputation, which requires industry level wage bill data, is also possible. However, the effect of external imputation is minimal and can diminishes the dispersion within industries ([Bajgar et al., 2020](#)). Hence, I plan to incorporate the external imputation in later versions of the draft as a further robustness check, while maintaining internal imputation as our primary imputation method.

A2.2 HAN–ORBIS Matching

This appendix section documents the construction of the crosswalk between OECD patent applicants and ORBIS firms. I use *REGPAT* as the application-level backbone (keys: `appln_id`, applicant name, country) and the OECD *Harmonised Applicant Names* (HAN) resource as a name spine (key: `HAN_id`) that consolidates applicant-name variants and is explicitly designed to bridge to business registers such as ORBIS (OECD Directorate for Science, Technology and Innovation, 2024b,a; Dernis and Squicciarini, 2021). Because the OECD HAN–ORBIS crosswalk is not publicly distributed, I develop my own procedure to link `HAN_id` to Bureau van Dijk firm identifiers (`BvDID`).

Scope and unit of observation All matching is performed *by country* (ISO2) within the European sample used in the paper. Within each country, the unit of observation is an applicant group (`HAN_id`) and an ORBIS entity (`BvDID`). The final crosswalk retains a unique one-to-one link between each `HAN_id` and `BvDID`.³¹

Preprocessing and blocking I standardize applicant and firm names using conservative rules that minimise the risk of false positives:

- **Canonicalisation:** uppercase, diacritic stripping, punctuation and repeated whitespace removal.
- **Tokenization:** alphanumeric tokenization preserving stopwords (I do *not* drop generic tokens such as *company*, *industries*).
- **Legal-form handling:** only a “safe” set of corporate suffixes is removed (e.g., GMBH, AG, SAS, SARL, SRL, SpA, SL, LTD, SA, Oy, ApS, AB, NV, BV); the list is country-aware and intentionally conservative.
- **Blocking:** candidates are generated only within tight blocks defined by country and the first 5 characters of the canonical name.

Three matching phases I proceed in three phases, each nested inside the previous failure set:

- (A) **Strict exact** — accept if the fully canonicalised strings are identical.
- (B) **Safe-suffix exact** — accept if strings are identical after removing only the legal-form tokens listed above.
- (C) **Fuzzy (controlled)** — within blocks, compute Jaro–Winkler similarity and accept pairs that satisfy all of: similarity ≥ 0.92 , length ratio in $[0.70, 1.40]$, and first-token agreement.

Uniqueness and tie-breaking To obtain a research-quality crosswalk, I impose a deterministic one-to-one mapping between `HAN_id` and `BvDID`. If a `HAN_id` (or `BvDID`) has multiple candidates, I keep the single best pair by (i) phase priority (strict > safe-suffix > fuzzy), then (ii) similarity score, (iii) length difference, and (iv) first-token agreement. All discarded candidates are logged; only unique pairs are used in the analysis.

³¹The matching is limited to EU/EEA countries covered by the ORBIS sample. Since larger firms are more likely to file patents and ORBIS over-represents larger firms, the omission of smaller non-covered firms is unlikely to materially affect the representativeness of the linked sample.

Diagnostics and validation I conduct extensive diagnostics to verify precision (including borderline-score audits, duplicate-structure checks, and token/suffix sanity checks) and tuned thresholds accordingly. These diagnostics confirm that accepted links are dominated by exact or near-exact agreement; borderline fuzzy cases constitute a small share and are excluded from the research file. The conservative design ensures that the matches I retain are highly reliable and unlikely to include false positives, even if this comes at the cost of leaving some true matches unmatched.

Coverage and composition The final one-to-one crosswalk (Run 3, v3) links approximately 243,000 applicant–firm pairs across European countries. Roughly one half of links are strict exact matches, about one fifth become exact once legal-form tokens are removed, and the remaining third are high-score fuzzy links; all links used in the paper are one-to-one by construction. Table A4 reports the composition.

Table A4: HAN–ORBIS Match Composition (EU sample, Run 3, v3)

Method	Count	Share
Strict exact	119,819	49.3%
Safe-suffix exact	46,636	19.2%
Fuzzy (Jaro–Winkler ≥ 0.92)	76,424	31.5%
Total	242,879	100%

A3 Production Function Estimation

I estimate country-industry-level production functions using a two-step GMM procedure to recover firm-level productivity and output elasticities. The approach follows the ACF (Akerberg et al., 2015) control function methodology and allows technology parameters to vary across industries and over time.

Production Function and Output Elasticity Consider the following production function, where Q_{it} is the output of firm i in time t , X_{it}^V is the variable input v , K_{it} is the capital,

$$Q_{it} = F(X_{it}^1, \dots, X_{it}^V, K_{it}; \beta) \exp(\omega_{it}), \quad (\text{A1})$$

and the log version

$$q_{it} = f(x_{it}^1, \dots, x_{it}^V, k_{it}; \beta) + \omega_{it} + \varepsilon_{it}. \quad (\text{A2})$$

I follow the specification of (De Loecker and Warzynski, 2012) where $f(\cdot)$ to be a *translog*³² production function with second order polynomial plus interaction, with labor being the variable input and capital being the dynamic input.

$$q_{it} = \beta_l l_{it} + \beta_k k_{it} + \beta_{ll} l_{it}^2 + \beta_{kk} k_{it}^2 + \beta_{lk} l_{it} k_{it}. \quad (\text{A3})$$

Under this specification of the production function, I can then calculate the output elasticity with respect to labor or capital as the following:

$$\hat{\theta}_{it}^L = \frac{\partial \ln F(\cdot)}{\partial \ln L_{it}} = \frac{\partial f(\cdot)}{\partial l_{it}} = \hat{\beta}_l + 2\hat{\beta}_{ll} l_{it} + \hat{\beta}_{lk} k_{it} \quad (\text{A4})$$

$$\hat{\theta}_{it}^K = \frac{\partial \ln F(\cdot)}{\partial \ln K_{it}} = \frac{\partial f(\cdot)}{\partial k_{it}} = \hat{\beta}_k + 2\hat{\beta}_{kk} k_{it} + \hat{\beta}_{lk} l_{it} \quad (\text{A5})$$

Table A5: Summary Statistics of Structural Parameters (Translog)

Parameter	N	Mean	SD	Median
β_l	1738	0.0084	0.3412	-0.0056
β_k	1738	0.3353	0.1839	0.3451
β_{ll}	1738	0.0598	0.0179	0.0610
β_{lk}	1738	-0.0377	0.0303	-0.0382
β_{kk}	1738	0.0179	0.0132	0.0167

³²To allow variations in technology (output elasticity) across firms and over time, I choose translog over the standard Cobb-Douglas, where output elasticity is fixed. The estimation is done at a 5-year window to allow for changes in structural parameter as well.

Two-step GMM The estimation procedure consists of two steps. The first stage is for purging out measurement errors ε . The second step uses moment conditions to estimate the β coefficients. The proxy variable, material costs in my case, is assumed to be a function of productivity and capital inputs. The inversion of the material demand function gives:

$$\omega_{it} = h_t(k_{it}, m_{it}, z_{it}),$$

where z_{it} are additional industry and time fixed effects that can affect material demand other than productivity and inputs. The production function then becomes

$$y_{it} = \underbrace{f(l_{it}, k_{it}) + \overbrace{h_t(m_{it}, k_{it}, \mathbf{z}_{it})}^{\text{Productivity } \omega_{it}}}_{\text{Expected Output } \hat{\phi}_{it}(l_{it}, k_{it}, m_{it}, z_{it})} + \varepsilon_{it}.$$

The residuals from the first stage is then

$$\hat{\varepsilon}_{it} = y_{it} - \hat{\phi}_{it}(l_{it}, k_{it}, m_{it}, z_{it}),$$

and the productivity can be expressed as a function of β 's

$$\omega_{it}(\beta) = \hat{\phi}_{it}(l_{it}, k_{it}, m_{it}, z_{it}) - f(l_{it}, k_{it}).$$

Productivity is assumed to follow a first-order Markov Process:

$$\omega_{it} = g(\omega_{it-1}) + \xi_{it},$$

where ξ_{it} is the productivity shock. The key identification assumption comes from the orthogonality between current productivity shock and current state variable (capital), as well as between current productivity shock and lagged free variable (labor). Namely, capital inputs are decided dynamically, while labor adjusts more freely and responds to contemporaneous productivity shock. Therefore, in the second stage, $\hat{\beta}$ are derived from Moments conditions:

$$E \left(\xi_t(\beta) \cdot \begin{pmatrix} K_t \\ L_{t-1} \end{pmatrix} \right) = 0.$$

Once $\hat{\beta}$ is estimated, I can then derive firm-level productivity $\hat{\omega}_{it}$ and output elasticity $\hat{\theta}_l$ and $\hat{\theta}_k$.

Endogenous Temperature Shock Motivated by [De Loecker and Syverson \(2021\)](#), I treat temperature shock as an active driver of productivity and incorporate it directly into the productivity process, to allow temperature to endogenously affect production process. Accordingly, I augment the Markov process to allow temperature (measured by cooling degree days, CDD) to shift productivity ³³:

$$\omega_{it} = g(\omega_{i,t-1}, \text{CDD}_{it}) + \xi_{it}.$$

I view CDD as plausibly exogenous with respect to firms' choices but observable/forecastable to producers when making input decisions. This specification avoids pushing systematic weather variation into the innovation term and lets me quantify the link between temperature and productivity within the same production-function framework.

³³Empirically, I find weak serial correlation in temperature shocks at the firm level: firm-level AR(1) estimates cluster near zero, and pooled regressions show only weak persistence. This supports the Markov-type assumption that CDD can be treated as an exogenous state variable in the productivity process.

Elasticity of Substitution Following Christensen et al. (1973) and the standard translog literature (Berndt and Wood, 1975), the (Hicks–)Allen–Uzawa elasticity of substitution between capital and labor is

$$\sigma_{KL,it} = \frac{\beta_{lk} + \theta_{K,it}\theta_{L,it}}{\theta_{K,it}\theta_{L,it}} = 1 + \frac{\beta_{lk}}{\theta_{K,it}\theta_{L,it}}, \quad (\text{A6})$$

where $\theta_{K,it}$ and $\theta_{L,it}$ are the output elasticities of capital and labor for firm i at time t , defined as in equation (A4). Under profit maximization with competitive input *and* output markets (so $w_j = p \cdot \text{MP}_j$), these output elasticities equal the corresponding *revenue* shares


The sign and magnitude of β_{lk} pin down the substitution regime. The threshold $\sigma_{KL,it} = 1$ is pivotal in the directed technical change (DTC) framework (Acemoglu, 2002): when $\sigma_{KL,it} > 1$ (stronger substitutability), the market-size effect dominates the price effect and innovation is directed toward the more abundant factor; when $\sigma_{KL,it} < 1$ (complementarity or weak substitutability), the price effect dominates and innovation is directed toward the scarcer factor.³⁴

On average across 1,738 country–industry pairs, I estimate $\hat{\beta}_{lk} = -0.038$ (Table A5). Using standard factor shares $\theta_{L,it} = 0.65$ and $\theta_{K,it} = 0.35$ implies $\sigma_{KL} \approx 1 + (-0.038)/(0.65 \times 0.35) \approx 0.83$, indicating that capital and labor are weak substitutes on average—below the Cobb–Douglas benchmark of unity. This aggregate masks important dynamics. Figure A2 shows a pronounced upward shift in the distribution of β_{lk} over time: the median rises from roughly -0.06 in 2000–2004 to near zero by 2015–2019, indicating steadily increasing substitutability and a local elasticity approaching the DTC threshold of one.³⁵

³⁴Under the translog specification in (A6), $\sigma_{KL,it} \geq 1$ if and only if $\beta_{lk} \geq 0$, because $\theta_{K,it}\theta_{L,it} > 0$. Capital and labor are *gross* substitutes when $\sigma_{KL,it} > 0$, i.e., when $\beta_{lk} > -\theta_{K,it}\theta_{L,it}$; they become gross complements when $\beta_{lk} < -\theta_{K,it}\theta_{L,it}$.

³⁵Because $\sigma_{KL,it}$ is local and varies with $\theta_{K,it}\theta_{L,it}$, the interpretation here is conditional on observed shares; the sign mapping $\sigma_{KL,it} \geq 1 \Leftrightarrow \beta_{lk} \geq 0$, however, does not depend on shares.

Figure A2: Rising Trend in Capital–Labor Substitutability



Notes: The figure shows the distribution of estimated β_{LK} across all country–industry pairs for four consecutive five-year windows. Boxes mark the interquartile range with the median shown by a diamond; whiskers extend to $1.5 \times \text{IQR}$. The horizontal dashed line at zero corresponds to the Cobb–Douglas locus where $\sigma_{KL} = 1$. The upward shift over 2000–2019 indicates increasing substitutability, consistent with conditions under which DTC predicts market-size effects favor capital-biased (labor-saving) innovation.

Industry Labor Intensity Summary

Table A6: Labor Elasticity (LI) — Summary Statistics (BCDF sectors)

Min	Q1	Median	Mean	Q3	Max
0.25	0.38	0.45	0.45	0.51	0.61

Table A7: Top 15 Labor-Intensive Industries (BCDF)

Industry (NACE2)	Description	Labor Elasticity
43	Specialised construction activities	0.61
33	Repair and installation of machinery and equipment	0.59
18	Printing and reproduction of recorded media	0.59
42	Civil engineering	0.58
19	Manufacture of coke and refined petroleum products	0.53
13	Manufacture of textiles	0.52
26	Manufacture of computer, electronic and optical products	0.52
25	Manufacture of fabricated metal products, except machinery and equipment	0.51
30	Manufacture of other transport equipment	0.49
32	Other manufacturing	0.49
15	Manufacture of leather and related products	0.48
29	Manufacture of motor vehicles, trailers and semi-trailers	0.48
14	Manufacture of wearing apparel	0.48
27	Manufacture of electrical equipment	0.47
20	Manufacture of chemicals and chemical products	0.43

Table A8: Bottom 15 Labor-Intensive Industries (BCDF)

Industry (NACE2)	Description	Labor Elasticity
35	Electricity, gas, steam and air conditioning supply	0.25
12	Manufacture of tobacco products	0.31
11	Manufacture of beverages	0.31
16	Manufacture of wood and of products of wood and cork	0.35
24	Manufacture of basic metals	0.36
41	Construction of buildings	0.37
10	Manufacture of food products	0.38
23	Manufacture of other non-metallic mineral products	0.38
22	Manufacture of rubber and plastic products	0.39
17	Manufacture of paper and paper products	0.4
21	Manufacture of basic pharmaceutical products and pharmaceutical preparations	0.42
28	Manufacture of machinery and equipment n.e.c.	0.43
31	Manufacture of furniture	0.43
20	Manufacture of chemicals and chemical products	0.43
27	Manufacture of electrical equipment	0.47

A4 Model of Directed Technical Change

This appendix presents a parsimonious aggregate model in the quality-ladder fashion of [Acemoglu et al. \(2012\)](#) to formalize two mechanisms linking climate shocks to innovation. I adapt this canonical directed technical change framework to the specific case of asymmetric, labor-biased climate damage. First, the model delivers a qualitative result for the direction of innovation: a labor-biased heat shock shifts research effort toward labor-saving (capital-augmenting) technologies when factors are sufficiently substitutable. Second, it shows how endogenous innovation partially offsets direct productivity damage through accumulated quality improvements. Throughout, I emphasize these qualitative predictions and maintain results conditional on the elasticity of substitution $\sigma > 1$.³⁶

A4.1 Setup

Production structure. A unique final good is produced from two sectoral bundles aggregated by a constant-elasticity-of-substitution (CES) technology:

$$Y_t = \left[\alpha Y_{K,t}^\rho + (1 - \alpha) Y_{L,t}^\rho \right]^{1/\rho}, \quad 0 < \alpha < 1, \rho < 1, \sigma = \frac{1}{1-\rho} > 0, \quad (\text{A7})$$

where σ is the elasticity of substitution between the capital-intensive and labor-intensive bundles. Each sectoral bundle aggregates a continuum of machine varieties with Cobb–Douglas weights over an input factor and variety-specific quality:

$$Y_{K,t} = K_t^{1-\alpha} \int_0^1 A_{Ki,t}^{1-\alpha} x_{Ki,t}^\alpha di, \quad (\text{A8})$$

$$Y_{L,t} = [\textcolor{red}{D}(T_t) L_t]^{1-\alpha} \int_0^1 A_{Li,t}^{1-\alpha} x_{Li,t}^\alpha di, \quad (\text{A9})$$

where K_t denotes the capital stock employed in the K -bundle, L_t is production labor, and $\textcolor{red}{D}(T_t) \in (0, 1]$ represents labor efficiency scaled by climate conditions. I capture labor-biased heat exposure through $\textcolor{red}{D}(T_t)$: higher temperatures reduce effective labor input in the L -bundle while leaving capital undamaged.³⁷

Innovation and quality ladders. A unit mass of scientists allocates effort between capital-augmenting and labor-augmenting research with shares $s_{K,t}$ and $s_{L,t} = 1 - s_{K,t}$. Average quality in each sector evolves according to a simple quality ladder:

$$A_{x,t} = (1 + \gamma \eta_x s_{x,t}) A_{x,t-1}, \quad A_{x,t} \equiv \int_0^1 A_{xi,t} di, \quad x \in \{K, L\}, \quad (\text{A10})$$

where $\gamma > 0$ is the step size of quality improvement and $\eta_x \in (0, 1)$ is the arrival rate of innovation in sector x . Scientists build on the existing technology frontier: today’s innovation scales yesterday’s average quality.

³⁶This condition is consistent with the pattern of factor substitutability in my sample: using a translog production function specification, I estimate an elasticity of substitution that increases over the study period and exceeds unity in recent years (see Appendix Section A3 for details). This time-varying pattern suggests growing substitutability between capital and labor, supporting the applicability of the $\sigma > 1$ case to contemporary climate-induced labor productivity shocks.

³⁷This labor-biased damage specification captures the direct physiological effects of heat stress on worker productivity, consistent with emerging evidence on heat-induced productivity losses in exposed sectors ([Graff Zivin and Neidell, 2014](#)).

Intermediate varieties and pricing. Each variety i in sector $x \in \{K, L\}$ is produced at unit cost $c_{xi,t} = \psi/A_{xi,t}$ (normalizing $\psi = \alpha^2$). With Dixit–Stiglitz preferences over varieties (demand elasticity $1/(1 - \alpha)$), monopolistic producers charge a constant markup $p_{xi,t} = \psi/(\alpha A_{xi,t})$. Cost minimization within each sectoral bundle yields variety demand

$$x_{xi,t} = [\alpha p_{xi,t}]^{-\frac{1}{1-\alpha}} A_{xi,t} F_{x,t}, \quad F_{K,t} = K_t, \quad F_{L,t} = \textcolor{red}{D}(T_t) L_t, \quad (\text{A11})$$

where $F_{x,t}$ represents the effective factor supply in sector x . Crucially, heat-induced reductions in labor efficiency $\textcolor{red}{D}(T_t)$ directly shrink the effective labor supply $F_{L,t}$, thereby reducing demand for labor-complementary machines in the L -sector: firms employing fewer effective workers require fewer labor-augmenting technologies. This channel creates the market-size mechanism for directed technical change. Under symmetry, the sectoral outputs satisfy

$$Y_{x,t} = [\alpha P_{x,t}]^{-\frac{\alpha}{1-\alpha}} A_{x,t} F_{x,t}, \quad x \in \{K, L\}, \quad (\text{A12})$$

where $P_{x,t}$ denotes the sectoral price index. The dual cost structure of (A7) implies the standard within-CES relative price condition

$$\frac{P_{K,t}}{P_{L,t}} = \frac{\alpha}{1 - \alpha} \left(\frac{Y_{K,t}}{Y_{L,t}} \right)^{-1/\sigma}, \quad (\text{A13})$$

which governs the relative demand for the two bundles.

A4.2 Direction of Innovation

Expected profit from a successful innovation in sector x is proportional to the product of the arrival rate, next-period quality, and the effective market size (factor supply times sectoral price adjusted for demand elasticity). Equating expected returns across sectors and using (A13) to eliminate relative prices yields the scientist indifference condition:

$$\frac{s_{K,t}}{s_{L,t}} = \left[\frac{1 + \gamma \eta_L (1 - s_{K,t})}{1 + \gamma \eta_K s_{K,t}} \cdot \frac{A_{K,t-1}}{A_{L,t-1}} \cdot \frac{K_t}{\textcolor{red}{D}(T_t) L_t} \right]^{-\varphi}, \quad \varphi = (1 - \alpha)(1 - \sigma). \quad (\text{A14})$$

This indifference locus governs the allocation of research effort. The key parameter is $\varphi = (1 - \alpha)(1 - \sigma)$: when $\sigma > 1$, I have $\varphi < 0$ and hence the exponent $-\varphi$ is positive. Heat exposure that lowers $\textcolor{red}{D}(T_t)$ raises the factor ratio $K_t/[\textcolor{red}{D}(T_t) L_t]$ —capital becomes relatively abundant compared to heat-impaired labor—thereby increasing the bracketed term and, given the positive exponent, tilting research allocation toward capital-augmenting technologies.

The economic intuition operates through a market-size effect. When heat reduces effective labor supply $\textcolor{red}{D}(T_t) L_t$ while leaving capital K_t undamaged, the K -sector enjoys a relatively larger effective factor base. Because machines are complementary to their respective factors in production (A8)–(A9), this asymmetric damage creates stronger demand for capital-augmenting technologies: firms shift toward capital-intensive production methods that are less vulnerable to heat-induced labor productivity losses. When factors are substitutes at the aggregate level ($\sigma > 1$), this market-size channel dominates any opposing price effects. The final-good sector reallocates toward the capital-intensive bundle (equations A7 and A13), amplifying demand for K -augmenting machines and making innovations that economize on heat-vulnerable labor more profitable. Scientists respond by directing effort toward labor-saving research.

Formally, equation (A14) characterizes how innovation incentives responds to climate shocks. When $\sigma > 1$, reductions in labor efficiency tilt research effort toward capital-augmenting technologies, with the magnitude of reallocation depending on the factor substitutability $\sigma - 1$ and the interiority of the research allocation $s_K^*(1 - s_K^*)$. This formalizes the empirical observation that industries with greater exposure to labor-biased heat shocks exhibit more patenting in labor-saving domains.

A4.3 Innovation as a Buffer Against Direct Damage

The preceding analysis shows that heat shocks redirect innovation toward labor-saving technologies. I now demonstrate how this endogenous response partially offsets the direct productivity damage from reduced labor efficiency. Consider the log-differential of aggregate output (A7). Define the usage-based weights $\omega_{K,t} = \alpha(Y_{K,t}/Y_t)^\rho$ and $\omega_{L,t} = (1 - \alpha)(Y_{L,t}/Y_t)^\rho$, which sum to unity and measure each bundle's contribution to final output. Standard CES algebra yields

$$d \ln Y_t = \omega_{K,t} d \ln Y_{K,t} + \omega_{L,t} d \ln Y_{L,t}. \quad (\text{A15})$$

Differentiating the sectoral production functions (A8)–(A9) and using the equilibrium pricing structure gives

$$d \ln Y_{K,t} = d \ln K_t + \frac{1}{1 - \alpha} d \ln A_{K,t} + (\text{price terms}), \quad (\text{A16})$$

$$d \ln Y_{L,t} = d \ln [D(T_t)L_t] + \frac{1}{1 - \alpha} d \ln A_{L,t} + (\text{price terms}). \quad (\text{A17})$$

The coefficient $1/(1 - \alpha)$ on quality improvements reflects how innovations amplify output through the monopolistic competition structure: each quality improvement raises not only direct productivity but also the effective variety of machines available, as seen from (A12). The price terms reflect endogenous sectoral reallocation through (A13) and can be absorbed into second-order adjustments. Substituting (A16)–(A17) into (A15) and isolating the climate and technology components yields

$$\Delta \ln Y_t \approx \underbrace{\omega_{L,t} \Delta \ln D(T_t)}_{\text{Direct heat damage}} + \underbrace{\frac{\omega_{K,t}}{1 - \alpha} \Delta \ln A_{K,t} + \frac{\omega_{L,t}}{1 - \alpha} \Delta \ln A_{L,t}}_{\text{Technology contribution}} + (\text{factors}), \quad (\text{A18})$$

where the factor terms collect $d \ln K_t$ and $d \ln L_t$.

Define the composite technology index

$$\mathcal{G}_t \equiv \frac{\omega_{K,t}}{1 - \alpha} \ln A_{K,t} + \frac{\omega_{L,t}}{1 - \alpha} \ln A_{L,t}. \quad (\text{A19})$$

Then the technology contribution in (A18) simplifies to $\Delta \mathcal{G}_t$ (up to second-order terms). Equation (A18) decomposes productivity growth into a direct negative effect from the labor-efficiency loss—captured by $\omega_{L,t} \Delta \ln D(T_t) < 0$ when heat reduces D —and a positive offsetting term $\Delta \mathcal{G}_t > 0$ driven by quality improvements from innovation.

The key insight is that $\mathcal{G}_t$ embeds the innovation response from (A10): when scientists reallocate toward capital-augmenting research in response to heat shocks (as shown in equation A14), the resulting improvements in $A_{K,t}$ enter the composite index $\mathcal{G}_t$ and boost aggregate productivity. In the context of heat exposure, this means that innovations enabling firms to substitute capital for heat-vulnerable labor partially compensate for the direct productivity losses from reduced labor efficiency. Quantitatively, the magnitude

of $\Delta\mathcal{G}_t$ depends on the innovation response encoded in equation (A14) and the usage weight $\omega_{K,t}/(1-\alpha)$ reflecting how much the economy relies on capital-intensive production.

In short, the decomposition (A18) formalizes the mitigation mechanism: directed technical change allows the economy to adapt to labor-biased climate shocks by accumulating compensating technologies, thereby attenuating the net productivity loss relative to a counterfactual without innovation.

A5 Additional Figures and Tables

A5.1 Temperature and Firm Productivity

Table A9: Effect of Temperature Bins on Productivity

	TFP (1)
<30	-0.0191 (0.0382)
30-40	0.0470*** (0.0106)
40-50	0.0005 (0.0211)
50-60	0.0158*** (0.0046)
70-80	0.0261*** (0.0084)
80-90	0.0942*** (0.0342)
>90	0.0497 (0.0347)
Labor intensity	-0.3279*** (0.0794)
<30 $\times$ Labor intensity	0.0425 (0.0846)
30-40 $\times$ Labor intensity	-0.1099*** (0.0247)
40-50 $\times$ Labor intensity	-0.0046 (0.0491)
50-60 $\times$ Labor intensity	-0.0369** (0.0145)
70-80 $\times$ Labor intensity	-0.0592*** (0.0207)
80-90 $\times$ Labor intensity	-0.2196*** (0.0652)
>90 $\times$ Labor intensity	-0.1108* (0.0665)
Observations	9,064,059
Firm FE	✓
Country $\times$ Year FE	✓
Sector $\times$ Year FE	✓

Notes: Estimating equation uses firm FE, country-year FE, and sector-year FE. Standard errors clustered by firm and country-year. *** p<0.01, ** p<0.05, * p<0.1.

Table A10: CDD on Productivity

Labor Intensity Group	TFP	
	Low (1)	High (2)
CDD	-0.0001*** (5.35×10^{-5})	-0.0004* (0.0002)
Precipitation	-0.0001 (0.0001)	-4.63×10^{-5} (0.0002)
Observations	6,534,983	2,178,768
Firm FE	✓	✓
Country×Year FE	✓	✓
Sector×Year FE	✓	✓

Notes: Estimating equation uses firm FE, country-year FE, and sector-year FE. Standard errors clustered by firm and country-year. CDD measures Cooling Degree Days above 85. *** p<0.01, ** p<0.05, * p<0.1.

A5.2 Heat-Induced Market Share Dynamics

Table A11: Market Share and Temperature Bins: Interaction with Labor Intensity

	Market Share (log) (1)
<30	-0.0749* (0.0421)
30-40	-0.0521*** (0.0189)
40-50	-0.0687 (0.0427)
50-60	-0.0541*** (0.0187)
70-80	0.0004 (0.0124)
80-90	-0.0374* (0.0201)
>90	0.0640*** (0.0187)
Labor intensity	0.2170 (0.1766)
<30 × Labor intensity	0.0817 (0.0954)
30-40 × Labor intensity	0.1354*** (0.0441)
40-50 × Labor intensity	0.2027** (0.0974)
50-60 × Labor intensity	0.1470*** (0.0434)
70-80 × Labor intensity	-0.0249 (0.0225)
80-90 × Labor intensity	0.0263 (0.0385)
>90 × Labor intensity	-0.2309*** (0.0307)
Observations	9,064,059
Firm FE	✓
Country×Year FE	✓
Sector×Year FE	✓

Notes: Dependent variable is $\ln(\text{Market Share})$. Regressors are temperature-bin indicators (baseline $[60, 70)^\circ F$), the continuous labor-intensity index, and bin×labor-intensity interactions; precipitation is included as a control. Fixed effects: firm; country–year; sector–year. Markets are defined by country-NACE-4 industries. Standard errors are two-way clustered by market and country–year. Temperature-bin main effects and interactions are expressed in percent; in the hottest bin the interaction is negative and precisely estimated, indicating that higher labor intensity amplifies heat-induced share losses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A5.3 Technology Adjustment: Output Elasticities

Table A12: Output Elasticities (Capital and Labor) vs Extreme Degree-Days ($> 85^{\circ}\text{F}$) — Firm Level

	log(Capital Output Elasticity) (1)	log(Labor Output Elasticity) (2)
CDD	0.0045** (0.0018)	-0.0009*** (0.0003)
Precipitation	-0.0039 (0.0028)	0.0005** (0.0002)
Observations	9,064,059	9,064,059
Firm FE	✓	✓
Sector×Year FE	✓	✓
Country×Year FE	✓	✓

Notes: Replacing *TMAX* bins with extreme degree-days ($> 85^{\circ}\text{F}$) delivers the same directional pattern: capital OE rises with heat, labor OE falls. Same fixed effects and clustering as Table 8; coefficients scaled $\times 100$. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A5.3.1 Factor Ratio

Table A13: Effect of daily maximum temperature on log capital-labor ratio

	Capital/Labor Costs (1)
<30	-0.0184 (0.0176)
30-40	0.0186 (0.0163)
40-50	0.0160** (0.0082)
50-60	0.0009 (0.0090)
70-80	0.0042 (0.0081)
80-90	0.0357* (0.0183)
>90	0.0922*** (0.0348)
Precipitation	-0.0026 (0.0028)
Observations	9,064,059
Firm FE	✓
Sector×Year FE	✓
Country×Year FE	✓

Notes: Dependent variable is log capital-labor ratio. Coefficients measure the effect of one additional day in each temperature bin relative to $[60, 70)^{\circ}F$. Specification includes firm fixed effects, country-year fixed effects, and industry-year fixed effects. Standard errors clustered at firm and country-year level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A5.3.2 Industry-Level Technology Response

To complement the firm-level estimates in Figure 6, I aggregate output elasticities to the *industry* level and re-estimate the same bin specification. Let i index firms in industry I (NACE2), country c , and year t . Define the sales-share average elasticity,

$$\overline{OE}_{Ict}^m \equiv \sum_{i \in (I,c,t)} w_{ict} OE_{it}^m, \quad w_{ict} = \frac{\text{SALES}_{ict}}{\sum_{j \in (I,c,t)} \text{SALES}_{jct}}, \quad m \in \{K, L\}, \quad (\text{A20})$$


and then estimate the analogue of (6) with $\log \overline{OE}_{Ict}^m$ on the left-hand side, the same temperature-bin regressors, precipitation, and the same high-dimensional fixed effects as in the firm specification—country $\times$ year and sector $\times$ year—clustering standard errors by country–industry and year. Coefficients are scaled $\times 100$.

By construction, $\overline{OE}_{Ict}^m$ embeds (i) *within-firm technology adjustment* and (ii) *between-firm reallocation* via w_{ict} . A first-order change in the log industry elasticity can be written (Olley–Pakes style) as

$$\Delta \log \overline{OE}_{Ict}^m \approx \underbrace{\sum_i w_{ict} \Delta \log OE_{it}^m}_{\text{within-firm}} + \underbrace{\sum_i \Delta w_{ict} (\log OE_{it}^m - \log \overline{OE}_{Ict}^m)}_{\text{between-firm reallocation}}, \quad (\text{A21})$$

so any amplification at the industry level relative to firm estimates reflects reallocations that favor producers with higher capital elasticities.

Figure A3: Effect of $TMAX$ bins on industry average capital elasticity



Notes: Each panel reports semi-elasticities (coefficients $\times 100$) from bin regressions of *log* industry-level elasticities on counts of daily maximum temperature bins, omitting $[60, 70)^\circ F$. Specifications include country $\times$ year and sector $\times$ year fixed effects; standard errors clustered by country–industry and year. Panel (a) uses the sales-share weighted industry average; Panel (b) splits industries by ex ante labor-intensity (least vs. most labor-intensive). Shaded bands show 95% confidence intervals.

Panel (a) of Figure A3 shows that relative to $[60, 70)^\circ F$, additional *very hot* days ($\geq 90^\circ F$) are associated with higher industry-average *capital* elasticity. The point estimates are economically meaningful: firm-level semi-elasticities in Figure 6 are on the order of 0.05, whereas industry-level coefficients in Figure A3 approach 1.0. The gap is consistent with substantial *between-firm* reallocation toward more capital-intensive producers when heat shocks strike.

Panel A3(b) sharpens this interpretation. The amplification is concentrated in *ex ante* labor-intensive industries, which display steep increases in capital elasticity at high temperatures; by contrast, capital-

intensive industries show muted responses ³⁸ This pattern matches Facts 1 and 2 (3.1–3.2): labor-intensive firms suffer larger heat-induced productivity losses and lose market share to more capital-intensive peers, raising the industry’s average capital elasticity even if within-firm adjustments are modest.

³⁸Examples of least labor-intensive industries: Extraction of crude petroleum and natural gas (06), Manufacture of coke and refined petroleum products (19), and Manufacture of basic chemicals (20); Examples of most labor-intensive industries: Mining of metal ores (07), Manufacture of leather and related products (15), and Civil engineering (42).

A5.4 Directed Innovation: Patent Evidence

A5.4.1 Alternative Patent Measure: Count vs Share

Table A14: Effects of Degree Days on Patents (Region–Industry, Poisson, Lag-2)

	Labor-saving patents (1)	Labor-saving patent share (2)
Lag2 Degree Days ($\geq 85F$)	-0.0069*** (0.0016)	-0.0064*** (0.0025)
Labor Elasticity	-2.797*** (0.1440)	-5.267*** (0.2868)
Lag2 Degree Days ($\geq 85F$) $\times$ Labor Elasticity	0.0144*** (0.0034)	0.0104* (0.0054)
Observations	76,400	76,400
Region $\times$ Sector FE	✓	✓
Priority Year FE	✓	✓

Notes: Poisson at the region–industry–year level with region-sector and year fixed effects; standard errors clustered by region-year and industry. PPML coefficients are semi-elasticities. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A5.4.2 Lag Robustness

Table A15: Lag Robustness—CDD $\times$ LI (Labor-Saving DV, Region–Industry)

	Labor-saving patents			
	(1)	(2)	(3)	(4)
Lag2 Degree Days ($\geq 85F$)	-0.0069*** (0.0016)			
Labor Elasticity	-2.797*** (0.1440)	-2.771*** (0.1440)	-2.720*** (0.1442)	-2.711*** (0.1446)
Lag2 Degree Days ($\geq 85F$) $\times$ Labor Elasticity	0.0144*** (0.0034)			
Lag3 Degree Days ($\geq 85F$)		-0.0058*** (0.0015)		
Lag3 Degree Days ($\geq 85F$) $\times$ Labor Elasticity		0.0133*** (0.0034)		
Lag4 Degree Days ($\geq 85F$)			-0.0052*** (0.0016)	
Lag4 Degree Days ($\geq 85F$) $\times$ Labor Elasticity			0.0120*** (0.0034)	
Lag5 Degree Days ($\geq 85F$)				-0.0049*** (0.0016)
Lag5 Degree Days ($\geq 85F$) $\times$ Labor Elasticity				0.0120*** (0.0034)
Observations	76,400	75,840	75,295	74,691
Region $\times$ Sector FE	✓	✓	✓	✓
Priority Year FE	✓	✓	✓	✓

Notes: Each column is a separate Poisson regression with CDD lag $K \in \{2, 3, 4, 5\}$ interacted with LI. All specifications include region-sector and year fixed effects; standard errors clustered by region-year and industry. The interaction coefficient remains positive and significant across lags. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A5.4.3 Long-Difference Specifications

Table A16: Heat and Labor-Saving Innovation: Long Difference (Region–Industry)

	Labor-saving patents(2000–2020) (1)
$\Delta DegreeDays(2000 - -2020)$	-0.0125*** (0.0023)
Labor Elasticity	-2.011*** (0.1693)
$\Delta DegreeDays(2000 - -2020) \times$ Labor Elasticity	0.0268*** (0.0051)
Observations	3,392
Region FE	✓
Industry FE	✓

Notes: This table reports long-difference estimates at the region-industry level. The dependent variable is the change in labor-saving patents from 2000 to 2020. $\Delta DegreeDays(2000 - 2020)$ measures the change in cumulative degree days over the period. Labor Elasticity is measured at the industry level in the initial period (2000). Region and industry fixed effects control for differential regional and sectoral growth trends, respectively. Standard errors clustered at the region level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

A5.5 Damage Mitigation from Innovation

Table A17: Firm-level mitigation: TFP on CDD $\times$ LS count (lag = 2)

	log TFP (1)
CDD	-0.0005** (0.0002)
LS count (t-k)	-0.3508*** (0.1165)
CDD $\times$ LS count (t-k)	0.0005*** (0.0001)
Observations	6,096,992
Firm FE	✓
Country $\times$ Year FE	✓
Sector $\times$ Year FE	✓

Notes: Outcome: log TFP. Temperature exposure: CDD defined as cumulative degree days above 85 °F. LS count measures the number of LS patents filed by the firm in the preceding two years. All coefficients expressed in percentage terms. Fixed effects: firm, country $\times$ year, and sector $\times$ year. Standard errors clustered by firm and country-year. Precipitation included when available. *** p<0.01, ** p<0.05, * p<0.1.