Generalist Inventors and Induced Innovation*

Lior Shabtai§ Yudan Ying‡

November 14, 2026

[Click Here For Most Recent Version]

Abstract

This paper provides evidence that workers with broader experience play a key role in firms' response to changing investment opportunities. Using patent records on 68,000 inventors at U.S. manufacturing firms (1992–2021) and a reduced-form shift-share design, we find that broader inventors increase green patenting more than narrower inventors following energy-cost shocks that raise firms' returns to green innovation. The induced green patents are more likely to involve new collaborators and are less scientifically important than these inventors' usual green patents, while their economic value shows no change. The findings are not explained by prior green experience, lower observed inventor ability or team match quality, cross-firm mobility, or technological distance from prior work. The evidence suggests that firms draw on versatile inventors to undertake projects that become viable after energy-cost shocks.

*We are grateful to Efraim Benmelech, Nicolas Crouzet, Carola Frydman, Diego Känzig, Joel Mokyr, and Dimitris Papanikolaou for their guidance and support. We also thank Vivek Bhattacharya, Nicola Bianchi, Sean Higgins, Ryan Hill, Sarit Markovich, David Matsa, Filippo Mezzanotti, Michael Powell, Bryan Seegmiller, Jeroen Swinkels, Lulu Wang, and Yishay Yafeh, as well as seminar participants at Kellogg School of Management. Lior Shabtai gratefully acknowledges support from the Crown Family Israel Center for Innovation at Kellogg School of Management.

§Northwestern University; lior.shabtai@kellogg.northwestern.edu.

‡Northwestern University; yudan.ying@kellogg.northwestern.edu.

1 Introduction

How do firms use their workforce to respond to changing investment opportunities? Firms reallocate workers internally toward activities whose returns have risen (Giroud and Mueller, 2015; Tate and Yang, 2015; Cestone et al., 2024), and a firm's ability to do so may depend on its workers' ability to move between jobs. Broader experience across activities may make workers more versatile and better equipped for such moves. While a large literature examines how the generality of workers' skills affects their mobility across firms and occupations (Becker, 1962; Lazear, 2009; Gathmann and Schönberg, 2010; Frydman, 2019), prior research has not established how workers' breadth of experience shapes their internal redeployment and subsequent output quality.

Innovation offers a setting well suited for studying this question. First, patent records link inventors to firms and allow us to observe their patent output and the scientific and economic value of their inventions. Second, patent histories capture inventors' prior work in specific areas and the breadth of their technical experience. Third, patent records identify the technological field of each invention, and we can therefore examine how firms draw on inventors with different experience when the returns to developing particular technologies change. Specifically, we use energy-cost shocks to obtain exogenous variation in firms' returns to green innovation.

In this paper, we study whether broader inventors increase green patenting more when their firms' energy costs rise, and how the quality of their green innovation changes. We construct a panel linking the patenting histories of 68,000 inventors to about 1,200 U.S. Compustat manufacturing firms from 1992 to 2021. An inventor's breadth of technical knowledge and experience captures how dispersed their prior patents are across technology fields. Comparing the innovation outcomes of inventors who patent with the same firm in the same year, we find that broader inventors play a key role in firms' response to changing investment opportunities: the green-patenting response to energy-cost shocks is stronger among them, but the scientific value of their green patents declines relative to their usual work. The economic value of induced green patents, by contrast, shows no change.

Expansion into green innovation may reflect higher returns to developing such technologies or changes in firms' and their inventors' innovative capacity. To study how firms respond to changing investment opportunities, we need variation in returns that is orthogonal to changes in their inventive capabilities. We employ a reduced-form shift-share design that combines firms' predetermined fuel-cost exposures with exogenous

fuel-specific price changes to construct firm-level energy-cost shocks. The identifying variation comes from the shifts: exogenous price changes generated by oil supply news shocks identified by Känzig (2021). These shocks affect both oil and natural gas prices, with no pass-through to coal or electricity prices. The shares are firms' predetermined fuel-cost exposures, constructed by combining firms' sales shares across the industries in which they operate with industry-level energy intensity and fuel use.[1] These predetermined exposures translate the common price shifts into differential changes in firms' energy costs. We first show that higher energy costs increase firms' green patenting: a one-standard-deviation shock, corresponding to a 5% increase in energy costs for the median firm, raises firms' green patenting by 3.7%, with no change in their other patenting activity.

At the inventor level, we measure exposure to these shocks through the firms with which inventors patent. We use firm-by-year fixed effects to compare inventors facing the same firm-level shock to innovation incentives and to absorb other factors that vary across firms and years and may affect innovation outcomes. We include inventor fixed effects to absorb persistent differences across inventors, such as prior education, innate ability, and other fixed determinants of performance. We also control for inventors' green patent share, lagged patent stock, and career age.

We start with the quantity of induced green innovation. If broader experience makes inventors more versatile, firms may draw more heavily on them when the returns to green innovation rise. In line with this prediction, a one-standard-deviation increase in the energy-cost shock raises green patenting by 2% more for an inventor who is one standard deviation broader.[2] There is no corresponding differential increase in non-green patenting, indicating that the response does not reflect a general increase in inventive activity. To assess whether firms reassign broader inventors to the induced green projects, we examine whether their additional green patents involve new co-inventors. To the extent that reassignment to new projects changes the colleagues with whom an inventor works, patents involving new co-inventors are informative about this margin. A one-standard-deviation increase in the shock raises the probability that a broader inventor files a green patent with a new co-inventor by 1.8% of its mean, with no corresponding

[1] We fix industry energy intensity and fuel use at their 1991 values and each firm's segment mix at its first observed composition.

[2] Breadth of technical knowledge and experience is measured as one minus the Herfindahl–Hirschman index (HHI) of the technology-field shares in an inventor's prior patent portfolio. For illustration, moving from prior patenting evenly distributed across two technology fields to four increases breadth from 0.50 to 0.75, equivalent to approximately 1.2 standard deviations.

change for green patents without a new co-inventor.

We consider two dimensions of patent value: scientific importance and economic value. Scientific importance is measured using the patent-text-based measure developed by Kelly et al. (2021). It captures both novelty relative to earlier patents and influence on subsequent innovation. Following Kelly et al. (2021), we define patents in the top decile of this measure within their industry and filing year as *breakthroughs*. Separating breakthroughs from other patents is closely related to Mokyr's (1990) distinction between macroinventions, which introduce radical ideas without clear precedent, and microinventions, which improve existing technologies and reduce costs. This distinction is especially relevant in our setting because microinventions are more likely than macroinventions to respond to changes in prices and incentives (Mokyr, 1990). To measure economic value to the firm, we use the estimates of Kogan et al. (2017), based on firms' stock-market reactions to patent grants.

We next study how energy-cost shocks affect the quality of green innovation. A one-standard-deviation shock lowers the breakthrough share of green patents by 7.1% of its mean for an inventor who is one standard deviation broader, with no change in the economic value of these patents. Across the entire distribution of scientific importance, the decline in the share of top-decile patents (i.e., breakthroughs) is offset by an increase in the share of patents between the 50th and 70th percentiles, while the share in the bottom half shows no estimated change. Thus, the decline in scientific value reflects a step down from breakthroughs rather than a general shift toward the lower end of the scientific importance distribution. This decline contrasts with the cross-sectional pattern, where we show that broader inventors' patents are more scientifically important than those of narrower inventors.

Dynamic specifications show that both quantity and quality effects concentrate one year after the shock, with no evidence of pre-trends in a pre-shock placebo year and smaller, less-precise estimates over the subsequent three years. The findings remain qualitatively unchanged across alternative measures of breadth, shock horizons, exposure definitions, quality measures, and econometric specifications, and after accounting for co-inventor team size and composition.

We interpret these findings as firms assigning broader inventors to green projects that become viable following the shock but are of lower scientific value than those inventors' usual work. Higher energy costs raise the returns to green innovation, making projects worthwhile that firms would not previously have pursued. Broader inventors' versatil-

ity allows firms to draw on them to staff these additional projects. Although we do not observe assignments directly, our sample inventors patent almost exclusively with firms, where inventive activity is typically directed through corporate R&D projects (Nagaoka and Walsh, 2009).[3] Furthermore, the concentration of additional green patents in work with new collaborators supports this interpretation. The quality results indicate that this expansion involves work less likely to produce breakthroughs than broader inventors' usual green output. The shift toward upper-middle-tier patents, with no estimated increase in the bottom half of the distribution, is consistent with firms pursuing additional projects near their investment threshold: projects that become profitable after the shock, but are somewhat less scientifically important than those previously undertaken. The estimated effect on the economic value of these patents is small and statistically insignificant. Thus, while of lower scientific value, these projects may still be valuable to the firm by improving existing production technologies or reducing energy costs.

We consider five alternative explanations and find no evidence that the results reflect observably weaker inventor-team matches, prior green experience, observed inventor ability, cross-firm mobility, or larger technological pivots. First, because the induced patents involve new collaborators, the decline could reflect weaker inventor matches rather than weaker projects. We find no evidence for that: predetermined team characteristics predict no decline in broader inventors' breakthrough probabilities after the shock, while their actual breakthrough share falls. Second, the findings could be explained by broader inventors having more prior green experience. Across several measures, interacting prior green experience with the shock leaves the breakthrough-share decline unchanged and the quantity effect positive and statistically significant, though slightly attenuated. Third, selection on inventor ability could explain the findings. Interactions with predetermined education measures from résumé data, available for 43.5% of the inventor panel, leave both effects unchanged, with no evidence of selection on lower observed ability. Fourth, while the shocks may induce broader inventors to move to firms with lower-quality green opportunities rather than be reassigned internally, the results persist among inventors who patent with only one firm throughout their career. Finally, the decline could reflect a *pivot penalty*: scientific value may fall when inventors work in fields farther from their prior knowledge (Hill et al., 2025). Accounting for technological distance between new green patents and the inventor's prior patenting leaves the decline in

[3]Over their observed careers, 85.8 percent of our sample inventors patent exclusively with firms, and about 95 percent of their patents are assigned to firms. Section 2.5 provides further details on their corporate affiliation.

breakthrough share unchanged.

Taken together, our results show that broader inventors increase green patenting more following energy-cost shocks, but their additional patents are less scientifically important than their usual green output. We interpret these findings as firms assigning versatile workers to newly viable projects, consistent with an expansion of microinventions that improve existing technologies and reduce costs (Mokyr, 1990). Workers with broader experience thus play a key role in how firms respond to changing investment opportunities.

We contribute to several strands of literature. First, our findings relate to research on the allocation of resources within firms. Firms and business groups reallocate capital and labor internally toward units with stronger opportunities (Stein, 1997; Maksimovic and Phillips, 2002; Giroud and Mueller, 2015; Cestone et al., 2024). Within internal labor markets (Doeringer and Piore, 1985), workers move through internal job changes and promotions (Baker et al., 1994; Lazear and Oyer, 2004; Benson et al., 2019; Huitfeldt et al., 2023; Minni, 2026) and are redeployed across units and tasks in response to shocks (Tate and Yang, 2015; Cestone et al., 2016; Huneeus et al., 2021; Adhvaryu et al., 2022). In the context of innovation, firms direct inventors across projects and units (Conti et al., 2014; Stadler et al., 2022), and Kim et al. (2024) find a positive correlation between inventor breadth and inter-unit redeployment. While this literature documents reallocation and its consequences, it provides limited evidence on how workers' prior breadth shapes the quantity and quality of output induced by changing opportunities. Our findings highlight the role of broader inventors in firms' response to higher innovation returns and are consistent with firms redeploying them to newly viable projects of lower scientific value than these inventors' usual work.

Second, we contribute to the literature on general skills and the structure of human capital (Becker, 1962). It examines the specificity and portability of experience across industries, occupations, and tasks (Neal, 1995; Poletaev and Robinson, 2008; Kambourov and Manovskii, 2009; Gathmann and Schönberg, 2010), and models how skill portfolios shape job assignment (Gibbons and Waldman, 2004; Lazear, 2009; Yamaguchi, 2012). Broader skill portfolios are valuable when work requires multiple capabilities or occupational fit is uncertain (Lazear, 2005; Silos and Smith, 2015), while workers face larger wage penalties when their occupations are farther from their specialized training (Eckardt, 2026). Among managers, broader experience and more general skills are associated with higher pay and mobility (Custódio et al., 2013; Frydman, 2019). Much of this work assesses skill portfolios through wages and career outcomes. Instead, we examine how

prior experience shapes the quantity and quality of work following changes in firms' investment opportunities, highlighting worker versatility as a source of flexibility for firms.

Third, we contribute to research on the production of new knowledge. This literature studies how recombination and prior experience shape innovation (Weitzman, 1996, 1998; Fleming, 2001; Fleming and Sorenson, 2004; Kaplan and Vakili, 2015; Arts and Fleming, 2018; Jones, 2023; Hill et al., 2025), how demand affects the amount and composition of innovative activity (Acemoglu and Linn, 2004; Finkelstein, 2004; Dranove et al., 2022), how financing affects the novelty of drug development (Krieger et al., 2022), how the effect of inventor rewards on patent inventiveness varies with inventors' specialization (He et al., 2026), and how competition affects the quality of scientific research (Hill and Stein, 2025). We add evidence on internal labor allocation as a channel linking higher innovation returns to changes in output quality. Our results suggest that firms draw on workers' versatility to expand innovation in newly profitable activities, but assign them to projects less likely to produce breakthroughs.

Finally, we contribute to the literature on induced and directed innovation in energy-saving and climate-mitigation technologies. Prices and policy redirect innovation toward energy-saving and clean technologies (Newell et al., 1999; Popp, 2002; Johnstone et al., 2010; Acemoglu et al., 2012, 2016, 2023; Calel and Dechezleprêtre, 2016). These responses are path dependent, with firms tending to innovate in technologies in which they have accumulated experience (Aghion et al., 2016, 2024), and induced clean innovation is concentrated among inventors already active in the domain (Dugoua and Gerarden, 2025). We show that induced green innovation depends not only on inventors' prior green experience, but also on the breadth of their experience across technologies.

The remainder of the paper proceeds as follows. Section 2 presents the data, innovation-quality measures, breadth measure, and firm-level energy-cost shock, followed by sample construction and summary statistics. Section 3 presents the empirical design and firm-level response to energy-cost shocks. Section 4 presents the inventor-level results. Section 5 discusses mechanisms and alternative explanations. Section 6 concludes.

2 Data and Measurement

This section describes the data sources and estimation sample, the measures of innovation quality and prior knowledge, and the construction of the firm-level energy-cost shock.

Appendix B provides additional construction details and conventions.

2.1 Data sources and sample construction

Firm data. Compustat North America provides firm-level financial data, and Compustat Business Segments provides information on firms' SIC-coded business lines and segment sales. We use segment sales to weight industry-level energy use when constructing each firm's exposure to changes in oil and natural gas prices. Our estimation sample consists of publicly listed US firms with manufacturing activity and covers 1992 to 2021.

Energy data. We combine data from the Manufacturing Energy Consumption Survey (MECS) and the NBER-CES Manufacturing Industry Database to measure industry energy use. The 1991 MECS reports energy input use at a substantially more granular industry level than earlier waves.[4] We therefore begin the estimation window in 1992, the year after its release. We convert the reported fuel quantities into expenditure shares using EIA prices for each energy input. The NBER-CES database provides energy expenditure as a share of industry sales for four-digit SIC industries, which we also measure in 1991. We combine 1991 MECS industry-level energy input shares and 1991 NBER-CES industry-level energy intensity with firms' segment sales to construct firm-level exposure to changes in energy input costs. We use monthly oil prices and the EIA Citygate natural-gas price,[5] and coal and industrial-electric-power producer price indexes from the BLS. Appendix B reports the industry coverage of each source and the matching procedure, and Section 2.4 describes how we combine these inputs with the oil supply news shocks.

Patent data. Patent data are obtained from PatentsView and cover patents filed with and granted by the US Patent and Trademark Office (USPTO). The data report patent numbers, filing dates, citations, disambiguated inventor identities and assignees, and technology classifications. We use technology classifications from the Cooperative Patent Classification (CPC), the joint USPTO–EPO system for classifying patent technologies. We

[4]The number of reported SIC categories increased from 30 in 1988 (20 two-digit groups and 10 four-digit industries) to 62 in 1991 (20 two-digit groups, 2 three-digit groups, and 40 four-digit industries).

[5]The EIA Citygate price measures the price that local distribution companies and utilities pay for gas from the pipeline transmission system and covers the full 1992–2021 analysis window. The alternative Henry Hub spot-price benchmark, the wholesale price of gas traded at the Henry Hub in Louisiana, begins in 1997. The two series move closely together: the correlation between their monthly real prices is 0.95 over their 1997–2021 overlap. Both delivered and hub prices are commonly used in research on natural-gas markets and price pass-through (Avalos et al., 2016; Fowlie et al., 2018; Hausman, 2019).

define a patent as green if it receives a CPC Y02 classification. Y02 covers technologies tagged for climate-change mitigation or adaptation and is widely used to measure green innovation (e.g., Calel and Dechezleprêtre, 2016); Appendix Table A1 reports results using alternative definitions. We merge two sets of patent-quality measures by patent number: the text-based scientific-value measures of Kelly et al. (2021), hereafter KPST, and the economic-value measures of Kogan et al. (2017), hereafter KPSS. KPSS assigns patents to CRSP firms, which we map to Compustat identifiers and use to link patenting activity to firms. Subsection 2.2 describes these measures.

Résumé data. Revelio Labs provides individual-level résumé data compiled from public online professional profiles, including employment histories and education. We link inventors to Revelio employees at their sample firms using a firm-level crosswalk and name matching, validating matches against employment-spell overlap when dates are available.[6] The matched data cover 29,864 of the 68,606 inventors in the main green-count regression (43.5 percent). We use these data to construct predetermined proxies for human capital and ability in tests of alternative explanations to our results.

2.2 Innovation quality measures

We measure two dimensions of patent quality: scientific value, from patent text and citations, and economic value to the firm, from stock-market reactions to patent grants.

KPST measures of the scientific value of innovation. Kelly et al. (2021) measure the scientific value of innovation using patent text. They compare the text of each patent with that of patents filed in the years before and after it. Backward similarity measures similarity to existing knowledge, and forward similarity measures the similarity of subsequent innovation to the focal patent. A scientifically important patent departs from existing knowledge and influences subsequent innovation. KPST measure importance as the log ratio of forward to backward similarity. We use their five-year measures for patents filed through 2016, the last filing cohort with a forward window.

We assess a patent's relative scientific value by comparing it with patents filed in the same year and technological area.[7] This accounts for differences in patent language, tech-

[6] Appendix B.6 describes the matching procedure.

[7] We use the Algorithmic Links with Probabilities (ALP) CPC-to-NAICS crosswalk of Goldschlag et al. (2020) to assign each patent to a four-digit NAICS industry.

nological development, and patenting intensity across industries and cohorts. Following KPST, we define a breakthrough as a patent in the top decile of the forward-to-backward similarity ratio within its NAICS4-by-year cell. This notion of breakthroughs, capturing both novelty and subsequent influence, is closely related to Mokyr's (1990) distinction between macroinventions, in which radical new ideas emerge without clear precedent, and microinventions, which improve, adapt, and streamline existing techniques. We also use the patent's percentile rank within its industry-year cell as a continuous measure of scientific value.

Citation measures. For each patent, we use USPTO citation records to count citing patents granted within five years of its grant. Using a fixed window avoids the mechanical truncation of lifetime citation counts for recent patents. A patent is *highly cited* when its five-year citation count reaches the top decile of its NAICS4 by filing-year cell and receives at least one citation. As a continuous intensity measure we use the log of one plus the five-year count.

KPSS measure of economic value. Kogan et al. (2017) estimate a patent's economic value from the excess stock-market return of the patent-owning firm around its grant date. We use their real-dollar patent-value measure, which is available through 2020.

2.3 Breadth of technical experience

We use inventors' patent histories to measure the breadth of their experience and knowledge across technology fields. Working across fields exposes inventors to a wider range of technical knowledge, so we interpret their prior patenting as informative about the breadth of experience and knowledge they have accumulated. Our measure is therefore conceptually similar to measures used for workers (Gathmann and Schönberg, 2010; Lise and Postel-Vinay, 2020), managers (Custódio et al., 2013; Frydman, 2019), and inventors (Jones, 2009; Arts and Fleming, 2018; Melero and Palomeras, 2015; Ma et al., 2024) that rely on prior experience.

We construct breadth using the CPC main groups assigned to an inventor's prior patents,[8] excluding class codes related to climate change mitigation or adaptation (Y02).

[8]Related work constructs patent technology vectors and firm technology portfolios at the IPC main-group level (Hegde et al., 2023; Aharonson and Schilling, 2016; Park and Yoon, 2017).

Unlike standard CPC sections, which organize technologies into underlying technical disciplines, Y02 tags identify a technology's relevance to climate change mitigation or adaptation. Excluding Y02 ensures that breadth strictly captures an inventor's diversity across substantive technical domains, rather than treating prior green experience as an additional technical discipline.[9] We allocate each patent equally across its remaining non-Y02 main groups. Let s_{ict} denote technology field c's share of inventor i's accumulated prior patenting through year t. We define breadth as one minus the Herfindahl concentration of these shares:

$$\text{Breadth}_{it} = 1 - \sum_{c} s_{ict}^2,$$

This measure has a similar form to the originality and generality measures of Hall et al. (2001). It equals zero when prior patenting is concentrated in one field and rises as it is spread more evenly across fields. Breadth enters the specifications lagged one year. Appendix C.1 shows that the results are qualitatively unchanged when past patenting is depreciated and under alternative CPC aggregation levels.

2.4 The energy-cost shock

Our measure of firm-level energy cost changes combines energy price movements induced by oil supply news shocks from Känzig (2021) with predetermined differences in firms' exposure across energy inputs:

$$\text{EnergyCostShock}_{f,t-1} = \sum_{k} Exposure_{fk} \times \widehat{\Delta p}_{k,t-1},$$

The measure has a reduced-form shift-share structure: it combines predetermined firm exposure shares with common, time-varying predicted energy-price changes (Adão et al., 2019; Goldsmith-Pinkham et al., 2020; Borusyak et al., 2025). Here, $\widehat{\Delta p}_{k,t-1}$ is the estimated real price change of energy input k attributable to oil supply news shocks, and $Exposure_{fk}$ is firm f's predetermined cost exposure to that input. As we show below, oil and natural gas prices respond to the oil supply news shock, whereas coal and electricity prices do

[9] For example, Y02P 70/10, the CPC tag for heat recovery in manufacturing, is assigned to both Kimberly-Clark's U.S. Patent 6,910,283, filed in 2003, which recovers heat from a tissue dryer's exhaust air, and Eastman Chemical's U.S. Patent 7,179,881, also filed in 2003, which transfers heat from hot to incoming PET pellets. Both inventions are designed to lower energy use in manufacturing. Yet their substantive classifications are D21F 5/00, for drying paper webs, and C08G 63/00, for solid-state polymerization of polyesters, respectively. Outside Y02, they share no classification. Including Y02 would therefore treat a common energy-saving application as shared technical knowledge.

not. The baseline measure therefore uses only the oil and natural gas components.

We first describe the oil supply news shock, explain why it provides exogenous variation relevant to our setting, and show that oil and natural gas prices respond while coal and electricity prices do not. We then discuss the two components of the energy-cost shock: the resulting energy price changes and firms' predetermined exposure to each energy input.

The oil supply news shock. We use the monthly structural shock to expected future oil supply identified by Känzig (2021), which is recovered from high-frequency changes in oil futures prices around OPEC production announcements. Känzig (2021) shows that these announcement surprises are not forecastable by macroeconomic or financial variables and are uncorrelated with measures of global demand, productivity, uncertainty, and monetary and fiscal shocks, so that the series isolates supply-driven changes in energy costs rather than price movements driven by global energy demand.

Energy price pass-through. We estimate the response of the real prices of oil, natural gas, coal, and electricity to this monthly structural shock. For each energy input k, we estimate local projections through 24 months, following Känzig (2021):

$$\Delta_h y_{k,m} = a_{k,h} + \psi_{k,h} Z_m + \phi_{k,h} y_{k,m-1} + \varepsilon_{k,m+h},$$

where $\Delta_h y_{k,m} = y_{k,m+h} - y_{k,m-1}$, $y_{k,m}$ is the log real price of energy input k in month m, Z_m is the structural oil supply news shock, and $\psi_{k,h}$ is the price response at horizon h. We estimate these projections separately for each energy input over 1992 to 2020 and use Newey-West standard errors, which account for heteroskedasticity and time-series correlation. We normalize the responses so that the shock raises the real oil price by 10 percent on impact. Figure 1 shows immediate pass-through to oil and gradual pass-through to natural gas: a shock that immediately increases oil prices by 10 percent increases natural gas prices by 6 percent at its peak, after ten to twelve months.[10] We find no response in coal or electricity prices. We therefore use only predicted changes in oil and natural gas

[10]The response of natural gas prices to oil supply news is consistent with substitution between oil and natural gas among dual-fuel industrial and power-sector users. The magnitude and timing of our estimated pass-through, about one fifth of the oil response on impact and rising to roughly two thirds after twelve months, fall within the range documented in prior studies of the oil-gas price relationship (Villar and Joutz, 2006; Brown and Yücel, 2008; Hartley et al., 2008). More broadly, evidence that oil market disturbances affect natural gas prices is consistent with Zamani (2016) and Jadidzadeh and Serletis (2017).

prices in constructing the energy-cost shock.

We use the estimated impulse responses to construct predicted annual changes in the real price of each energy input. We first calculate predicted monthly changes and then aggregate them to the calendar-year level. In month m, the predicted change combines the current oil supply news shock with the remaining effects of shocks realized over the preceding twelve months:

$$\widehat{\Delta p}_{k,m} = \sum_{h=0}^{12} \psi_{k,h} Z_{m-h},$$

where $\widehat{\Delta p}_{k,m}$ is the predicted price change in energy input k in month m, with $m \in t$. A shock realized h months earlier, Z_{m-h}, is weighted by the estimated price response of energy input k to that shock at horizon h, $\psi_{k,h}$. We then construct the annual predicted price change by averaging these monthly predictions over the twelve months of calendar year t:

$$\widehat{\Delta p}_{k,t} = \frac{1}{12} \sum_{m \in t} \widehat{\Delta p}_{k,m}.$$

For example, a positive oil supply news shock in March of year t contributes to the annual measure for year t through its estimated effect on predicted prices from March ($h = 0$) through December ($h = 9$). Its contribution is

$$\frac{1}{12} \left(\sum_{h=0}^{9} \psi_{k,h} \right) Z_{\text{Mar},t}.$$

The remaining pass-through in months 10 through 12 enters the annual measure for the following calendar year.

Predetermined firm energy input exposure. Because we do not observe each firm's energy input mix directly, we use Compustat Segments data to identify its operating segments and proxy the firm's input mix with the energy use of the corresponding industries. We weight each industry by the segment's share of firm sales. We measure firm f's exposure to energy input k as

$$Exposure_{fk} = \sum_{s} w_{fs} \, \eta_{j(s)} \, e_{j(s)k},$$

where w_{fs} is segment s's share of firm f's manufacturing sales, $\eta_{j(s)}$ is energy expenditure as a share of sales in the segment's industry, and $e_{j(s)k}$ is that industry's expenditure share

for energy input k. Thus, $Exposure_{fk}$ is firm f's predicted spending on energy input k as a share of sales, based on its segment composition.[11] This makes exposure predetermined with respect to subsequent shocks and prevents the measure from incorporating firms' own input switching or reallocation across business lines in response to energy prices. It also creates variation in exposure among firms in the same industry, because their manufacturing segments use energy inputs differently. Among our sample firms, 75.6 percent of the cross-firm variation in oil and natural gas exposure remains after accounting for NAICS3 industry fixed effects.

2.5 Sample construction

Firm sample. Starting from US-incorporated firms in Compustat North America and their annual financial records, we apply the following sample restrictions. We retain firms that can be linked to patent assignees through the KPSS–CRSP–Compustat crosswalk and report at least one manufacturing segment with positive sales that can be matched to the industry energy-use data. This construction measures manufacturing activity using firms' operating segments rather than their header SIC, and therefore retains firms that have a nonmanufacturing main business but report a manufacturing segment. Because the paper's central analysis examines how energy-cost shocks affect the composition of green innovation, we further restrict the sample to firms that file at least one green patent between 1992 and 2021. The resulting sample comprises 1,245 green-patenting manufacturing firms and 21,012 firm-years. Appendix B reports the full sample construction.

Inventor sample. Starting from USPTO patents filed by our sample firms, we use PatentsView's algorithmically disambiguated inventor identities to construct the inventor sample. We assign each inventor to the firm with which they patent, using contemporaneous and prior patenting only (Card et al., 2013; Song et al., 2019; Akcigit et al., 2024). We interpret these inventor–firm links as employment relationships for assignment of the energy-cost shock. The sample inventors are corporate scientists: about 95 percent of their patents are assigned to firms, 77 percent of them patent with a single firm over their observed careers, and the median inventor files 7 patents over a career

[11] We match segments to industries at the finest available SIC level and fix segment weights at the firm's first observed business composition. To gauge the scale of the resulting variation, a one standard deviation shock corresponds to a roughly 5 percent increase in energy expenditure for a firm with median exposure. In the largest shock years, 1998, 2008, and 2015, the cross-firm spread reaches about 10 percent of the typical energy bill.

of 13 years, compared with 2 patents for the median USPTO inventor. Appendix C.3 shows that the results are unchanged when we exclude the 14.2% of inventors who ever patent with a non-corporate entity elsewhere in their full recorded PatentsView career, or when we restrict the sample to inventor–firm links corroborated by Revelio employment records. Section 4 provides further support for interpreting these inventors as corporate scientists: following firm-level shocks, the increase in green patenting among broader inventors is concentrated in patents involving new co-inventors, consistent with inventive activity organized through corporate R&D teams. Because the paper's central analysis examines how energy-cost shocks affect the composition of green innovation across inventors with different prior knowledge, we retain inventors who file at least one green patent between 1992 and 2021. After applying the lag and shock requirements, the final inventor sample comprises 650,688 inventor-years covering 68,606 inventors. Appendix B provides the full inventor–firm assignment procedure and sample construction.

2.6 Summary statistics

Table 1 reports summary statistics for the firm panel. Green patents account for 10 percent of patents in the average patenting firm-year, confirming that green innovation is a meaningful part of the inventive activity of the firms we study. The median firm patents with 13 inventors in a year, while the mean firm patents with 98.5. This within-firm-year variation allows us to compare inventors with different knowledge profiles who face the same firm-level energy-cost shock. The firms are predominantly engaged in manufacturing. Manufacturing segments account for 95.6 percent of sales on average and 100 percent at the median. Natural gas accounts for 21.6 percent of energy expenditure on average, with a median of 16.4 percent. This makes the response of natural gas prices to oil supply news economically relevant for the firms in our sample.

Table 2 reports summary statistics for the inventor panel. Over their observed patenting careers, inventors file 17.95 patents on average and 7 patents at the median. They file 3.83 green patents on average and 2 at the median. Career age averages 14.62 years and has a median of 13 years. Seventy-seven percent of inventors patent with a single sample firm, and the median mover patents with two firms. These patent histories provide the basis for measuring prior experience and knowledge, while the limited mobility of inventors across firms supports assigning firm-level energy-cost shocks. Appendix B.5 examines further aspects of the distributions of the three key variables: breadth, firm

energy-cost exposure, and the energy-cost shock.

Cross-sectional benchmark. We first document that broader inventors, in the cross-section, file more patents, and their patents are of higher scientific value. Figure 2 plots the relationship between residualized measures of inventors' breadth and patent counts (Panel A), breakthrough shares (Panel B), and highly cited shares (Panel C). We residualize all measures on firm-by-year fixed effects, within-year deciles of lagged patent stock, career age, and career age squared. The coefficients from a fitted linear regression imply that a one-standard-deviation increase in breadth is associated with 4.4% more patents, a 10.4% higher breakthrough share, and a 5.1% higher share of highly cited patents, relative to their respective means.[12] While these descriptive associations may reflect selection of inventors into more and higher-quality projects, they provide a benchmark for the subsequent analysis: broader inventors generally produce higher-quality innovation, yet following shocks that raise the returns to green innovation they file patents that are of lower scientific value.

3 Empirical Design and Firm-Level Results

This section presents the empirical design and establishes that energy-cost shocks induce firm-level green innovation. We first explain how the design uses predetermined differences in firms' exposure to energy-input price changes. We then estimate the firm-level response and present tests supporting the identifying assumptions.

3.1 Firm-level analysis

Design. The literature on induced innovation dates to Hicks (1932), who argued that a rise in the price of a factor of production induces innovations that economize on that factor. Acemoglu (2002) develops a model of directed technical change in which market incentives shape the direction of innovation. In the context of green innovation, previous research has shown that higher energy prices induce energy efficient innovation (Popp, 2002), and higher fuel prices shift auto industry patenting toward clean technologies (Aghion et al., 2016).[13] This evidence motivates our use of energy-cost shocks to generate firm-level changes in the incentives for green innovation.

[12] Figure A1 repeats the exercise for green patenting alone and finds a similar pattern.

[13] See Hémous and Olsen (2021) and Dechezleprêtre and Hémous (2022) for extensive literature reviews.

Our design exploits annual variation in the returns to green innovation at the firm level, captured by the energy-cost shock measure developed in Section 2.4. The identification strategy relies on a reduced-form shift-share design (Borusyak et al., 2025): predetermined firm-level exposure to energy inputs, combined with yearly energy input price changes generated by exogenous oil supply news shocks. The design therefore draws its quasi-experimental variation from the oil supply news shocks, while fixed exposure shares translate those common shocks into differential changes in firms' energy costs, rather than relying on the cross-sectional relationship between baseline exposure and patenting. At the inventor level, we compare inventors who differ in their pre-existing breadth but face the same shock because they patent with the same firm in the same year. This analysis addresses our two questions: whether broader inventors increase green patenting more following the shock, and how the quality of their green output changes.

As a first step, we show that the energy-cost shock raises green patenting by the firms with which our sample inventors patent, but does not change their non-green patenting. We estimate the following equation separately for the number of green and non-green patents filed by firm f in year t:

$$y_{ft} = \exp\left[\beta\, Shock_{f,t-1} + \gamma' X_{f,t-1} + \alpha_f + \delta_{j(f),t}\right] + \varepsilon_{ft}, \tag{1}$$

y_{ft} is the number of green patents filed by firm f in year t, or the number of non-green patents in the corresponding specification. $Shock_{f,t-1}$ is the firm's annual energy-cost shock, lagged one year so that it precedes the innovation outcome, and z-scored to have mean zero and unit standard deviation. The vector $X_{f,t-1}$ includes lagged log assets, liquidity, leverage, profitability, and R&D intensity, and accounts for differences in firms' financial resources and research capacity that may affect how firms' patenting responds to the energy-cost shock. The firm fixed effects, α_f, control for time-invariant differences across firms. The industry-by-year fixed effects, $\delta_{j(f),t}$, control for industry-specific changes in innovation opportunities and energy costs. ε_{ft} is the error term. We estimate the patent-count specifications by PPML.[14] Standard errors are heteroskedasticity robust and clustered by firm.

For Equation 1 to identify the effect of energy costs on firm innovation, conditional on

[14] PPML is consistent without requiring the outcome to follow a Poisson distribution and accommodates fixed effects (Cohn et al., 2022; Chen and Roth, 2024). Poisson estimators are commonly used for count outcome variables (Aghion et al., 2016; Azoulay et al., 2019; Naaraayanan et al., 2026).

firm fixed effects, industry-by-year fixed effects, and lagged financial controls, the Energy Cost Shock must be orthogonal to omitted changes in firm patenting:

$$\mathbb{E}\left[Shock_{f,t-1}\,\varepsilon_{ft} \mid X_{f,t-1}, \alpha_f, \delta_{j(f),t}\right] = 0,$$

This assumption allows firms with different baseline energy exposure to differ systematically in their energy use, innovation activity, and financial characteristics. Nonetheless, Appendix Figure A2 shows that, within industry, firms' predetermined exposure is not systematically related to their size, leverage, liquidity, R&D intensity, or patenting activity, measured in the year before the firm's exposure baseline year. Without conditioning on industry, more exposed firms have lower leverage and profitability and higher liquidity and R&D intensity, differences that reflect industry composition. The exception is profitability, which remains negatively associated with exposure after conditioning on industry. We control for lagged profitability in every specification (Equation 1), and the magnitude is modest: a one-standard-deviation increase in profitability corresponds to a 0.6 percentage point decline in exposure, about 1 percent of its mean. Because energy intensity is measured at the industry level, the industry-by-year fixed effects absorb differences in overall energy exposure across industries, so that identification comes from comparisons across firms in the same industry and year that differ in their predetermined segment composition and fuel mix. Importantly, Känzig (2021) shows that the high-frequency surprises underlying the oil supply news series are serially uncorrelated, are not forecasted by past macroeconomic and financial variables, and are uncorrelated with measures of global demand, productivity, uncertainty, financial, monetary, and fiscal policy shocks. Furthermore, we discuss below a set of tests that provide support for our identification design based on exogenous shifts (Adão et al., 2019; Borusyak et al., 2025).

Results. Table 3 presents the results from estimating equation (1). Columns (1) and (2) show that a one-standard-deviation shock, corresponding to roughly a 5% increase in energy costs for the median firm, raises green patenting by 3.7% without financial controls and 3.8% with them. On the other hand, columns (3) and (4) show that the estimated effect on non-green patenting is essentially zero and statistically insignificant. Therefore, as hypothesized, the shock induces green innovation rather than an overall increase in firm R&D activity.

Identification and validity tests. First, we examine the timing of the firm-level response. If the Energy Cost Shock captures an unanticipated change in incentives, a shock realized after the patenting outcome should not predict earlier firm green patenting. We re-estimate Equation 1 for outcomes from two years before to three years after the baseline outcome year t, holding the event shock fixed at $t-1$. For the placebo, contemporaneous, and baseline outcomes in columns (1)–(3), financial controls are dated one year before each outcome. For the post-shock outcomes in columns (4)–(6), financial controls remain fixed at $t-1$, and we additionally control for every shock realized between $t-1$ and the outcome. Column (1) of Table A2 shows that the placebo, a shock dated one year after the outcome, is small and statistically indistinguishable from zero. The contemporaneous horizon in column (2), where the tested shock falls in the same calendar year as the outcome, is positive and statistically significant, while somewhat smaller than the main effect in Table 3. The positive contemporaneous effect is consistent with evidence that innovation can respond within the year to changes in expected returns (Aghion et al., 2016; Hill et al., 2025).[15] The effect is largest in year t, the year after the Energy Cost Shock (column (3)), and decays to small, statistically insignificant estimates in the three following years (columns (4)–(6)).

Second, following a procedure developed by Adão et al. (2019), we show evidence suggesting that the exposure shares are not driving our identification. We replace the realized sequence of energy-price shifts, Δp_{kt}, with randomly generated shifts, hold firms' predetermined exposure shares fixed, reconstruct the Energy Cost Shock, and re-estimate Equation 1 across 1,000 trials. We find that 2.9% of the coefficients are statistically significant at the 1% level and 10.3% at the 5% level (see Appendix Figure A3), rejection rates comparable to those reported in other shift share designs relying on exogeneity of the shifts.[16] The placebo distribution supports our identifying assumption that the estimates are driven by the realized energy-price shifts, Δp_{kt}, rather than by the exposure shares alone.

Third, we show evidence that the firm-level response is specific to energy prices rather than a general sensitivity of energy-exposed firms to aggregate macro-level shocks. If the exposure measure captured a more general tendency for firms to respond to cyclical or financial conditions, exposed firms would respond to other aggregate shocks as well.

[15] The ratio of the contemporaneous to one-year-lagged fuel-price elasticity is 0.91 in Aghion et al. (2016), compared with 0.88 in our firm-level estimates.

[16] Placebo rejection rates at the 5% level range from 11.5% in Bazzi et al. (2023) and 16.1% in Derenoncourt (2022) to 55% in the setting studied by Adão et al. (2019).

We therefore replace the predicted energy-price changes with four alternative aggregate shocks, retaining the same predetermined exposure measure and construction.[17] Appendix Table A3 shows that none generates a statistically significant change in firm green patenting. These placebo shocks indicate that green patenting by exposed firms responds to the energy-cost shocks specifically, and not to other aggregate news or macroeconomic shocks.

4 Inventor-Level Results

This section presents the inventor-level analysis. We examine whether inventors with broader prior experience increase green patenting more following a firm-level energy-cost shock, and how the scientific value of their output changes.

4.1 Empirical design

The inventor-level analysis compares inventors who patent with the same firm in the same year and therefore face the same energy-cost shock and all other firm-year conditions that may affect green innovation, such as the allocation of research and development resources and firm innovation strategy.

We start by studying the effect of the shock on the quantity of green innovation produced by broader inventors. We estimate the following specification for the annual green and non green patent counts of inventor i:

$$y_{it} = \exp\left[\beta\left(Shock_{f,t-1} \times Breadth_{i,t-1}\right) + \theta Breadth_{i,t-1} + \gamma' X_{i,t-1} + \alpha_i + \alpha_{ft}\right] + \varepsilon_{it}, \tag{2}$$

where y_{it} is the number of green patents filed by inventor i in year t, or the number of non-green patents in the corresponding specification. $Shock_{f,t-1}$ is the firm's annual energy-cost shock, lagged one year so that it precedes the patenting outcome, and z-scored to have mean zero and unit standard deviation. $Breadth_{i,t-1}$ measures the dispersion of the inventor's prior patenting across technology fields, capturing the breadth of their experience and knowledge. The coefficient β captures how the response to a one-standard-deviation shock varies with breadth. We report magnitudes per one-standard-deviation

[17] We consider innovations in geopolitical risk (Caldara and Iacoviello, 2022), monetary-policy and central-bank-information shocks (Jarociński and Karadi, 2020), and fiscal news (Ramey and Zubairy, 2018).

increase in breadth. If broader experience makes inventors more versatile across activities, we expect $\beta > 0$. The vector $X_{i,t-1}$ includes the lagged share of the inventor's prior patents classified as green, the log of one plus the inventor's lagged patent stock, and career age and its square. ε_{it} is the error term.

The firm-by-year fixed effects, α_{ft}, absorb the direct effect of the shock and all other conditions common to inventors patenting with the same firm in the same year. The inventor fixed effects, α_i, absorb time-invariant differences across inventors, and linear career age is absorbed jointly by the fixed effects. We estimate the equation by PPML.[18] Standard errors are clustered two-way by firm and inventor because inventors at the same firm share energy-cost shocks, and each inventor's outcomes may be correlated over time (Cameron et al., 2011; Abadie et al., 2023).

The estimation equation for the quality outcomes is analogous to Equation 2:

$$\begin{aligned} y_{it} = \beta \left(Shock_{f,t-1} \times Breadth_{i,t-1} \right) + \theta Breadth_{i,t-1} \\ + \gamma' X_{i,t-1} + \alpha_i + \alpha_{ft} + \varepsilon_{it}, \end{aligned} \tag{3}$$

where y_{it} is each measure of the quality of green patents filed by inventor i in year t. A negative β indicates that the quality of green patenting declines more for broader inventors following the shock. To study changes in the composition of green innovation, we restrict the sample to inventor-years in which the inventor files at least one green patent and estimate the equation by OLS. The controls and fixed effects are the same as in Equation 2, and Subsection 2.2 describes the construction of the quality measures.

4.2 Main results

Quantity results. Table 4 presents the results from estimating Equation 2. Consistent with breadth capturing inventor versatility, the effect of the shock on green patenting increases with breadth. Focusing on columns (2) and (4), which include inventor fixed effects, a one standard deviation increase in the shock raises green patenting by 2% more for an inventor whose breadth is one standard deviation greater (column (2)), with a small, negative, and statistically insignificant corresponding effect on non-green patenting (column (4)). Columns (1) and (3) omit inventor fixed effects and therefore also draw on comparisons across inventors. The estimated effect on green patenting remains posi-

[18] In the inventor-level green-patent sample, 74.8% of inventor-years have no green patent and the count is right-skewed.

tive and statistically significant at 1.4% (column (1)), while the non-green effect remains close to zero (column (3)). The similar estimates indicate that the result does not rely on fixed differences across inventors, while the larger estimate with inventor fixed effects is consistent with a stronger within-inventor effect.

Inventor-team composition. To assess whether firms reassign broader inventors to the induced green projects, we examine whether their additional green patents involve new co-inventors. To the extent that reassignment to new projects changes the colleagues with whom an inventor works, patents involving new co-inventors are informative about this margin. We classify a co-inventor as new if the pair had never appeared together on an earlier patent, at any firm. Table 5 estimates Equation (2) using separate indicators for filing at least one patent with a new co-inventor and filing at least one patent without one. Because fewer than 4% of inventor-years contain multiple green patents within either group, these indicators closely capture the corresponding patenting margins. A one-standard-deviation increase in the shock raises the probability of filing a green patent with a new co-inventor by 1.8% of its mean for an inventor whose breadth is one standard deviation higher. The corresponding estimate for green patents without a new co-inventor is close to zero. There is no corresponding increase in either group of non-green patents. Appendix Table A4 shows the same pattern using patent counts estimated by PPML. The concentration of the increase in green patents involving new collaborators is consistent with firms reassigning broader inventors to the induced green projects.

Quality results. Table 6 reports the effect of the shock on breakthrough innovation, our primary measure of scientific value, defined as the share of green patents in the top decile of the KPST importance distribution. A larger shock reduces the share of green patents that are breakthroughs more for broader inventors. In column (2), which includes inventor fixed effects, a one standard deviation increase in the shock reduces the breakthrough share by 7.1% of its mean for an inventor whose breadth is one standard deviation greater. When inventor fixed effects are omitted in column (1), the estimated effect is attenuated to 4.6% of the outcome mean and is only marginally statistically significant. As in the quantity result, this pattern suggests that the effect is concentrated within inventors rather than resulting only from persistent differences in breakthrough rates across broader and less broad inventors. Column (3) additionally controls for the inventor's lagged breakthrough share, and the estimate remains essentially unchanged at 7.2% of the outcome mean. This

indicates that the result is not explained by differences in inventors' recent innovation quality. Columns (4)–(6) report the corresponding estimates for non-green patents. The estimated effects range from -1.4% to -0.8% of the outcome mean and are small relative to the green-patent estimates. The decline in breakthrough innovation is therefore concentrated in green patenting.

Appendix Table A5 shows that the decline in breakthrough share is robust to alternative measures of breakthrough innovation, including breakthrough counts, the probability of filing any breakthrough patent, and continuous scientific importance. The effect is also negative, though less precisely estimated, for the highly cited share. The estimated effect on the KPSS measure of economic value is small and statistically indistinguishable from zero. Taken together, following stronger energy-cost shocks, we find a decline in the scientific value of broader inventors' green patents but no statistically significant change in their economic value.

We next use a quintile-based specification to examine the shock response across the breadth distribution, and in particular among the broadest inventors. We divide inventors into five equally sized quintiles based on their lagged breadth. For each outcome, we re-estimate Equation 2 after replacing the continuous breadth measure and its interaction with the shock with indicators for the second through fifth breadth quintiles and interactions between each indicator and the shock. The narrowest quintile is omitted as the reference group, so that each interaction coefficient measures the difference in the shock response between that quintile and the narrowest quintile. Figure 3 reports these estimates. The estimated quantity response is positive in every quintile above the narrowest and rises from about 2% in the second quintile to about 4% in the fourth and fifth quintiles, twice the 2% effect for a one-standard-deviation difference in breadth reported in Table 4. The estimated breakthrough response is negative in every higher quintile and ranges from 14% to 17% of the outcome mean in the third through fifth quintiles, roughly twice the 7.1% effect in Table 6.

Having established that broader inventors' induced green patents are less likely to be breakthroughs, we turn to study how their green output shifts across the entire importance distribution. Figure 4 shows that a one standard deviation increase in the shock reduces the share of green patents in the top decile by 0.90 percentage points for an inventor whose breadth is one standard deviation greater. It increases the share in the 50–70th percentile band by 0.74 percentage points. Notably, the estimated effects in the bottom half of the distribution are small and statistically indistinguishable from zero. This find-

ing indicates a step down from breakthrough inventions to upper-middle-tier patents, rather than a shift toward the least scientifically important patents.

4.3 Dynamics, placebos and robustness

Dynamic effects and placebo test. Identification requires that, conditional on the controls and fixed effects, energy-price shifts are orthogonal to other determinants of differences in green innovation between broader and narrower inventors in the same firm-year. One implication is that the shifts should not predict differential changes in their innovation outcomes before they occur. The dynamic analysis examines this implication and traces the effect of a shock at $t-1$ on outcomes from $t-2$ to $t+3$. Table 7 reports the estimates for green patenting in Panel A and the breakthrough share in Panel B. In both panels, the estimates at the placebo and contemporaneous horizons, $t-2$ and $t-1$, are statistically indistinguishable from zero. The quantity increase and the decline in the breakthrough share emerge in year t, the year after the energy-cost shock. In the years that follow, the quantity estimates remain positive. At $t+1$ and $t+3$, they are marginally statistically significant and roughly half the magnitude of the effect at t, while the estimate at $t+2$ is close to zero. For the breakthrough share, the estimate at $t+1$ is close to zero, the estimate at $t+2$ is positive but statistically insignificant, and the estimate at $t+3$ is negative and statistically significant at roughly one-third the magnitude of the effect at t. The effects are strongest one year after the shock, consistent with prior research showing that energy-related innovation responds within one to two years to changes in energy prices or policy (Johnstone et al., 2010; Dugoua and Gerarden, 2025).

Robustness. Appendix C reports additional checks on the construction of breadth and the energy-cost shock, sample selection, and outcome measurement. Subsection C.1 reconstructs breadth using different levels of the CPC hierarchy, the technology classes of cited patents, a Rao–Stirling index that accounts for technological distance across fields, within-field standardization, and allowing for knowledge to depreciate. It also addresses how jointly filed patents enter an inventor's knowledge stock by weighting patents by the inverse of team size or restricting the measure to patents filed by small teams, and controls for team size and co-inventors' breadth, prior green patenting, and prior breakthrough rates. Subsection C.2 varies the horizon and fuel composition used to construct the energy-cost shock and constructs firm exposure using either the firm's largest manufacturing segment or its firm-level SIC rather than its sales-weighted segment structure.

Subsection C.3 restricts the sample to inventors with only corporate patents and to inventors whose employment at the patenting firm is confirmed in résumé data. It also examines alternative quantity outcomes and estimators, fixed effects and firm controls, breakthrough definitions, and definitions of green innovation, and re-estimates the team-composition results using patent counts and PPML. Across these analyses, broader inventors increase green patenting without a corresponding increase in non-green patenting, while the scientific value of their induced green patents declines.

Taken together, the descriptive and causal results distinguish inventors' overall output from the additional green innovation induced by a rise in energy costs. In the cross section, broader inventors produce more patents overall and have higher shares of breakthroughs and highly cited patents. Following an increase in energy costs, however, broader inventors increase green patenting more strongly, and the additional green patents they produce are less likely to be breakthroughs relative to their own average green output. This pattern is consistent with broader experience making inventors versatile and more readily redeployable across activities and with higher returns bringing additional, less path-breaking projects into production. We develop this interpretation in Section 5.

5 Mechanisms

In this section, we argue that reassignment of broader inventors to newly viable green projects provides a consistent explanation for the quantity, team-composition, and quality results. We then provide evidence that poorer inventor matches, prior green experience, selection on inventor ability, cross-firm mobility, and larger technological pivots do not account for the decline in scientific value.

Firms respond to new investment opportunities by adjusting their workforce across activities (Giroud and Mueller, 2015; Cestone et al., 2024). Which workers are redeployed depends on the portability of their skills. In the skill-weights framework of Lazear (2009), accumulated skills are more portable when they are useful across a wider range of jobs. Broader experience across technical domains can build knowledge useful across the firm's activities, making inventors more versatile. Following an energy-cost shock, broader inventors increase green patenting but not their other patenting, and the increase in green patenting is concentrated in patents involving new collaborators. Because corporate in-

novation activity is typically directed through R&D projects (Nagaoka and Walsh, 2009), we view the finding that broader inventors' additional green patents involve new collaborators as indicating that firms expand green innovation by assigning these inventors to projects that become viable due to higher energy costs.

The quality results provide further evidence on the nature of the green work associated with this expansion. When higher energy costs raise the returns to green innovation, firms may find it worthwhile to pursue projects that previously fell below their profitable investment threshold. The estimated decline in the breakthrough share is present when we compare broader and narrower inventors within the same firm-year and becomes larger when we follow the same inventor over time. This pattern is consistent with firms assigning broader inventors to newly viable green projects of somewhat lower scientific value. Figure 4 shows that the composition of their green output shifts away from breakthroughs and toward upper-middle-tier patents, with no estimated increase in the bottom half of the scientific-importance distribution. This shift from breakthroughs toward upper-middle-tier patents is consistent with the project-selection explanation: higher returns make additional projects near the firm's existing investment threshold viable. Mokyr (1990) emphasizes that microinventive effort responds to prices and incentives, whereas macroinventions have a less predictable connection to economic conditions. This distinction helps interpret our results: firms assign versatile inventors to newly viable projects even when this work is less path-breaking than those inventors' usual green output.

Lower scientific value need not imply lower economic value to the firm. In fact, the estimated effect on the KPSS measure of economic value is small and statistically indistinguishable from zero (Appendix Table A5, column (7)). Furthermore, Mokyr (1990) emphasizes that macroinventions and microinventions are complementary: microinventions reduce production costs and help make major advances practical and productive, while, without new macroinventions, continued refinement of existing techniques eventually runs into diminishing returns.

5.1 Alternative explanations

We next examine five alternative explanations that could account for the lower scientific value of broader inventors' induced green patents.

Inventor-team match quality. One concern is that the lower scientific value reflects poorer matches among the inventors assigned to the induced projects, rather than the scientific importance of the projects themselves. This possibility arises because the increase in green patenting is concentrated in patents involving new collaborators. While we cannot directly distinguish project quality from match quality, we show that the patent quality predicted by the assigned team's observable characteristics does not decline. Appendix Section D explains how we predict each patent's breakthrough probability using a rich set of predetermined characteristics of the full team and its members, and Appendix Table A6 reports the results. For an inventor whose breadth is one standard deviation higher, a one-standard-deviation increase in the shock raises the predicted breakthrough share by 0.6% of its mean, an estimate statistically indistinguishable from zero, while the actual breakthrough share falls by 6.2% of its mean. In addition, Appendix Table A7 shows that the breakthrough share declines by a statistically significant 6.4% for green patents involving a new co-inventor and by a statistically insignificant 4.2% for green patents without one. However, the two estimates are not statistically distinguishable from one another. Thus, we find no evidence that the lower scientific value reflects the formation of observably weaker inventor matches.

Prior green innovation experience. The estimated effect of breadth could instead reflect inventors' prior experience in green innovation. If broader inventors are more likely to have worked on green technologies before the shock, their larger quantity response and lower breakthrough share could be driven by prior green experience rather than versatility across activities. Table 8 tests this possibility by re-estimating the quantity and quality specifications after adding, one at a time, an interaction between the energy-cost shock and an alternative measure of prior green experience. Column (1) uses the lagged green share, columns (2) and (3) use indicators for at least one or two prior green patents, column (4) uses $\log(1 + \text{green patent stock})$, and column (5) uses lagged proximity to green knowledge, measured by cosine similarity following Jaffe (1986).[19] The specifications otherwise include the controls and fixed effects described in Equations 2 and 3.

Across all five specifications, a one-standard-deviation increase in the energy-cost shock raises green patenting by 1.2 to 1.9 percent more for an inventor whose breadth is one standard deviation greater, while lowering the breakthrough share by 6.9 to 7.7

[19] Appendix B.3 describes the green-proximity measure. Across the relevant quantity and quality samples, green share and green proximity have correlations with breadth between 0.02 and 0.10, while the prior-green indicators and green patent stock have correlations between 0.29 and 0.36.

percent of its mean. The positive quantity response associated with four of the five measures of prior green experience is consistent with path dependence in directed innovation, whereby prior knowledge in a technological domain facilitates subsequent innovation in that domain (Aghion et al., 2016; Dugoua and Gerarden, 2025). Notably, allowing each measure of prior green experience to interact with the shock leaves both breadth effects qualitatively unchanged. Prior green experience therefore does not account for the quantity and quality effects associated with breadth.

Selection on inventor ability. Our results could reflect selection on inventor ability rather than the quality of the projects that broader inventors patent following a higher energy-cost shock. If broader inventors who redeploy into green innovation after an energy-cost shock have lower ability, their induced patents may be of lower scientific value. We use Revelio résumé data, and match 29,864 of the 68,606 inventors in the main green-count regression (43.5 percent), to test this possibility. These data provide information on educational attainment, which we use as human-capital proxies. Table 9 reports the full-sample estimates in column (1) and the estimates for the matched sample in column (2). Columns (3)–(6) re-estimate the specification after adding, one at a time, interactions between the shock and indicators for a degree from a top-20 university, a degree from a top-50 university, a master's degree, and a PhD.[20] The specifications otherwise include the controls and fixed effects described in Equations 2 and 3. Both the quantity and quality results remain unchanged. Across columns (2)–(6), a one-standard-deviation increase in the shock raises green patenting by 2.7 to 2.8 percent more for an inventor whose breadth is one standard deviation higher, while lowering the breakthrough share by 9.6 to 10.0 percent of its mean. Thus, observed differences in inventor ability and human capital do not account for the stronger quantity response or the decline in breakthrough share among broader inventors.

Cross-firm mobility. Our results could also be explained by the movement of inventors across firms rather than the selection of projects they undertake within firms. If broader inventors move to firms with lower-quality green opportunities following an energy-cost

[20]The top-20 and top-50 indicators equal one if any degree recorded by Revelio is from a university in the corresponding top tier of the *U.S. News & World Report* 2025 National Universities ranking for degrees earned in the United States (https://www.usnews.com/best-colleges/rankings/national-universities), or of the *Times Higher Education* 2025 World University Rankings for degrees earned elsewhere (https://www.timeshighereducation.com/world-university-rankings/2025/world-ranking).

shock, the resulting decline in breakthrough share could reflect this re-sorting rather than a change in project selection within firms. Table 10 tests this possibility by re-estimating Equations 2 and 3 using samples of inventors with limited cross-firm mobility. Column (1) reports the full-sample estimates, column (2) restricts the sample to inventor-years with no recorded firm switch in the current year or preceding two years, and column (3) restricts it to inventors who patent with only one firm throughout their observed careers.[21] The specifications otherwise include the controls and fixed effects described in Equations 2 and 3. Both the quantity and quality results remain unchanged. Across the three samples, a one-standard-deviation increase in the shock raises green patenting by 1.8 to 2.0 percent more for an inventor whose breadth is one standard deviation higher, while lowering the breakthrough share by 7.8 to 8.3 percent of its mean. The quantity effect remains essentially unchanged, while the decline in breakthrough content is modestly larger in the restricted samples. Thus, the movement of inventors across firms does not account for the stronger quantity response or the decline in breakthrough share among broader inventors.

A pivot penalty. Finally, our results could reflect a decline in scientific value because an energy-cost shock induces broader inventors to move farther from their established knowledge. Hill et al. (2025) document, across scientists and inventors, that work conducted farther from prior knowledge is less likely to be highly cited, which they describe as a pivot penalty. Table 11 tests this possibility by re-estimating Equation 3 after adding technological distance and its interaction with the shock. Column (1) reports the baseline estimate. Columns (2)–(4) use three measures of pivot size. Column (2) replicates the citation-based measure in Hill et al. (2025), defined as one minus the cosine similarity between the CPC main-group mix of a new green patent's backward citations and the inventor's citation-based technology mix over the preceding three years. Column (3) uses the same cosine-similarity measure but compares the new green patent with the inventor's full prior citation history. Column (4) instead compares the CPC main groups assigned to the new green patent with the CPC main groups assigned to the inventor's full prior patent portfolio. The specifications otherwise include the controls and fixed effects described in Equation 3. The quality result remains unchanged: a one-standard-deviation increase in the shock lowers the breakthrough share by 5.8 to 7.5 percent of its mean for

[21] Relative to the baseline samples used in the quantity and quality regressions, respectively, the no-recent-switch restriction retains 84.1 and 88.5 percent of observations, and the never-switcher restriction retains 63.8 and 67.1 percent.

an inventor whose breadth is one standard deviation greater. The interaction between pivot size and the shock is statistically insignificant for the two citation-based measures. It is positive and marginally significant for the own-CPC measure, the opposite of the negative relationship implied by a pivot penalty. Thus, technological distance from prior work does not explain the decline in breakthrough share among broader inventors.

The decline in scientific value is therefore not explained by observably weaker inventor-team matches, prior green innovation experience, observed differences in inventor ability, cross-firm re-sorting, or larger technological pivots. These tests lend support to our interpretation that firms assign broader inventors to newly viable green projects that are less path-breaking.

6 Conclusion

This paper examines the role of workers' breadth of experience in firms' response to changing investment opportunities. Using patent records and firm-level energy-cost shocks that raise the returns to green innovation, we trace changes in the quantity and scientific value of inventors' output.

In the cross section, inventors whose prior patenting is more dispersed across technology fields file more patents and produce inventions of greater scientific value. Following exogenous increases in energy costs, however, we observe a larger increase in the number of green patents involving broader inventors, without a decline in their non-green patenting. The increase is concentrated in patents involving new co-inventors. These additional green patents are less likely to be of high scientific value than these inventors' usual green work, shifting away from breakthroughs toward upper-middle-tier patents, with no estimated increase in the lower half of the quality distribution.

Our findings are consistent with firms drawing on broader inventors to staff green projects that would have remained below their investment threshold under a smaller exogenous increase in energy costs. The lower scientific value of these inventors' additional green output is consistent with an expansion of microinventions that improve existing technologies and reduce costs. The economic value of these patents, by contrast, shows no change. Observable team characteristics do not predict the decline in breakthrough share, and the estimated effects for patents with and without new co-inventors are not statistically distinguishable. The decline is also not explained by inventors' prior green experience, observed inventor ability, cross-firm mobility, or technological distance from

an inventor's prior work.

The results indicate that workers' breadth of experience facilitates firms' adjustment to changing investment opportunities. When an activity becomes more valuable, firms decide both whether to invest more and which workers to assign to newly viable projects. Broader experience makes workers more versatile across activities, but their contribution to firms' adjustment involves additional projects of lower scientific value than their usual work.

References

Abadie, A., S. Athey, G. W. Imbens, and J. M. Wooldridge (2023). When should you adjust standard errors for clustering? *The Quarterly Journal of Economics 138*(1), 1–35.

Acemoglu, D. (2002). Directed technical change. *The Review of Economic Studies 69*(4), 781–809.

Acemoglu, D., P. Aghion, L. Barrage, and D. Hémous (2023). Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution. Working Paper 31657, National Bureau of Economic Research.

Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hémous (2012). The environment and directed technical change. *American Economic Review 102*(1), 131–166.

Acemoglu, D., U. Akcigit, D. Hanley, and W. Kerr (2016). Transition to clean technology. *Journal of Political Economy 124*(1), 52–104.

Acemoglu, D. and J. Linn (2004). Market size in innovation: Theory and evidence from the pharmaceutical industry. *The Quarterly Journal of Economics 119*(3), 1049–1090.

Adão, R., M. Kolesár, and E. Morales (2019). Shift-share designs: Theory and inference. *Quarterly Journal of Economics 134*(4), 1949–2010.

Adhvaryu, A., N. Kala, and A. Nyshadham (2022). Management and shocks to worker productivity. *Journal of Political Economy 130*(1), 1–47.

Aghion, P., A. Bergeaud, M. De Ridder, and J. Van Reenen (2024). Lost in transition: Financial barriers to green growth. Working paper, March.

Aghion, P., A. Dechezleprêtre, D. Hémous, R. Martin, and J. Van Reenen (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy 124*(1), 1–51.

Aharonson, B. S. and M. A. Schilling (2016). Mapping the technological landscape: Measuring technology distance, technological footprints, and technology evolution. *Research Policy 45*(1), 81–96.

Akcigit, U., J. G. Pearce, and M. Prato (2024). Tapping into talent: Coupling education and innovation policies for economic growth. *Review of Economic Studies 91*(1), 1–31.

Arts, S. and L. Fleming (2018). Paradise of novelty—or loss of human capital? exploring new fields and inventive output. *Organization Science 29*(6), 1074–1092.

Avalos, R., T. Fitzgerald, and R. R. Rucker (2016). Measuring the effects of natural gas pipeline constraints on regional pricing and market integration. *Energy Economics 60*, 217–231.

Azoulay, P., C. Fons-Rosen, and J. S. Graff Zivin (2019). Does science advance one funeral at a time? *American Economic Review 109*(8), 2889–2920.

Baker, G., M. Gibbs, and B. Holmström (1994). The internal economics of the firm: Evidence from personnel data. *Quarterly Journal of Economics 109*(4), 881–919.

Bazzi, S., A. Ferrara, M. Fiszbein, T. Pearson, and P. A. Testa (2023). The other great migration: Southern whites and the new right. *Quarterly Journal of Economics 138*(3), 1577–1647.

Becker, G. S. (1962). Investment in human capital: A theoretical analysis. *Journal of Political Economy 70*(5, Part 2), 9–49.

Benson, A., D. Li, and K. Shue (2019). Promotions and the Peter Principle. *Quarterly Journal of Economics 134*(4), 2085–2134.

Borusyak, K., P. Hull, and X. Jaravel (2025). A practical guide to shift-share instruments. *Journal of Economic Perspectives 39*(1), 181–204.

Brown, S. P. A. and M. K. Yücel (2008). What drives natural gas prices? *The Energy Journal 29*(2), 45–60.

Caldara, D. and M. Iacoviello (2022). Measuring geopolitical risk. *American Economic Review 112*(4), 1194–1225.

Calel, R. and A. Dechezleprêtre (2016). Environmental policy and directed technological change: Evidence from the European carbon market. *The Review of Economics and Statistics 98*(1), 173–191.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics 29*(2), 238–249.

Card, D., J. Heining, and P. Kline (2013). Workplace heterogeneity and the rise of west german wage inequality. *Quarterly Journal of Economics 128*(3), 967–1015.

Cestone, G., C. Fumagalli, F. Kramarz, and G. Pica (2016). Insurance between firms: The role of internal labor markets. CEPR Discussion Paper 11336, Centre for Economic Policy Research, London.

Cestone, G., C. Fumagalli, F. Kramarz, and G. Pica (2024). Exploiting growth opportunities: The role of internal labour markets. *Review of Economic Studies 91*(5), 2676–2716.

Chen, J. and J. Roth (2024). Logs with zeros? some problems and solutions. *Quarterly Journal of Economics 139*(2), 891–936.

Cohn, J. B., Z. Liu, and M. I. Wardlaw (2022). Count (and count-like) data in finance. *Journal of Financial Economics 146*(2), 529–551.

Conti, R., A. Gambardella, and M. Mariani (2014). Learning to be edison: Inventors, organizations, and breakthrough inventions. *Organization Science 25*(3), 833–849.

Custódio, C., M. A. Ferreira, and P. Matos (2013). Generalists versus specialists: Lifetime work experience and chief executive officer pay. *Journal of Financial Economics 108*(2), 471–492.

Dechezleprêtre, A. and D. Hémous (2022). Directed technical change and environmental economics. In *The Economics of Creative Destruction*.

Derenoncourt, E. (2022). Can you move to opportunity? evidence from the great migration. *American Economic Review 112*(2), 369–408.

Doeringer, P. B. and M. J. Piore (1985). *Internal Labor Markets and Manpower Analysis*. Armonk, NY: M.E. Sharpe.

Dranove, D., C. Garthwaite, and M. Hermosilla (2022). Does consumer demand pull scientifically novel drug innovation? *The RAND Journal of Economics 53*(3), 590–638.

Dugoua, E. and T. D. Gerarden (2025). Induced innovation, inventors, and the energy transition. *American Economic Review: Insights 7*(1), 90–106.

Eckardt, D. (2026). Training specificity and occupational mobility: Evidence from German apprenticeships. *Econometrica 94*(3), 741–766.

Finkelstein, A. (2004). Static and dynamic effects of health policy: Evidence from the vaccine industry. *The Quarterly Journal of Economics 119*(2), 527–564.

Fleming, L. (2001). Recombinant uncertainty in technological search. *Management Science 47*(1), 117–132.

Fleming, L. and O. Sorenson (2004). Science as a map in technological search. *Strategic Management Journal 25*(8-9), 909–928.

Fowlie, M., M. Greenstone, and C. Wolfram (2018). Do energy efficiency investments deliver? evidence from the weatherization assistance program. *Quarterly Journal of Economics 133*(3), 1597–1644.

Frydman, C. (2019). Rising through the ranks: The evolution of the market for corporate executives, 1936–2003. *Management Science 65*(11), 4951–4979.

Gathmann, C. and U. Schönberg (2010). How general is human capital? A task-based approach. *Journal of Labor Economics 28*(1), 1–49.

Gibbons, R. and M. Waldman (2004). Task-specific human capital. *American Economic Review 94*(2), 203–207. Papers and Proceedings.

Giroud, X. and H. M. Mueller (2015). Capital and labor reallocation within firms. *Journal of Finance 70*(4), 1767–1804.

Goldschlag, N., T. J. Lybbert, and N. J. Zolas (2020). Tracking the technological composition of industries with algorithmic patent concordances. *Economics of Innovation and New Technology 29*(6), 582–602.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020). Bartik instruments: What, when, why, and how. *American Economic Review 110*(8), 2586–2624.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2001). The NBER patent citations data file: Lessons, insights and methodological tools. Working Paper 8498, National Bureau of Economic Research.

Hartley, P. R., I. Medlock, Kenneth B., and J. E. Rosthal (2008). The relationship of natural gas to oil prices. *The Energy Journal 29*(3), 47–65.

Hausman, C. (2019). Shock value: Bill smoothing and energy price pass-through. *Journal of Industrial Economics 67*(2), 242–278.

He, W., D. Prud'homme, N. Han, and K. G. Huang (2026). Inventor rewards, specialization, and innovation performance. *Journal of Management 52*(8), 3201–3248.

Hegde, D., K. Herkenhoff, and C. Zhu (2023). Patent publication and innovation. *Journal of Political Economy 131*(7), 1845–1903.

Hémous, D. and M. Olsen (2021). Directed technical change in labor and environmental economics. *Annual Review of Economics 13*, 571–597.

Hicks, J. R. (1932). *The Theory of Wages*. London: Macmillan.

Hill, R. and C. Stein (2025). Race to the bottom: Competition and quality in science. *The Quarterly Journal of Economics 140*(2), 1111–1185.

Hill, R., Y. Yin, C. Stein, X. Wang, D. Wang, and B. F. Jones (2025). The pivot penalty in research. *Nature 642*, 999–1005.

Huitfeldt, I., A. R. Kostøl, J. Nimczik, and A. Weber (2023). Internal labor markets: A worker flow approach. *Journal of Econometrics 233*(2), 661–688.

Huneeus, F., B. Larrain, M. Larrain, and M. Prem (2021). The internal labor markets of business groups. *Journal of Corporate Finance 69*, 102017.

Jadidzadeh, A. and A. Serletis (2017). How does the U.S. natural gas market react to demand and supply shocks in the crude oil market? *Energy Economics 63*, 66–74.

Jaffe, A. B. (1986). Technological opportunity and spillovers of R&D: Evidence from firms' patents, profits, and market value. *American Economic Review 76*(5), 984–1001.

Jarociński, M. and P. Karadi (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics 12*(2), 1–43.

Johnstone, N., I. Haščič, and D. Popp (2010). Renewable energy policies and technological innovation: Evidence based on patent counts. *Environmental and Resource Economics 45*(1), 133–155.

Jones, B. F. (2009). The burden of knowledge and the "death of the renaissance man": Is innovation getting harder? *The Review of Economic Studies 76*(1), 283–317.

Jones, C. I. (2023). Recipes and economic growth: A combinatorial march down an exponential tail. *Journal of Political Economy 131*(8), 1994–2031.

Kambourov, G. and I. Manovskii (2009). Occupational specificity of human capital. *International Economic Review 50*(1), 63–115.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review 111*(4), 1092–1125.

Kaplan, S. and K. Vakili (2015). The double-edged sword of recombination in breakthrough innovation. *Strategic Management Journal 36*(10), 1435–1457.

Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021). Measuring technological innovation over the long run. *American Economic Review: Insights 3*(3), 303–320.

Kim, K., I. Guler, and S. Karim (2024). Who gets redeployed? inventor characteristics and resource redeployment decisions. *Strategic Management Journal*. DOI: 10.1002/smj.3661.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017). Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics 132*(2), 665–712.

Krieger, J., D. Li, and D. Papanikolaou (2022). Missing novelty in drug development. *The Review of Financial Studies 35*(2), 636–679.

Lazear, E. P. (2005). Entrepreneurship. *Journal of Labor Economics 23*(4), 649–680.

Lazear, E. P. (2009). Firm-specific human capital: A skill-weights approach. *Journal of Political Economy 117*(5), 914–940.

Lazear, E. P. and P. Oyer (2004). Internal and external labor markets: A personnel economics approach. *Labour Economics 11*(5), 527–554.

Lise, J. and F. Postel-Vinay (2020). Multidimensional skills, sorting, and human capital accumulation. *American Economic Review 110*(8), 2328–2376.

Ma, S., W. Wang, and Y. Wu (2024). Steering labor mobility through innovation. Working paper, National Bureau of Economic Research.

Maksimovic, V. and G. Phillips (2002). Do conglomerate firms allocate resources inefficiently across industries? Theory and evidence. *Journal of Finance 57*(2), 721–767.

Melero, E. and N. Palomeras (2015). The Renaissance man is not dead! the role of generalists in teams of inventors. *Research Policy 44*(1), 154–167.

Minni, V. (2026). Making the invisible hand visible: Managers and the allocation of workers to jobs. *Quarterly Journal of Economics 141*(3), 1871–1920.

Mokyr, J. (1990). *The Lever of Riches: Technological Creativity and Economic Progress*. New York: Oxford University Press.

Naaraayanan, S. L., K. Sachdeva, and V. Sharma (2026). The real effects of environmental activist investing. *The Review of Financial Studies*. Advance article, published online August 14, 2026.

Nagaoka, S. and J. P. Walsh (2009). The R&D process in the U.S. and japan: Major findings from the RIETI–Georgia Tech inventor survey. Discussion Paper 09-E-010, Research Institute of Economy, Trade and Industry.

Neal, D. (1995). Industry-specific human capital: Evidence from displaced workers. *Journal of Labor Economics 13*(4), 653–677.

Newell, R. G., A. B. Jaffe, and R. N. Stavins (1999). The induced innovation hypothesis and energy-saving technological change. *The Quarterly Journal of Economics 114*(3), 941–975.

Park, Y. and J. Yoon (2017). Application technology opportunity discovery from technology portfolios: Use of patent classification and collaborative filtering. *Technological Forecasting and Social Change 118*, 170–183.

Poletaev, M. and C. Robinson (2008). Human capital specificity: Evidence from the Dictionary of Occupational Titles and displaced worker surveys, 1984–2000. *Journal of Labor Economics 26*(3), 387–420.

Popp, D. (2002). Induced innovation and energy prices. *American Economic Review 92*(1), 160–180.

Ramey, V. A. and S. Zubairy (2018). Government spending multipliers in good times and in bad: Evidence from US historical data. *Journal of Political Economy 126*(2), 850–901.

Silos, P. and E. Smith (2015). Human capital portfolios. *Review of Economic Dynamics 18*(3), 635–652.

Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. von Wachter (2019). Firming up inequality. *Quarterly Journal of Economics 134*(1), 1–50.

Stadler, C., C. E. Helfat, and G. Verona (2022). Transferring knowledge by transferring individuals: Innovative technology use and organizational performance in multiunit firms. *Organization Science 33*(1), 253–274.

Stein, J. C. (1997). Internal capital markets and the competition for corporate resources. *Journal of Finance 52*(1), 111–133.

Tate, G. and L. Yang (2015). The bright side of corporate diversification: Evidence from internal labor markets. *Review of Financial Studies 28*(8), 2203–2249.

Villar, J. A. and F. L. Joutz (2006, October). The relationship between crude oil and natural gas prices. Technical report, U.S. Energy Information Administration, Office of Oil and Gas.

Weitzman, M. L. (1996). Hybridizing growth theory. *American Economic Review 86*(2), 207–212.

Weitzman, M. L. (1998). Recombinant growth. *The Quarterly Journal of Economics 113*(2), 331–360.

Yamaguchi, S. (2012). Tasks and heterogeneous human capital. *Journal of Labor Economics 30*(1), 1–53.

Zamani, N. (2016). How the crude oil market affects the natural gas market: Demand and supply shocks. *International Journal of Energy Economics and Policy 6*(2), 217–221.

Figures

Figure 1: The effect of oil supply shocks on real energy prices, 1992–2020

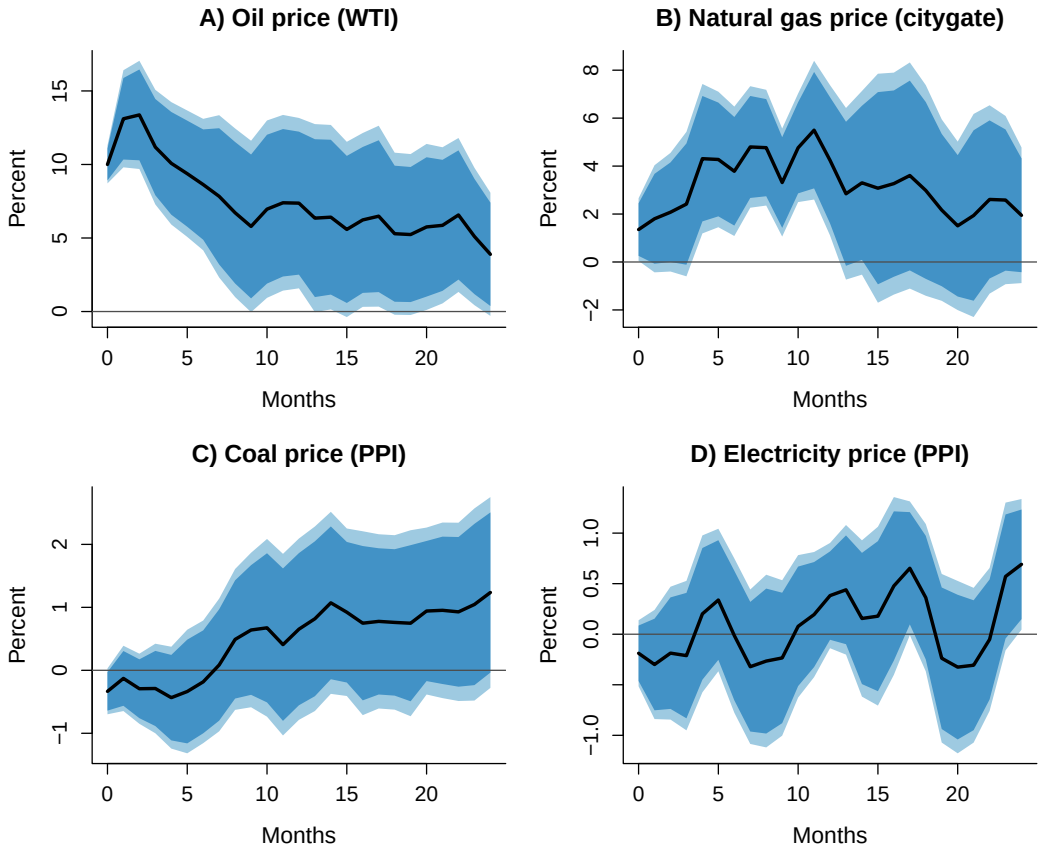

Notes: The figure shows local projection impulse responses of monthly real fuel prices to the Känzig (2021) oil supply news shock at horizons of 0 to 24 months, estimated on monthly data for 1992 to 2020. Prices are the WTI crude oil price (Panel A), the EIA citygate natural gas price (Panel B), and the BLS producer price indexes for coal (Panel C) and industrial electric power (Panel D), deflated by the CPI. Each regression controls for one lag of the log price level. Standard errors are Newey-West, with bandwidth equal to the response horizon plus one month. Responses are normalized so that the real oil price rises 10 percent on impact. Shaded areas are 90 and 95 percent confidence intervals.

Figure 2: Breadth and inventive output

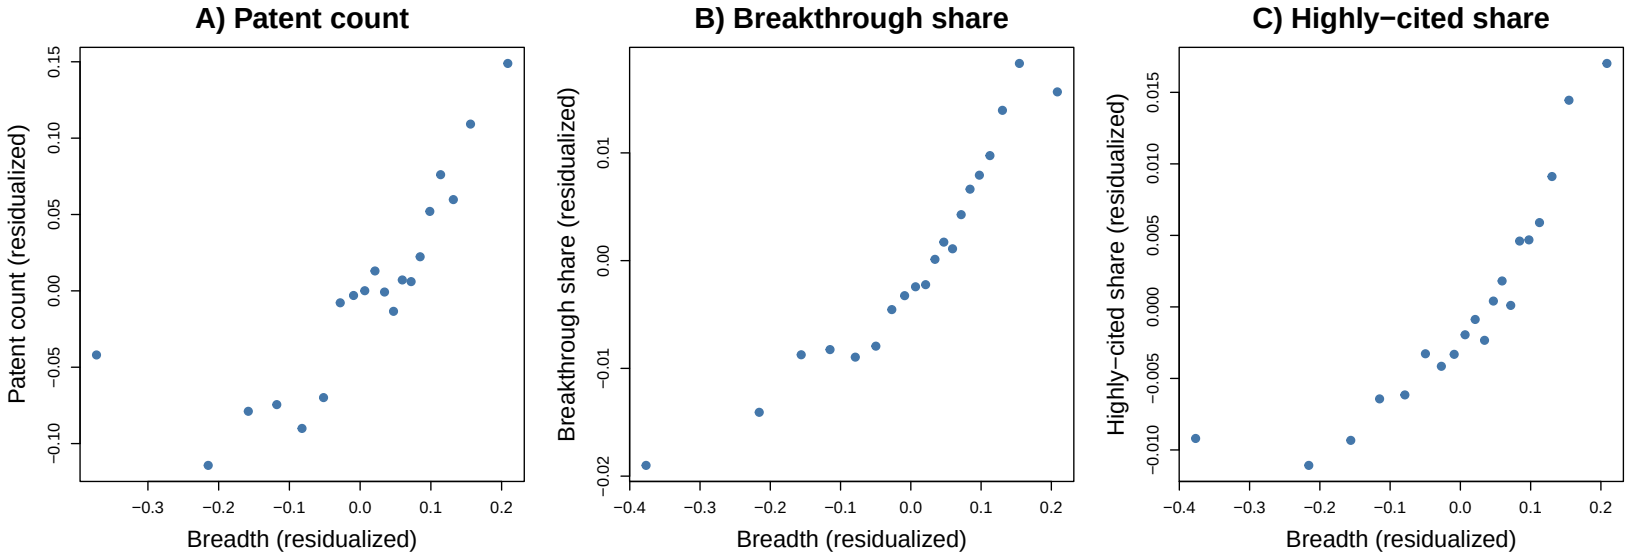

Note: The figure plots residualized annual patent count (Panel A), breakthrough share (Panel B), and highly cited share (Panel C) against residualized breadth. All variables are residualized on firm-by-year fixed effects, within-year deciles of lagged patent stock, career age, and career age squared. Points are means within twenty equal-sized bins of residualized breadth. The sample covers inventor-years from 1992 to 2021. Panel B is restricted to 1992–2016 due to outcome-data availability. Panels B and C are restricted to inventor-years with at least one patent.

Figure 3: The effect of the shock across the breadth distribution

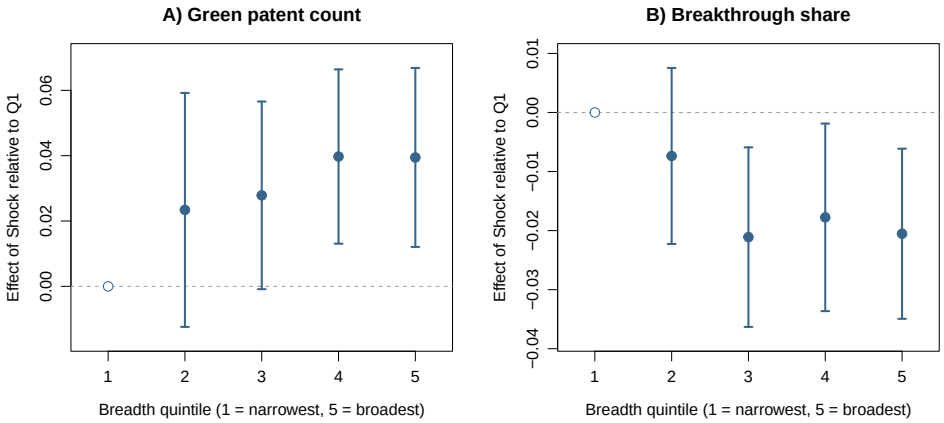

Note: This figure reports the estimated difference in the effect of a one-standard-deviation energy-cost shock between inventors in each breadth quintile and inventors in the narrowest quintile. Panel A reports PPML estimates for green patent counts, and Panel B reports OLS estimates for the breakthrough share of green patents. All specifications control for the lagged green patent share, lagged total patent stock, career age and its square, and include inventor and firm-by-year fixed effects. Confidence intervals are 95%, and standard errors are clustered two-way by firm and inventor.

Figure 4: The effect of the shock across the scientific importance distribution

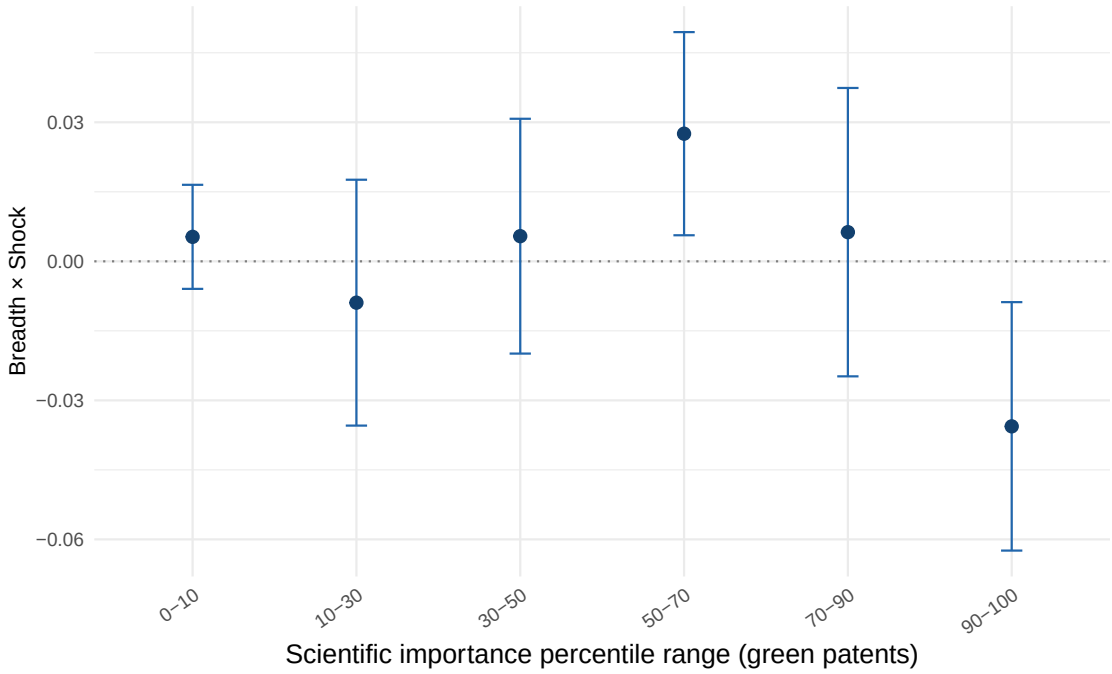

Note: The figure reports the coefficient on the interaction of lagged breadth and the lagged energy-cost shock in regressions for the share of an inventor's green patents in each indicated scientific-importance percentile range. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. The sample includes inventor-years from 1992 to 2016 with at least one green patent. All specifications control for lagged green patent share, lagged patent stock, career age and its square, and include inventor and firm-by-year fixed effects. Confidence intervals are 95%, and standard errors are clustered two-way by firm and inventor.

Tables

Table 1: Summary statistics: firms

	Mean	SD	p25	p50	p75
Panel A) Innovation					
Patent count	59.7	281	1	6	25
Green count	4.6	26.7	0	0	1
Green share	0.103	0.21	0	0	0.1
Patent stock	332.9	1,496.9	10.8	37.4	146.5
Inventors	95.4	406.3	2	12	46
Panel B) Financials and composition					
Total assets ($M)	5,569.8	17,629.3	101.2	530.5	2,789.3
Liquidity	0.291	0.276	0.06	0.193	0.464
Leverage	0.205	0.223	0.013	0.158	0.31
Profitability	-0.015	0.345	-0.05	0.101	0.161
R&D intensity	0.136	0.2	0.022	0.064	0.154
Manufacturing share of sales (%)	95.8	15.8	100	100	100
Panel C) Energy-expenditure shares (%)					
Oil share	2.1	1.8	1.7	1.7	2.2
Gas share	22.1	12.9	12.4	16.4	39.6
Electricity share	74.2	14.9	55.6	80.5	86
Coal share	1.6	2.4	0	0.8	3.1

Note: This table presents summary statistics for the firm analysis sample and includes 21,012 observations covering 1,245 firms. Panel A reports innovation outcomes, Panel B reports financials and composition, and Panel C reports energy cost exposure variables. Variable definitions are in Appendix A.

Table 2: Summary statistics: inventors

	Mean	SD	p25	p50	p75
Panel A) Inventor level					$N = 68,606$
Total patents	17.95	43.61	3.00	7.00	18.00
Green count	3.83	7.42	1.00	2.00	4.00
Green share	0.43	0.34	0.14	0.33	0.67
Career age (years)	14.62	9.49	7.00	13.00	21.00
Number of firms	1.30	0.63	1.00	1.00	1.00
Single-firm inventor	0.77	0.42	1.00	1.00	1.00
Panel B) Inventor-year level					$N = 650,688$
Green count	0.40	1.10	0.00	0.00	1.00
Total patents	1.89	5.36	0.00	1.00	2.00
Green share$_{t-1}$	0.22	0.30	0.00	0.08	0.33
Breadth (1−HHI)	0.78	0.21	0.74	0.86	0.92
Total patent stock	8.57	18.09	1.61	3.91	9.25
Career age (years)	11.51	8.49	5.00	10.00	17.00

Note: This table presents summary statistics for the inventor analysis sample. Panel A reports inventor-level variables and includes 68,606 inventors. Panel B reports inventor-year variables and includes 650,688 observations. Variable definitions are in Appendix A.

Table 3: The effect of energy-cost shocks on firm innovation

	Green patent count		Other patent count	
	(1)	(2)	(3)	(4)
Shock_{t-1}	0.0359*	0.0377**	0.0043	0.0056
	(0.0184)	(0.0188)	(0.0090)	(0.0092)
	[+3.7%]	[+3.8%]	[+0.4%]	[+0.6%]
Firm FE	Y	Y	Y	Y
Industry × year FE	Y	Y	Y	Y
Firm controls	N	Y	N	Y
Observations	21,012	21,012	21,012	21,012

Note: This table presents Poisson pseudo-maximum-likelihood (PPML) estimates of the firm-level response to the shock. The outcome is the firm's count of green patents in columns (1) and (2) and its count of other patents in columns (3) and (4). Shock is the lagged energy-cost shock, as defined in Section 2.4, standardized to zero mean and unit standard deviation. The bracketed figure below each standard error reports the implied percent change in the outcome per one standard deviation increase in the shock. The sample is green patenting firms in the firm panel with lagged financial controls available, so all four columns share one sample. All columns include firm and industry-by-year fixed effects; columns (2) and (4) add lagged firm controls: log size, liquidity, leverage, profitability, and R&D intensity. Standard errors (in parentheses) are clustered at the firm level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4: The effect of energy-cost shocks on the quantity of green innovation

	Green patent count		Other patent count	
	(1)	(2)	(3)	(4)
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0557*	0.0784***	0.0048	-0.0057
	(0.0308)	(0.0284)	(0.0174)	(0.0210)
	[+1.4%]	[+2.0%]	[+0.1%]	[−0.1%]
Breadth_{t-1}	-0.1540**	0.6456***	-0.1034	0.1903***
	(0.0639)	(0.0674)	(0.0693)	(0.0466)
Observations	585,967	583,525	585,967	584,112
Inventor FE	N	Y	N	Y
Firm×year FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y

Note: This table presents Poisson pseudo-maximum-likelihood (PPML) estimates of equation (2) using the inventor panel. The outcome is the count of the inventor's green patents in columns (1) and (2) and the count of non-green patents in columns (3) and (4). Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. The bracketed row reports the implied percent change in the outcome per one standard deviation increase in the shock, for an inventor one standard deviation above mean breadth. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include firm-by-year fixed effects; columns (2) and (4) add inventor fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 5: Energy-cost shocks and inventor-team composition

	Green patents		Other patents	
Co-inventors	(1) At least one new	(2) None new	(3) At least one new	(4) None new
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0104***	-0.0006	-0.0054	-0.0044*
	(0.0026)	(0.0021)	(0.0037)	(0.0025)
	[+1.8%]	[−0.1%]	[−0.4%]	[−0.4%]
Breadth_{t-1}	0.0411***	-0.0099*	0.0068	0.0591***
	(0.0068)	(0.0056)	(0.0121)	(0.0076)
Inventor FE	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y
Outcome mean	0.142	0.116	0.368	0.308
Adjusted R^2	0.131	0.223	0.173	0.313
Observations	583,525	583,525	583,525	583,525

Note: This table presents OLS estimates of equation (2) using the inventor panel, with indicators for filing at least one patent of a given type in year t as the outcome. Columns (1) and (2) split the inventor's green patents by team composition, and columns (3) and (4) do the same for non-green patents. A co-inventor is new if the inventor and the co-inventor had not appeared together on any earlier patent, at any assignee. Columns (1) and (3) use an indicator for filing at least one patent with two to fifteen inventors on which at least one co-inventor is new. Columns (2) and (4) use an indicator for filing at least one patent without a new co-inventor, which includes team patents on which every co-inventor had previously patented with the inventor and solo patents. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4, standardized to be zero-mean and have a standard deviation of 1. The bracketed row reports the implied change in the outcome per one-standard-deviation increase in the shock, for an inventor one standard deviation above mean breadth, as a percentage of the outcome mean. Appendix Table A4 reports the corresponding PPML estimates with patent counts as outcomes. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 6: The effect of energy-cost shocks on the scientific value of green innovation

	Green breakthrough share			Other breakthrough share		
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0228*	-0.0356***	-0.0363**	-0.0056*	-0.0033	-0.0032
	(0.0129)	(0.0136)	(0.0142)	(0.0030)	(0.0039)	(0.0039)
	[−4.6%]	[−7.1%]	[−7.2%]	[−1.4%]	[−0.8%]	[−0.8%]
Breadth_{t-1}	0.0723***	0.0145	0.0325*	0.0568***	0.0272***	0.0379***
	(0.0160)	(0.0168)	(0.0183)	(0.0129)	(0.0067)	(0.0071)
Observations	65,541	65,541	65,541	223,918	223,918	223,918
Adjusted R^2	0.201	0.415	0.426	0.090	0.246	0.252
Outcome mean	0.126	0.126	0.126	0.104	0.104	0.104
Inventor FE	N	Y	Y	N	Y	Y
Firm×year FE	Y	Y	Y	Y	Y	Y
Prior breakthrough share control	N	N	Y	N	N	Y
Controls	Y	Y	Y	Y	Y	Y

Note: This table presents OLS estimates of equation (3) using the inventor panel, the breakthrough-share analogue of the main quantity table. The outcome is the share of the inventor's green patents that are breakthroughs in columns (1) through (3), and the share of non-green patents that are breakthroughs in columns (4) through (6); a breakthrough is a patent in the top decile of the importance measure (Section 2.2). Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. The bracketed row reports the change in the outcome as a percent of its mean per one standard deviation increase in the shock, for an inventor one standard deviation above mean breadth. The Outcome mean row is the unconditional mean of each outcome. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared. Columns (1) and (4) include firm-by-year fixed effects; columns (2), (3), (5), and (6) add inventor fixed effects; and columns (3) and (6) additionally control for the inventor's lagged breakthrough share. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 7: Dynamics of the effect of energy-cost shocks

	Placebo	Contemporaneous	Baseline	Post-shock		
	(1)	(2)	(3)	(4)	(5)	(6)
	$t-2$	$t-1$	t	$t+1$	$t+2$	$t+3$
Panel A)	Dependent variable: Green patent count					
Breadth $\times$ Shock$_{t-1}$	-0.0159	0.0132	0.0784***	0.0453*	0.0099	0.0409*
	(0.0191)	(0.0236)	(0.0284)	(0.0231)	(0.0212)	(0.0224)
Breadth	0.6512***	0.6493***	0.6456***	0.7656***	0.7860***	0.7357***
	(0.0675)	(0.0671)	(0.0674)	(0.0729)	(0.0697)	(0.0750)
Observations	583,525	583,525	583,525	488,913	421,941	367,766
Panel B)	Dependent variable: Breakthrough share					
Breadth $\times$ Shock$_{t-1}$	-0.0103	0.0073	-0.0356***	0.0025	0.0156	-0.0136***
	(0.0083)	(0.0052)	(0.0136)	(0.0109)	(0.0111)	(0.0052)
Breadth	0.0144	0.0121	0.0145	0.0295**	0.0181	0.0125
	(0.0163)	(0.0166)	(0.0168)	(0.0137)	(0.0122)	(0.0148)
Adjusted R^2	0.415	0.415	0.415	0.400	0.398	0.395
Observations	65,541	65,541	65,541	52,982	45,144	38,953
Firm$\times$year FE	Y	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y

Note: This table presents estimates of equations (2) and (3) using the inventor panel, tracing the response across outcome horizons $t-2$ to $t+3$ relative to the baseline shock year, $t-1$. Panel A reports PPML estimates with the count of the inventor's green patents as the outcome. Panel B reports OLS estimates with the share of green patents in the top decile of the importance measure as the outcome (Section 2.2), which we define as a breakthrough patent. In columns (1) to (3), breadth and controls are measured one year before the outcome and the shock is measured in $t-1$, so the outcome is measured one year before the shock in column (1), in the same year in column (2), and one year after it in column (3). In columns (4) to (6), breadth, the shock, and controls are held at their $t-1$ values. Each regression also includes the interaction of breadth with every shock that occurs after $t-1$ and before the outcome year. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include firm-by-year and inventor fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 8: Alternative explanations: prior green innovation experience

Green patenting experience measure:	(1) Green share	(2) Prior green ≥ 1	(3) Prior green ≥ 2	(4) Log green stock	(5) Green proximity
Panel A)	Dependent variable: Green patent count				
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0734***	0.0486*	0.0547*	0.0476*	0.0761***
	(0.0275)	(0.0292)	(0.0282)	(0.0256)	(0.0260)
	[+1.9%]	[+1.2%]	[+1.4%]	[+1.2%]	[+1.9%]
$\text{Green experience}_{t-1} \times \text{Shock}_{t-1}$	0.0349**	0.0164***	0.0125***	0.0132***	-0.0034
	(0.0169)	(0.0047)	(0.0049)	(0.0041)	(0.0077)
Observations	583,525	583,525	583,525	583,525	583,525
Panel B)	Dependent variable: Breakthrough share				
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0365***	-0.0386***	-0.0344**	-0.0356**	-0.0348***
	(0.0134)	(0.0145)	(0.0139)	(0.0146)	(0.0131)
	[−7.3%]	[−7.7%]	[−6.9%]	[−7.1%]	[−6.9%]
$\text{Green experience}_{t-1} \times \text{Shock}_{t-1}$	-0.0103***	0.0020	-0.0011*	0.0002	-0.0025*
	(0.0034)	(0.0020)	(0.0007)	(0.0012)	(0.0014)
Outcome mean	0.124	0.124	0.124	0.124	0.124
Adjusted R^2	0.416	0.415	0.415	0.416	0.415
Observations	65,541	65,541	65,541	65,541	65,541
Inventor FE	Y	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y	Y

Note: This table presents estimates of equation (2) using the inventor panel, adding one measure of prior green innovation experience and its interaction with the shock alongside the breadth interaction. Panel A reports Poisson pseudo-maximum-likelihood (PPML) estimates with the count of the inventor's green patents as the outcome. Panel B reports OLS estimates with the share of green patents in the top decile of the importance measure as the outcome (Section 2.2), which we define as a breakthrough patent. Each column uses a different measure of green experience: in column (1) the share of the inventor's prior patents that are green, in columns (2) and (3) indicators for having filed at least one and at least two green patents through year $t-1$, in column (4) the log of one plus the inventor's lagged green patent stock, and in column (5) a lagged cosine-similarity measure of proximity to green knowledge following Jaffe (1986). That measure enters both on its own and interacted with the shock, and replaces prior green share as the control in these specifications. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. The "Outcome mean" row is the unconditional mean of the Panel B outcome. All columns control for the log of one plus the inventor's lagged patent stock, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 9: Alternative explanations: selection on inventor ability

	Baseline		Human-capital proxy			
	(1)	(2)	(3)	(4)	(5)	(6)
	Full sample	Résumé-matched sample	Top-20 university	Top-50 university	Master's degree	PhD
Panel A)	Dependent variable: Green patent count					
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0784***	0.1097***	0.1067**	0.1056**	0.1105***	0.1101**
	(0.0284)	(0.0424)	(0.0431)	(0.0439)	(0.0426)	(0.0438)
	[+2.0%]	[+2.8%]	[+2.7%]	[+2.7%]	[+2.8%]	[+2.8%]
Human-capital proxy $\times \text{Shock}_{t-1}$	–	–	0.0204**	0.0183**	0.0051	-0.0016
			(0.0080)	(0.0088)	(0.0126)	(0.0102)
Observations	583,525	269,212	269,212	269,212	269,212	269,212
Panel B)	Dependent variable: Breakthrough share					
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0356***	-0.0500**	-0.0501**	-0.0496**	-0.0479**	-0.0488**
	(0.0136)	(0.0207)	(0.0208)	(0.0213)	(0.0207)	(0.0209)
	[−7.1%]	[−10.0%]	[−10.0%]	[−9.9%]	[−9.6%]	[−9.7%]
Human-capital proxy $\times \text{Shock}_{t-1}$	–	–	0.0008	-0.0014	0.0095	-0.0038
			(0.0027)	(0.0043)	(0.0073)	(0.0070)
Outcome mean	0.124	0.124	0.124	0.124	0.124	0.124
Adjusted R^2	0.415	0.411	0.411	0.411	0.411	0.411
Observations	65,541	26,725	26,725	26,725	26,725	26,725
Inventor FE	Y	Y	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y	Y	Y

Note: This table presents estimates of equations (2) and (3) using the inventor panel. Panel A reports PPML estimates with the count of the inventor's green patents as the outcome. Panel B reports OLS estimates with the share of green patents in the top decile of the importance measure as the outcome (Section 2.2), which we define as a breakthrough patent. Column (1) uses the baseline sample. Column (2) restricts the sample to inventors matched to Revelio résumé data. Columns (3) to (6) add, respectively, an interaction between the shock and indicators for a degree from a top-20 university, a degree from a top-50 university, a master's degree, and a PhD, measured from Revelio education records (Appendix B.6). Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. The bracketed rows report the implied percentage change in green patent count in Panel A and the change in the breakthrough share as a percentage of its mean in Panel B, per one-standard-deviation increase in the shock for an inventor one standard deviation above mean breadth. The Outcome mean row is the unconditional mean of the Panel B outcome. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 10: Alternative explanations: cross-firm mobility

	(1) Baseline	(2) Same firm, last 3 years	(3) Same firm throughout observed career
Panel A)	Dependent variable: Green patent count		
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0784^{***}	0.0701^{***}	0.0739^{***}
	(0.0284)	(0.0242)	(0.0244)
	[+2.0%]	[+1.8%]	[+1.9%]
Breadth_{t-1}	0.6456^{***}	0.6323^{***}	0.5975^{***}
	(0.0674)	(0.0662)	(0.0705)
Observations	583,525	502,473	388,578
Panel B)	Dependent variable: Breakthrough share		
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0356^{***}	-0.0392^{**}	-0.0417^{**}
	(0.0136)	(0.0156)	(0.0166)
	[−7.1%]	[−7.8%]	[−8.3%]
Breadth_{t-1}	0.0145	0.0072	-0.0018
	(0.0168)	(0.0190)	(0.0200)
Outcome mean	0.124	0.125	0.130
Adjusted R^2	0.415	0.414	0.402
Observations	65,541	59,025	46,285
Inventor FE	Y	Y	Y
Firm×year FE	Y	Y	Y

Note: This table presents estimates of equations (2) and (3) using the inventor panel on samples that restrict cross-firm mobility. Panel A reports PPML estimates with the count of the inventor's green patents as the outcome. Panel B reports OLS estimates with the share of green patents in the top decile of the importance measure as the outcome (Section 2.2), which we define as a breakthrough patent. Column (1) uses the baseline sample. Column (2) retains inventor-years in which the inventor patents with the same firm in the preceding three years. Column (3) retains inventors who patent with the same firm throughout their observed career. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. The bracketed rows report the implied percentage change in green patent count in Panel A and the change in the breakthrough share as a percentage of its mean in Panel B, per one-standard-deviation increase in the shock for an inventor one standard deviation above mean breadth. The Outcome mean row is the unconditional mean of the Panel B outcome. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 11: Alternative explanations: technological pivots

		Pivot distance		
	(1) Baseline	(2) Cited CPC past 3 years	(3) Cited CPC all prior	(4) Own CPC
	Dependent variable: Breakthrough share			
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0356***	-0.0274***	-0.0331**	-0.0374***
	(0.0136)	(0.0103)	(0.0133)	(0.0136)
	[−7.1%]	[−5.8%]	[−7.0%]	[−7.5%]
Pivot distance × Shock	–	-0.0001	0.0021	0.0031*
		(0.0024)	(0.0016)	(0.0018)
Outcome mean	0.126	0.118	0.119	0.126
Adjusted R^2	0.415	0.440	0.435	0.416
Observations	65,541	46,001	53,530	65,541
Inventor FE	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y

Note: This table presents OLS estimates of equation (3) using the inventor panel. The outcome is the share of green patents in the top decile of the importance measure (Section 2.2), which we define as a breakthrough patent. Column (1) uses the baseline specification. Columns (2) to (4) add a measure of pivot distance and its interaction with the shock. Following Hill et al. (2025), column (2) measures pivot distance using the CPC main-group mix of the green patent's backward citations relative to the inventor's citation mix over the preceding three years. Column (3) uses the inventor's full prior citation history, and column (4) compares the CPC main groups assigned to the green patent with the inventor's full prior patent portfolio. Pivot distance enters in levels and interacted with the shock. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock and pivot-distance measures are standardized to be zero-mean and have a standard deviation of 1. The bracketed row reports the change in the outcome as a percentage of its mean per one-standard-deviation increase in the shock for an inventor one standard deviation above mean breadth. The Outcome mean row is the unconditional mean of the outcome. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A Variable Definitions

This appendix defines the main variables used in the paper. Variable names match the tables and text. For Compustat variables, annual data items appear in *italics*.

A.1 Knowledge measures

Breadth: one minus the Herfindahl index of an inventor's prior patents across non-Y02 CPC main groups, lagged one year. A patent assigned to multiple main groups is divided equally across them. (Source: PatentsView.)

Green share: share of an inventor's prior patents classified under CPC Y02, lagged one year. (Source: PatentsView.)

Log(1+patent stock): log of one plus the inventor's depreciated prior-patent count. The first and 99th percentiles of the count are winsorized before the transformation. (Source: PatentsView.)

Prior green ≥ 1 / ≥ 2: indicators equal to one if the inventor filed at least one / two green patents before the outcome year. (Source: PatentsView.)

Log(1+green patent stock): log of one plus the inventor's prior green-patent count. (Source: PatentsView.)

Green proximity: lagged, depreciated cosine similarity of the inventor's prior patent portfolio to the recent green-technology profile of the inventor's industry. (Source: PatentsView.)

Pivot distance: one minus the cosine similarity between the CPC main-group profile of a focal green patent and the inventor's prior technology profile. It is constructed three ways: (1) following Hill et al. (2025), from the CPC main-group profiles of the focal patent's backward citations and the inventor's citations over the preceding three years, (2) using the inventor's full prior citation history, and (3) from the focal patent's own CPC main-group profile and the inventor's full prior patent portfolio. (Source: PatentsView.)

A.2 Innovation outcomes

Green patent count: number of patents classified under CPC Y02 filed by the inventor or firm in a year. (Source: PatentsView.)

Non-green patent count: number of patents not classified under CPC Y02 filed by the inventor or firm in a year. (Source: PatentsView.)

Breakthrough share: share of an inventor's green or non-green patents in the top decile of the KPST importance measure within NAICS4-by-filing-year cells. Importance is the log ratio of forward to backward patent-text similarity. (Sources: Kelly et al., 2021; PatentsView.)
Importance percentile: mean NAICS4-by-filing-year percentile of KPST importance across an inventor's green patents in a year. (Sources: Kelly et al., 2021; PatentsView.)
Highly cited share: share of an inventor's green patents in the top decile of five-year forward citations within NAICS4-by-filing-year cells. (Source: PatentsView.)
Economic value: mean log real-dollar value of an inventor's green patents filed in a year. Patent values are estimated from firms' stock-market reactions around patent grants (Kogan et al., 2017).

A.3 Energy-cost shock

Shock: predicted change in a firm's oil and natural-gas costs as a share of sales, constructed by interacting predetermined input exposures with oil-supply-news-induced predicted price changes, lagged one year, and z-scored. (Sources: Känzig, 2021; EIA MECS; NBER-CES; Compustat Segments.)
Oil and natural-gas exposure: baseline cost of each input as a share of firm sales. It combines the firm's first-observed manufacturing-segment mix with 1991 industry energy intensity and input cost shares. (Sources: EIA MECS; NBER-CES; Compustat Segments.)

A.4 Firm characteristics

Size: logarithm of total assets (*at*). (Source: Compustat.)
Liquidity: cash and short-term investments (*che*) divided by total assets (*at*). (Source: Compustat.)
Leverage: total debt (*dltt*+*dlc*) divided by total assets (*at*). (Source: Compustat.)
Profitability: EBITDA (*oibdp*) divided by total assets (*at*). (Source: Compustat.)
R&D intensity: research and development expenditure (*xrd*) divided by total assets (*at*). (Source: Compustat.)
CAPX/assets: capital expenditure (*capx*) divided by total assets (*at*). (Source: Compustat.)
Manufacturing share: manufacturing-segment sales divided by total segment sales. (Source: Compustat Segments.)

A.5 Inventor characteristics

Career age: outcome year minus the inventor's first patenting year. (Source: PatentsView.)
Prior breakthrough: indicator for at least one prior patent in the top decile of the KPST importance measure. (Sources: Kelly et al., 2021; PatentsView.)
Top-20 university: indicator equal to one if any recorded degree is from a top-20 university. (Source: Revelio data.)
PhD: indicator equal to one if the inventor has a recorded doctoral degree. (Source: Revelio data.)
Firm seniority: highest seniority category across the inventor's active employer positions, using the most recent value observed through the prior year. (Source: Revelio data.)

B Data Construction Details

This appendix collects the second-order construction decisions, mechanical conventions, and accounting detail behind the measures in Section 2.

B.1 Industry and employer assignment

Industry coverage and matching. The 1991 MECS reports fuel use for 62 manufacturing SIC industries, including 40 four-digit, 2 three-digit, and 20 two-digit industries. We map each segment to the finest SIC industry level available in MECS. Energy intensity is matched at four-digit SIC, with a two-digit fallback when a segment lacks a four-digit match. The NBER-CES Manufacturing Industry Database provides energy expenditure as a share of the value of shipments for 459 four-digit SIC industries.

MECS cost shares. MECS reports fuel quantities, which we convert to Btu and then to fuel costs using 1991 EIA national-average industrial prices per MMBtu (residual fuel oil 2.30, distillate 4.50, LPG 5.00, natural gas 2.62, coal 1.45, and electricity 13.90). We calculate each industry's energy-input expenditure shares from these costs.

ALP CPC–NAICS crosswalk. We define a patent's main CPC code as its first-listed inventive CPC code. The USPTO describes this first classification as the inventive code

that most adequately represents the invention as a whole (USPTO MPEP §905). We use the ALP crosswalk of Goldschlag et al. (2020), which links the four-character subclass of this code to NAICS 2002 industries with probability weights estimated from algorithmic matching of patent text to industry descriptions. We use the version that excludes service industries and assign each patent the highest-weight industry, breaking ties lexically.

Firm sample construction. The CRSP–Compustat linking procedure maps 3.2 million KPSS patents, or 94.1 percent of the KPSS patent records, and assigns each mapped patent to a unique Compustat firm. Of 26,555 US-incorporated Compustat firms observed from 1992 to 2021, 8,403 have a usable manufacturing-segment exposure. After requiring that these firms have at least one USPTO patent that we can match through the KPSS crosswalk, the linked firm sample comprises 4,678 firms and 66,279 firm-year observations. Restricting to firms that file at least one green patent over 1992–2021 leaves 1,526 firms and 27,154 firm-years. Requiring nonmissing lagged financial controls leaves 1,419 firms and 22,823 firm-years. Finally, dropping firm-years in fixed-effect groups with no variation in green or other patenting, which do not contribute to the Poisson estimates, yields the final sample of 1,245 firms and 21,012 firm-years. A further 3.1 percent of otherwise eligible manufacturing firms lack a usable manufacturing segment. We exclude these firms rather than proxy their energy input mix using their Compustat header SIC, to avoid measurement error and definition inconsistency across firms. In 1992–2020, these excluded firms are much smaller, with mean total assets of $108 million, compared with $2.0 billion for firms with a usable manufacturing segment.

Inventor sample construction. PatentsView's algorithmically disambiguated inventor identities[22] identify 747,950 inventors who appear on 1,362,438 patents filed from 1992 to 2021 at firms in the linked firm sample. Retaining inventors who file at least one green patent over 1992–2021 and applying the lag and shock requirements yields the final inventor sample of 650,688 inventor-years covering 68,606 inventors.

Inventor–firm assignment. The firm assignment determines the energy-cost shock attached to each inventor. In patenting years, each inventor is assigned the shock of the

[22]PatentsView disambiguates inventors using a learned pairwise-similarity model based on patent attributes, including names, locations, co-inventors, assignees, and technology classifications, then resolves matches into inventor clusters. Across two staff-curated benchmarks and the Li et al. benchmark, the algorithm has estimated precision of 87–91 percent and recall of 91–96 percent (Monath et al., 2021; Binette et al., 2022).

firm with which they patent in year t. The assigned firm is the one the inventor patents with most in year t, with ties broken by the firm the inventor patented with most in prior years. To assign shocks in non-patenting years, we construct an inventor–firm panel over the interval from the inventor's first to last patent at a sample firm. We retain the assignment in gaps bounded by patents at the same firm. For a move from firm A to firm B, where the first patent at B is filed in year t, earlier gap years remain assigned to A and the assignment to B begins in $t-1$. Thus, an inventor who patents at firm A in 2000 and then at firm B in 2004, with no intervening patent, is assigned to A from 2000 through 2002 and to B from 2003 through 2004. We omit years with patenting outside the sample firms, unresolved-tie years, and gaps for which the patent history does not identify a sample-firm attachment. Of 2,154,168 patenting inventor-years at sample firms over 1992–2021, 98.1 percent involve a single firm. Of the remaining 40,295, 58.8 percent are resolved by the current-year rule and 32.0 percent by the prior-years rule, and the 3,692 still-tied inventor-years (0.17 percent) are dropped, leaving 2,150,476 assigned inventor-years. Table 10 shows that the results are robust to excluding inventors who move across firms, so they do not hinge on these assignment decisions.

B.2 Breadth and prior green experience

CPC hierarchy and level of aggregation. The Cooperative Patent Classification organizes patent technologies hierarchically into sections, classes, subclasses, main groups, and subgroups (USPTO MPEP §905). It contains 662 subclasses (e.g., photographic apparatus, G03B), 9,987 main groups (e.g., cameras, G03B 19/00, within G03B), and 210,347 subgroups (e.g., still-picture cameras, G03B 19/02, within G03B 19/00). A subclass covers a broad technology family. Main-group titles define distinct fields of subject matter within a subclass, and subgroups provide more specific subdivisions of a main group. Appendix Table A8 illustrates the distinction. Within photographic apparatus (G03B), cameras (G03B 19/00) and projectors (G03B 21/00) are distinct main groups, while still-picture cameras (G03B 19/02) and motion-picture cameras (G03B 19/18) are subgroups within cameras. In the patent universe used to construct our knowledge measures, patents carry 2.1 CPC subclasses, 4.1 main groups, and 7.4 subgroups on average. The corresponding medians are 2, 3, and 5. We use main groups as the baseline technology fields, following Verhoeven et al. (2016); Hegde et al. (2023); Hill et al. (2025). They retain substantial variation across technologies without splitting fields into many narrow subgroup categories. We show robustness to the coarser subclass level, as in

Fusillo (2023), and to the finer subgroup level, as in Fleming (2001); Arts and Fleming (2018).

B.3 Green specialization and proximity measures

The main analysis captures green experience with the lagged share of an inventor's prior patents that are green. We construct the following cosine-similarity-based green-specialization measure, which uses both prior green and non-green patents, for supplementary analysis. The measure captures the proximity of an inventor's prior experience to the green knowledge domain.

Construction. To measure specialization in green technologies, we first define the green knowledge domain. The CPC classifies technologies hierarchically into subclasses, main groups, and subgroups. We use CPC main groups, which distinguish meaningful technological areas while avoiding the very fine distinctions among subgroups. Appendix B describes the CPC hierarchy and gives examples at each level of aggregation.

For every four-digit NAICS industry j and filing year t, we construct a green knowledge profile from green patents filed during the preceding five years. We retain the most common CPC main groups until they account for 90 percent of the profile's mass, then normalize the profile to unit length. Each patent is represented by a unit-length vector of its CPC main groups, excluding Y02. Patent-level green proximity is the cosine similarity between this patent vector and its industry's green-knowledge profile.

A high value indicates that a patent draws on the same underlying technologies as recent green invention in its industry. This construction is similar to approaches that use cosine similarity between firms' patent portfolios to measure technological proximity between firms (Jaffe, 1986; Bloom et al., 2013).

Because CPC classifications vary in richness across technological areas, raw green-proximity scores are not directly comparable across industries. We convert each patent's proximity into a percentile rank among patents filed in the same four-digit NAICS industry and year. Following Hall et al. (2001), we set annual depreciation to 15 percent and measure each inventor's green specialization as the patent-stock-weighted average of those percentiles.

Validation of the proximity measure. Appendix Table A9 shows that green patents have a mean proximity percentile of 0.714, compared with 0.483 for non-green patents.

The difference is positive and highly statistically significant in each of the ten key industries presented, and overall. The green profile uses only patents filed in the prior five years, so a focal green patent does not confound the profile used to score it.

Appendix Table A10 provides a second check. Conditional on NAICS4-by-filing-year fixed effects, a one-standard-deviation increase in proximity is associated with a 2.5-percentage-point higher probability that a non-green patent receives a citation from a green patent in the following five years. This is 40 percent of the unconditional mean. The corresponding association for green patents is 3.3 percentage points, but only 8 percent of their much higher unconditional mean.

B.4 Persistence of predetermined energy-cost exposure

The exposure measure combines the business mix reported in each firm's first Compustat Segments observation with industry energy parameters constructed from 1991 data:

$$Exposure_{fk} = \sum_{s} w_{fs}\, \eta_{j(s)}\, e_{j(s)k}. \tag{2.4}$$

We hold these inputs fixed when estimating the firm's response to subsequent fuel-price shocks. The three components are a segment's share of firm manufacturing sales, w_{fs}, its industry's energy intensity, $\eta_{j(s)}$, and its industry's cost share for fuel k, $e_{j(s)k}$. This prevents the treatment measure from conflating the effect of higher energy costs with changes in firms' and industries' operations or energy use in response to those costs. Its interpretation is strengthened if the baseline exposure continues to capture persistent differences in firms' cost structures. We therefore examine the persistence of the three underlying inputs, namely firm segment mix, industry energy intensity, and industry fuel-cost shares, as well as the firm-level exposure terms they produce and the resulting predicted shock.

Figures A4, A5, and A6 examine the three components of Equation 2.4 separately. Figure A4 examines the segment-composition term, w_{fs}, over time, while keeping the industry terms fixed. The cosine similarity with the first-observed composition is 0.92 after five years, 0.86 after ten years, and 0.80 after twenty years at the two-digit SIC level. At the more detailed four-digit level, the corresponding similarities are 0.84, 0.73, and 0.61. Figure A5 examines the industry energy-intensity term, $\eta_{j(s)}$, over time: its Pearson correlation with the 1991 value is 0.94 in 2000, 0.89 in 2012, and 0.90 in 2018, and industries in the highest 1991 intensity quintile remain distinctly more energy intensive throughout.

Figure A6 examines the fuel-cost-share term, $e_{j(s)k}$, over time. Natural gas, electricity, and coal remain highly persistent, with 2010 rank correlations with their 1991 values around 0.8 or higher. Oil cost shares are somewhat less stable: their rank correlation falls from 0.90 in 1994 to 0.29 in 1998, when MECS transitioned from SIC to NAICS classifications, then rises to 0.56 in 2002, 0.60 in 2006, and 0.61 in 2010.

Figures A7 and A8 return to the treatment measure itself. Figure A7 varies w_{fs} over time while holding $\eta_{j(s)}$ and $e_{j(s)k}$ at their 1991 values, and reports the resulting persistence of $Exposure_{fk}$ for each fuel. Twenty years after a firm's first observation, the rank correlations between current and baseline exposure are 0.81 for oil, 0.83 for gas, 0.85 for coal, and 0.75 for electricity. Figure A8 rebuilds the predicted shock using $Exposure_{fk}$ calculated from each alternative base year's segment composition, again varying w_{fs} while holding the industry terms and the estimated fuel-price responses fixed. The resulting cross-firm rank correlation with the 1991-based shock exceeds 0.8 in every base year, including 0.81 when exposure is fixed in 2020. These correlations are calculated within calendar year, so they isolate cross-firm exposure differences rather than the common fuel-price path. Overall, the underlying components and the resulting exposure measure are persistent, supporting the use of predetermined exposure to capture durable cross-firm differences in energy-cost exposure.

B.5 Distributions of the key variables

This subsection describes the distributions of the central variables: inventor breadth, firm energy-cost exposure, and the energy-cost shock that combines the second with the oil supply news series.

Breadth. Figure A9 presents the distribution of breadth. Panel A plots the raw measure. The distribution is concentrated toward the upper end of the range, between 0.8 and 0.95, with dispersion across the range below. It displays the usual characteristics of a Herfindahl-based index: a mass point at zero for inventors whose prior patents fall in a single technology field, and smaller spikes at 0.50, 0.67, and 0.75, the values implied by prior patenting spread evenly over exactly two, three, or four fields. These are the familiar small-count properties of such indexes. Panel B plots breadth net of lagged patent stock and career age. The spikes largely disappear and the distribution becomes smooth and two-sided, so neither patenting volume nor career stage accounts for much of the variation used in the specifications. Thus, since we control for both in the regressions,

we argue that the breadth measure captures substantial variation in the dispersion of inventors' experience and knowledge across technological fields net of total patent stock and career age.

Energy-cost exposure. Figure A10 plots variation in energy-cost exposure across industries (Panel A) and within industries (Panel B). Panel A reports mean oil and natural-gas cost exposure by NAICS3 industry. Exposure varies substantially across manufacturing industries: paper manufacturing (1.91 percent of sales), nonmetallic mineral products (1.69 percent), petroleum and coal products (1.53 percent), and primary metals (1.49 percent) have the highest exposure. Food manufacturing (0.61 percent), chemicals (0.60 percent), plastics and rubber products (0.58 percent), and fabricated metal products (0.56 percent) follow, while computer and electronic products average 0.16 percent. Because exposure tracks industry composition in this way, industry-by-year fixed effects are essential to the design. Panel B plots exposure after subtracting its NAICS3 mean. The within-industry distribution has a standard deviation of 0.43 percentage points and an interquartile range of 0.26 percentage points, so firms in the same industry differ materially in their predetermined fuel exposure.

Energy-cost shock. Figure A11 describes the regressor used in the specifications, the product of predetermined exposure and the predicted change in fuel prices attributable to oil supply news. The distribution is concentrated near zero, with a mean of -0.003, a median of -0.001, a standard deviation of 0.046, and an interquartile range of 0.017. Forty-six percent of firm-years have a positive shock and 54 percent have a negative shock. Energy costs therefore both rise and fall over the sample. In a given year, the common oil supply news shock is translated into different firm-level changes in energy costs by predetermined fuel exposure.

B.6 Linking Sample Inventors to Résumé Data

This subsection describes the Revelio data and the procedure used to link them to the inventor panel. These data are used in the supplementary tests of selection on inventor ability in Section 5.

Résumé data. Revelio Labs provides individual-level résumé data compiled from public online professional profiles. The data include employment histories, education, oc-

cupation, and a Revelio-supplied seniority category. Employment records identify the employer, the start and end dates of the position, and the seniority category. Education records identify the institution, degree, and, when reported, field of study.

Matching. We first link the sample firms to Revelio companies. Revelio provides a crosswalk between its company identifier and Compustat `gvkey`. We supplement this mapping with unique exact matches on company names for firms without a `gvkey` match, dropping ambiguous matches. The resulting crosswalk matches 85 percent of the firms at which inventors in the baseline sample patent. Of those, 92 percent are linked through Revelio's `gvkey` crosswalk and 8 percent through the supplemental matching procedure.

We then identify candidate inventor–employee pairs within matched firms using exact first and last name matches after converting names to a common ASCII format, requiring middle initials to match when both are observed. In Step 1, we classify a candidate as a match when at least one of the inventor's patent filing years falls within a recorded employment spell at the firm, allowing the filing year to precede the recorded start date or follow the recorded end date by one year. In Step 2, when the employment spell has no usable dates, we retain the name match. We drop candidates with recorded employment dates that do not overlap with any of the inventor's filing years at the firm, as well as same-name matches that cannot be resolved using employment timing. This procedure matches 29,864 of the 68,606 inventors in the main green-count regression, or 43.5 percent.

Information coverage. Among the 26,038 matched inventors used to construct the human-capital measures, 23,832 (91.5 percent) have an education record identifying an educational institution, and 20,391 (78.3 percent) have a recorded highest-degree category. The highest recorded degree is a Bachelor's degree for 4,351 inventors (16.7 percent), a Master's degree for 6,164 (23.7 percent), an MBA for 1,402 (5.4 percent), and a doctoral degree for 8,222 (31.6 percent).[23] We observe that 17,707 matched inventors (68.0 percent) attended at least one institution in the United States, and 4,559 (17.5 percent) attended a top-20 institution as defined in Section 5. We observe a seniority category for 24,957 matched inventors (95.8 percent).

Human-capital measures. We use the matched records to construct three predetermined proxies for inventor human capital. *PhD* indicates whether the inventor holds a

[23]The remaining 252 recorded degrees are Associate or High School degrees.

doctoral degree. *Top-20 university* indicates whether the inventor earned any degree from a university in the top 20 of the *U.S. News & World Report* 2025 National Universities ranking for degrees earned in the United States, or in the top 20 of the *Times Higher Education* 2025 World University Rankings for degrees earned elsewhere. *Seniority* is the highest seniority category across the inventor's active positions in a given year.

C Robustness Checks

C.1 Breadth measure construction

Alternative breadth definitions. The baseline measure is the lagged dispersion of an inventor's patenting across non-Y02 CPC main groups, constructed as one minus the Herfindahl–Hirschman index (HHI) of patenting concentration across technology fields, as defined in Equation 2.3. Table A11 examines whether the findings depend on the hierarchical level at which technology fields are defined, the source of the classification information, whether differences among fields are weighted by technological distance, or whether breadth is measured relative to inventors in the same primary field. Columns (2) and (3) construct breadth at the coarser CPC-subclass and finer CPC-subgroup levels, respectively. Column (4) constructs breadth from the technology classes of the patents an inventor cites, rather than from the classes assigned to the patents she produces, following the construction of the patent-originality measure in Hall et al. (2001). Column (5) uses a Rao–Stirling index (Stirling, 2007), which extends the HHI-based measure by weighting each pair of fields by their technological distance. The Rao–Stirling index assigns greater breadth to portfolios spanning more distant fields. We measure distance as one minus the cosine similarity between fields' pre-sample co-classification profiles. Column (6) standardizes the baseline breadth measure within the inventor's primary CPC class and year, where the primary class is the one accounting for the largest share of the inventor's accumulated prior patenting. This addresses the concern that measured breadth may be mechanically higher in technological areas whose CPC classifications are more granular. The quantity effect remains positive and statistically significant at the 5 percent level in all six specifications, ranging from 1.2 to 2.0 percent. The estimated breakthrough effect is negative and statistically significant at the 5 percent level in all specifications except the one based on cited prior art, which is significant at the 10 percent level, ranging from a 5.0 to 7.1 percent decline relative to the outcome mean. Thus, the results are robust to

alternative choices in the construction of breadth.

Depreciation of prior knowledge. The baseline breadth measure weights all of an inventor's prior patents equally. As a robustness check, we rebuild the field stocks as depreciated stocks,

$$K_{ict} = \sum_{p \in P_{it}} \frac{\mathbf{1}\{p \text{ is classified in } c\}}{\text{number of non-Y02 main groups on } p} + (1-\delta)K_{ic,t-1},$$

so that the accumulated knowledge embodied in older patents receives less weight than newer ones. Table A12 reports the undepreciated baseline in column (1) and annual depreciation rates of 10%, 15%, and 20% in columns (2)–(4), respectively. In Panel A, the quantity effect ranges from 2.0 to 2.2%, while in Panel B, the breakthrough-share effect ranges from a 6.6 to 7.1% decline. The results are therefore not driven by the assumed rate at which prior technical knowledge depreciates.

Attribution and team size. The baseline measure of breadth attributes all technology classes on a jointly filed patent to each listed inventor, regardless of team size. This could overstate the breadth of an inventor's individual experience and knowledge, especially when patents involve large teams. Table A13 addresses this concern in three ways. Column (2) weights each patent's contribution by the inverse of team size. Column (3) reconstructs breadth using only patents with at most three inventors, the number of inventors on the median patent.[24] Columns (4)–(7) control for an inventor's mean team size. Column (4) additionally controls for the average breadth of co-inventors on the inventor's patents filed in year t and its interaction with the energy-cost shock. Column (5) analogously controls for their average green share and its interaction with the shock. Column (6) controls for co-inventors' average prior quality, measured by their share of breakthrough patents through year $t-1$, and its interaction with the shock. Column (7) includes all three co-inventor measures and their interactions. Across these specifications, the focal inventor's breadth interaction remains positive for green-patent quantity, with an implied effect ranging from 1.1 to 2.1 percent, and negative for breakthrough share, with an implied decline ranging from 3.6 to 7.1 percent.

[24] Among patents used to construct inventor histories, the median team has three inventors, 67.1% have at most three inventors, and the 90th-percentile team has six.

C.2 Shock construction and identification

Alternative firm energy-cost exposure measures. The baseline shock measures each firm's exposure using the sales-weighted composition of its manufacturing segments in the first year in which the firm appears in the Compustat Segment data. The finance literature has raised concerns about using segment data to measure firms' activities (Villalonga, 2004; Chevalier, 2004). Here we show that the results do not hinge on the specific segment structure used to construct firm-level energy-cost exposure. Column (1) of Table A14 reports the baseline estimates. Column (2) instead constructs exposure using the industry-level energy intensity and fuel composition of the firm's largest manufacturing segment. Column (3) reproduces the baseline estimates on the sample for which the SIC-based exposure used in column (4) is available, and column (4) constructs exposure using the firm's Compustat SIC. Thus, the alternatives in columns (2) and (4) avoid averaging exposure across a firm's reported segments and instead assign the firm the exposure of a single industry: the industry of its largest segment or its firm-level SIC industry. Across these constructions, the quantity effect remains positive and statistically significant at the 5 percent level, ranging from 1.4 to 2.1 percent. The breakthrough effect remains negative and statistically significant at the 5 percent level, ranging from a 6.5 to 7.4 percent decline. Thus, the results are not driven by the use of detailed Compustat segment data to construct firms' energy-cost exposure.

Alternative energy-cost shock definitions. The baseline shock combines firms' predetermined exposure to oil and natural gas with the two fuels' estimated price responses to the oil-supply news shock over a 12-month horizon. Columns (2)–(5) of Table A15 show that the result does not hinge on this construction. Columns (2) and (3) use six- and 18-month aggregation horizons, respectively. Column (4) uses the natural-gas component of the baseline shock alone. Column (5) restricts the sample to inventors who patent with firms whose manufacturing segments account for more than 50% of reported segment sales.[25] The effect of the shock on broader inventors remains positive and statistically significant for green-patent quantity and negative and statistically significant for breakthrough share in every specification. The implied quantity effect ranges from 1.7 to 2.2%, while the breakthrough-share effect ranges from a 5.8 to 7.4% decline.

[25] This restriction retains over 93% of observations in each baseline sample.

C.3 Sample restrictions and alternative specifications

Inventor affiliation. Our research design interprets inventors as workers of the firms with which they patent, as inferred from PatentsView. A concern is that some inventors may instead work at a university or government laboratory while conducting research for a corporation. We therefore use PatentsView's complete disambiguated assignee records to identify inventors who, at some point in their observed careers, patent with a university, research institute, or government entity.[26] In the main green-count regression sample, 14.2% of inventors have at least one such patent.

Table A16 re-estimates the headline quantity and quality specifications after excluding, in column (2), the roughly 11% of inventors with a non-corporate patent between their first and last corporate patents and, in column (3), the 14.2% of inventors with any non-corporate patent over their full observed career. Column (4) reports the estimates for inventors matched to Revelio employment records, whose observed employment at the patenting firm overlaps their patenting activity. Across these samples, the results are qualitatively unchanged. The implied quantity effect ranges from 2.1% to 2.8%, while the breakthrough-share effect ranges from a 7.1% to 9.9% decline. These results, together with the fact that 85.8% of inventors patent only with corporate assignees over their observed careers, support interpreting the inventors in our sample as workers of the firms with which they patent.

Alternative quantity outcomes. The baseline quantity specification estimates Equation 2 using Poisson pseudo-maximum likelihood (PPML), which accommodates the large number of zero patent counts. Table A17 examines three alternative green innovation quantity outcomes in the same specification. Column (1) uses an indicator for filing at least one green patent, column (2) uses the logarithm of one plus the green patent count, and column (3) uses the share of an inventor's annual patents that are green. The estimated effect of a one-standard-deviation shock for inventors who are one standard deviation broader is positive and statistically significant when the outcome is an indicator for filing at least one green patent, the logarithm of one plus the green patent count, or the share of an inventor's annual patents that are green. The implied effects on the filing

[26] PatentsView assigns government entities to assignee types 6–9, but generally classifies universities as companies. We therefore combine government assignee types with a pre-specified search of assignee organization names for university, college, research-institute, national-laboratory, and related terms, including university, institute of technology, research foundation, national laboratory, NIH, NASA, Department of Energy, and Department of Defense.

indicator and green share are 1.1% and 1.5% of their outcome means, respectively. The small 0.2% log-count effect is consistent with documented attenuation when expected counts are low (Cohn et al., 2022). Columns (4) and (5) show no corresponding increase in non-green patenting.[27] Thus, the positive quantity effect is evident on both the extensive and intensive margins of green innovation and does not depend on the exponential conditional-mean specification used by PPML.

Team-composition counts. Appendix Table A4 reestimates the team-composition results in Table 5 using patent counts and PPML. A one-standard-deviation increase in the shock raises green patenting with a new co-inventor by 2.8% more for an inventor whose breadth is one standard deviation higher, compared with 0.3% for green patenting without a new co-inventor. The corresponding estimates remain 2.5% and -0.2% on a common sample.[28] There is no corresponding increase in non-green patenting. Thus, the team-composition result does not depend on using indicators rather than patent counts.

Alternative fixed effects and firm controls. Table A18 replaces firm-by-year fixed effects with firm and NAICS4-by-year fixed effects, the direct effect of the shock, and lagged firm financial controls. With inventor fixed effects, a one-standard-deviation increase in the shock raises green patenting by 1.4% more for an inventor whose breadth is one standard deviation higher and reduces green breakthrough share by 6.6% of its mean, with both effects statistically significant and neither corresponding non-green effect statistically significant.

Alternative breakthrough definitions. The baseline quality outcome is the breakthrough share of an inventor's green patents, defined as the share that falls in the top decile of the technological-importance measure of Kelly et al. (2021) among patents filed in the same NAICS4 industry and year. Table A19 examines whether the result depends on the comparison group used to evaluate a patent's importance or the threshold used to define a breakthrough. Column (1) reports the baseline top-decile cutoff within NAICS4-by-year cells. Columns (2)–(4) change the comparison group, using NAICS3-by-year cells and cells that compare a green patent only with other green patents within NAICS4-

[27] Column (6) is mechanically the mirror image of column (3), because green and non-green patent shares sum to one. It is included for completeness.

[28] PPML uses only inventors who ever file a patent of a given kind, so the two columns otherwise draw on different inventors.

by-year or NAICS3-by-year cells. Columns (5)–(7) retain the baseline NAICS4-by-year comparison group while varying the cutoff from the top 5 percent to the top 20 percent. The estimated effect remains negative and statistically significant in every specification, ranging from a 4.1 to 10.9 percent decline relative to the outcome mean.

Alternative green-patent definitions. The baseline outcome defines a patent as green if it was assigned any CPC Y02 code by the patent office. Table A1 examines six alternative definitions that focus more directly on firms' energy use and production processes. Column (2) restricts the outcome to patents carrying at least one code from four energy-related Y02 subclasses: climate-change mitigation in buildings (Y02B), information and communication technologies that reduce energy use (Y02D), energy generation, transmission, or distribution (Y02E), and production or processing of goods (Y02P). Columns (3) and (4) use broad and narrow sets of technologies selected with assistance from a large language model.[29] Column (5) applies the "clean" electricity definition of Dugoua and Gerarden (2025). A patent is "clean" if it carries at least one code for renewable, nuclear, or enabling electricity technologies, but no "grey" or "dirty" code.[30] In their classification, "grey" technologies include efficiency improvements to combustion processes and electricity generation from biomass or waste, while "dirty" technologies involve fossil-fuel combustion. Column (6) relaxes the definition in column (5) by retaining both "clean" and "grey" technologies, while continuing to exclude "dirty" technologies. Column (7) restricts the baseline definition to patents classified as process innovations, or as both process and product innovations, by Tham et al. (2026), to better capture technologies that may reduce energy costs in production processes. Across the six alternative definitions, the estimated effect of a one-standard-deviation increase in the energy-cost shock for inventors who are one standard deviation broader ranges from 1.2 to 4.4 percent for green innovation quantity. The corresponding effect on the share of green patents that are breakthroughs remains negative, ranging from a 3.7 to 9.8 percent decline. Thus, our results are robust to outcome definitions focused more directly on firms' energy use and production processes.

[29] We provided a large language model with the official titles and descriptions of the Y02 main groups and asked it to classify whether the technologies in each group could reduce a firm's own oil or natural-gas use. The narrow definition retains groups classified as directly related, while the broad definition additionally includes groups classified as potentially related.

[30] For Dugoua and Gerarden's complete clean, grey, and dirty CPC and IPC code lists, see their Online Appendix Tables A.1–A.3.

D Predicted Breakthrough Share from Inventor-Team Characteristics

This appendix examines whether the decline in the breakthrough share of broader inventors' green patents can be explained by the observable characteristics of the inventor matches assigned to those patents. We first estimate how team characteristics predict breakthrough innovation using one observation per multi-inventor patent in our sample. We estimate

$$\text{Breakthrough}_{pft} = X'_{pft}\gamma + \alpha_{ft} + \varepsilon_{pft}, \tag{4}$$

where p denotes a patent assigned to firm f and filed in year t, and α_{ft} denotes firm-by-year fixed effects. The vector X_{pft} contains characteristics of the patent's inventor team, with members' experience and collaboration histories measured using patents filed strictly before year t. These include team size, the share of inventor pairs with a prior joint patent, the number of distinct CPC main groups in team members' prior patent portfolios, and members' mean prior patent count, breadth, green share, breakthrough rate, and career age. We also include the shares of members filing their first patent, new to the firm, and with prior green experience, together with the mean pairwise technological similarity and the share of technologically distant pairs. To capture differences in team structure that may not be reflected in these averages, we additionally include the within-team variances of members' prior patent counts, breadth, green shares, breakthrough rates, and career ages, as well as the variance of pairwise technological similarity.

We estimate Equation (4) using green patents. Using the estimated coefficients, we assign each green patent the breakthrough probability predicted by its team characteristics. Because our quality specification is estimated at the inventor-firm-year level, we average these predicted probabilities across the green patents filed by each inventor in a given firm and year. We then re-estimate Equation (3) using this predicted breakthrough share as the outcome. If the decline in the actual breakthrough share reflects observably weaker inventor matches, the predicted breakthrough share should also decline for broader inventors following the shock. We also estimate the prediction model using green and non-green patents together and apply the resulting coefficients to the green patents.

Appendix Table A6 reports the results. For an inventor whose breadth is one standard deviation higher, a one-standard-deviation increase in the shock reduces the actual breakthrough share by 6.2% of its mean. The breakthrough share predicted by the complete team's characteristics instead rises by 0.6% of its mean and is statistically indistin-

guishable from zero. The estimate is unchanged when the prediction model is estimated using green and non-green patents together. The team characteristics explain 8.1% of the within-firm-year variation in green patents' breakthrough status, but they do not predict the decline in breakthrough share following the shock. The decline therefore does not reflect broader inventors being assigned to teams that are observably less likely to produce breakthrough patents.

References

Adão, R., M. Kolesár, and E. Morales (2019). Shift-share designs: Theory and inference. *Quarterly Journal of Economics 134*(4), 1949–2010.

Arts, S. and L. Fleming (2018). Paradise of novelty—or loss of human capital? exploring new fields and inventive output. *Organization Science 29*(6), 1074–1092.

Binette, O., S. A. York, E. Hickerson, Y. Baek, S. Madhavan, and C. Jones (2022). Estimating the performance of entity resolution algorithms: Lessons learned through PatentsView.org. *arXiv preprint arXiv:2210.01230*.

Bloom, N., M. Schankerman, and J. Van Reenen (2013). Identifying technology spillovers and product market rivalry. *Econometrica 81*(4), 1347–1393.

Caldara, D. and M. Iacoviello (2022). Measuring geopolitical risk. *American Economic Review 112*(4), 1194–1225.

Chevalier, J. (2004). What do we know about cross-subsidization? Evidence from merging firms. *Advances in Economic Analysis & Policy 4*(1), Article 3.

Cohn, J. B., Z. Liu, and M. I. Wardlaw (2022). Count (and count-like) data in finance. *Journal of Financial Economics 146*(2), 529–551.

Dugoua, E. and T. D. Gerarden (2025). Induced innovation, inventors, and the energy transition. *American Economic Review: Insights 7*(1), 90–106.

Fleming, L. (2001). Recombinant uncertainty in technological search. *Management Science 47*(1), 117–132.

Fusillo, F. (2023). Green technologies and diversity in the knowledge search and output phases: Evidence from European patents. *Research Policy 52*(2), 104727.

Goldschlag, N., T. J. Lybbert, and N. J. Zolas (2020). Tracking the technological composition of industries with algorithmic patent concordances. *Economics of Innovation and New Technology 29*(6), 582–602.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2001). The NBER patent citations data file: Lessons, insights and methodological tools. Working Paper 8498, National Bureau of Economic Research.

Hegde, D., K. Herkenhoff, and C. Zhu (2023). Patent publication and innovation. *Journal of Political Economy 131*(7), 1845–1903.

Hill, R., Y. Yin, C. Stein, X. Wang, D. Wang, and B. F. Jones (2025). The pivot penalty in research. *Nature 642*, 999–1005.

Jaffe, A. B. (1986). Technological opportunity and spillovers of R&D: Evidence from firms' patents, profits, and market value. *American Economic Review 76*(5), 984–1001.

Jarociński, M. and P. Karadi (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics 12*(2), 1–43.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review 111*(4), 1092–1125.

Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021). Measuring technological innovation over the long run. *American Economic Review: Insights 3*(3), 303–320.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017). Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics 132*(2), 665–712.

Monath, N., C. Jones, and S. Madhavan (2021). Disambiguating patent inventors, assignees, and their locations in PatentsView. PatentsView technical documentation.

Ramey, V. A. and S. Zubairy (2018). Government spending multipliers in good times and in bad: Evidence from US historical data. *Journal of Political Economy 126*(2), 850–901.

Stirling, A. (2007). A general framework for analysing diversity in science, technology and society. *Journal of the Royal Society Interface 4*(15), 707–719.

Tham, W. W., S. Baslandze, E. Sojli, and L. Liu (2026). Foundational processes and growth. *Journal of Financial Economics 180*.

Verhoeven, D., J. Bakker, and R. Veugelers (2016). Measuring technological novelty with patent-based indicators. *Research Policy 45*(3), 707–723.

Villalonga, B. (2004). Diversification discount or premium? New evidence from the business information tracking series. *Journal of Finance 59*(2), 479–506.

Figures

Figure A1: Breadth and green inventive output

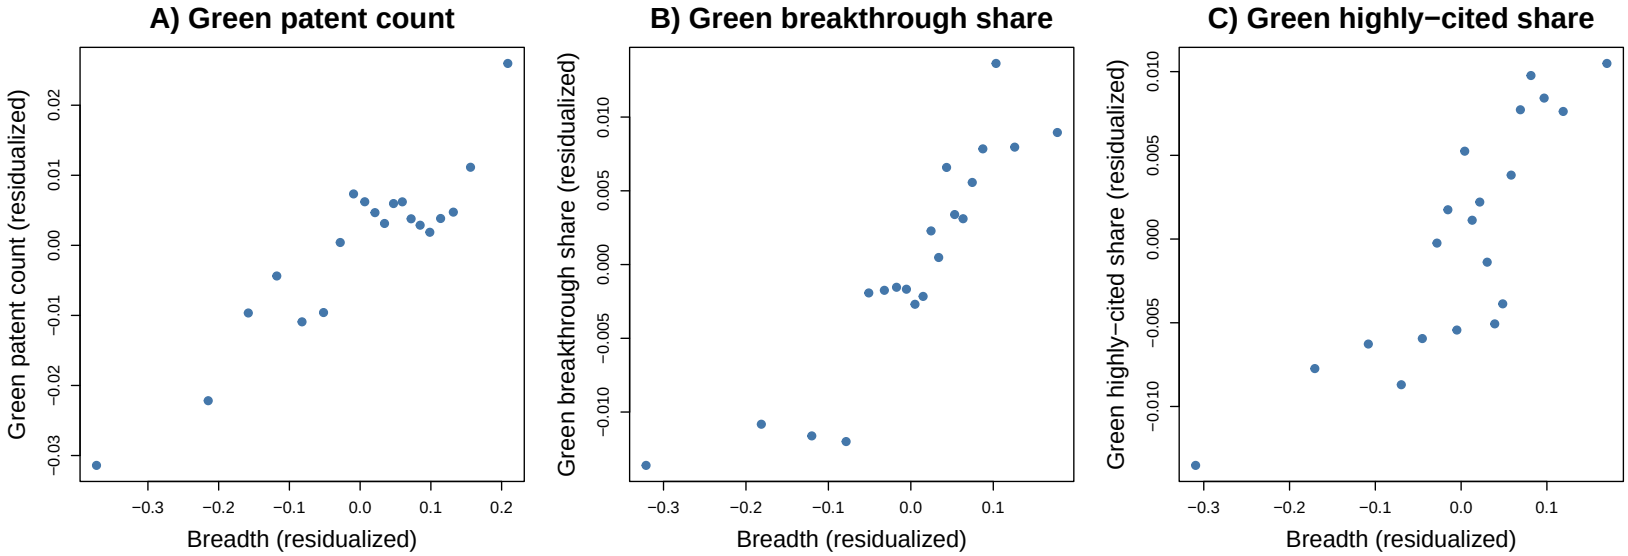

Note: The figure plots residualized annual green patent count (Panel A), breakthrough share of green patents (Panel B), and highly cited share of green patents (Panel C) against residualized breadth. All variables are residualized on firm-by-year fixed effects, within-year deciles of lagged patent stock, career age, and career age squared. Points are means within twenty equal-sized bins of residualized breadth, and lines are fitted linear relationships. The sample covers inventor-years from 1992 to 2021. Panel B is restricted to 1992–2016 due to outcome-data availability. Panels B and C are restricted to inventor-years with at least one green patent.

Figure A2: Predetermined exposure and pre-sample firm characteristics

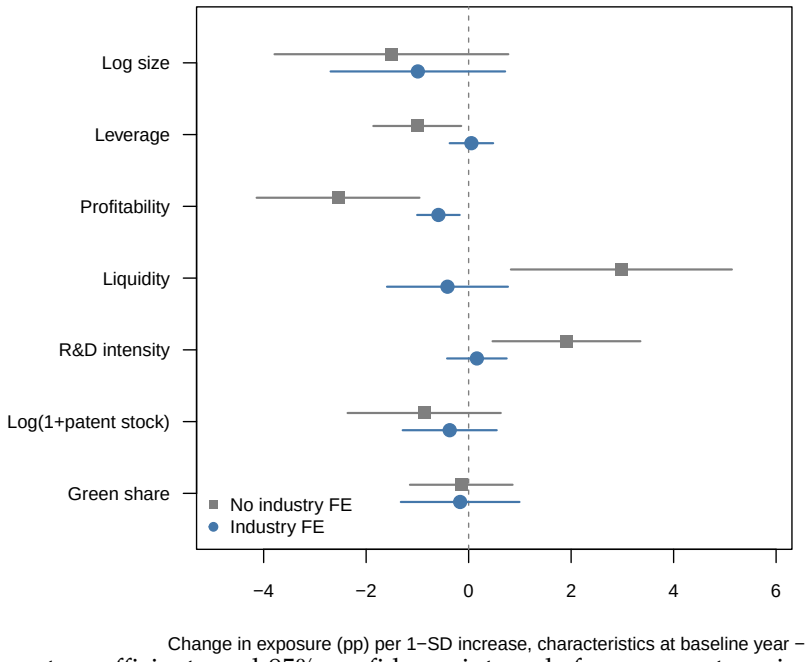

Note: The figure reports coefficients and 95% confidence intervals from separate univariate regressions of firms' predetermined oil-and-natural-gas exposure on each of seven firm characteristics, measured in the year before the firm's exposure baseline year (the year of its first observed segment composition): log assets, leverage, profitability, liquidity, R&D intensity, log of one plus the firm's patent stock, and the firm's green share of patents. All explanatory variables are standardized to have mean zero and unit standard deviation, so each coefficient reports the change in exposure, in percentage points, from a one-standard-deviation increase in the corresponding characteristic. Grey squares report estimates without industry fixed effects; blue circles add three-digit-NAICS fixed effects. The sample is the estimation-sample firms (Section 2.5) with a valid three-digit NAICS code and a non-missing value of the relevant characteristic in the year before their baseline year. Standard errors are clustered by three-digit NAICS industry.

Figure A3: Firm green patenting coefficients from random placebo shifts

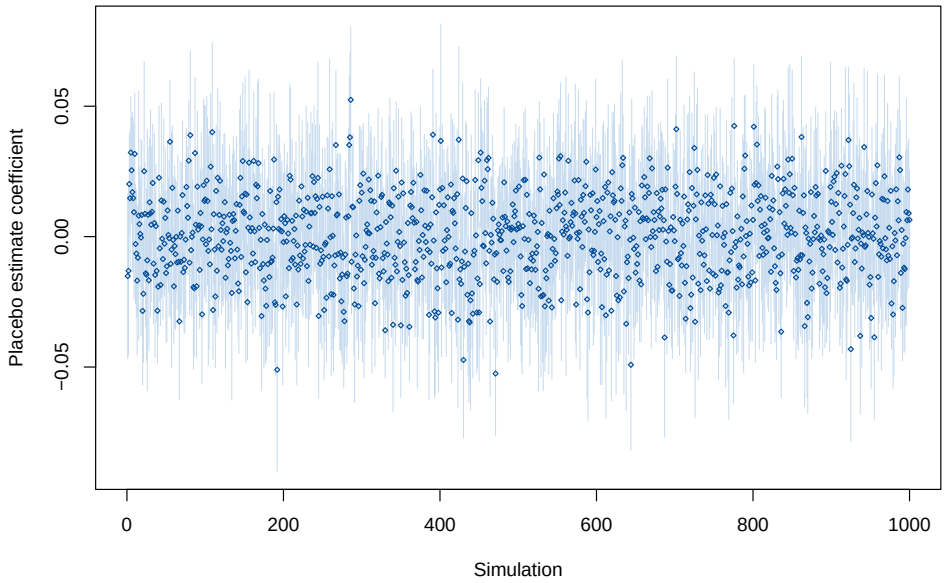

Note: The figure reports coefficients and 95% confidence intervals from 1,000 firm-level PPML regressions that hold firms' predetermined oil and natural-gas exposure fixed and replace the price shifts in the Energy Cost Shock with randomly generated shifts. Each shift is drawn from a normal distribution with mean zero and variance five, following Adão et al. (2019). The outcome is the firm's green patent count. All specifications include firm and NAICS3-by-year fixed effects and lagged financial controls. Standard errors are clustered by firm.

Figure A4: Persistence of firms' segment mix

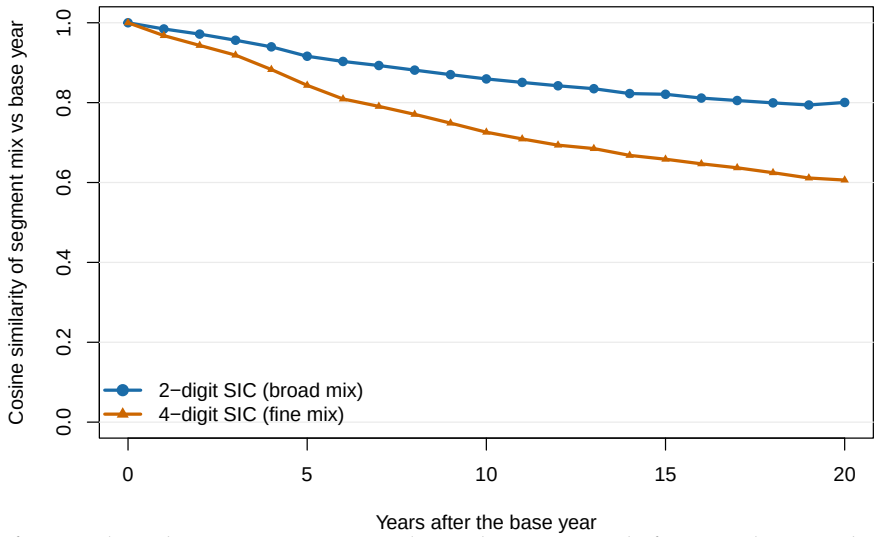

Note: The figure plots the mean cosine similarity between each firm's sales-weighted SIC segment mix in the indicated year and its mix in the firm's first observed year. The blue line (dots) aggregates segment SIC codes to two digits and the orange line (triangles) to four digits.

Figure A5: Persistence of industry energy intensity

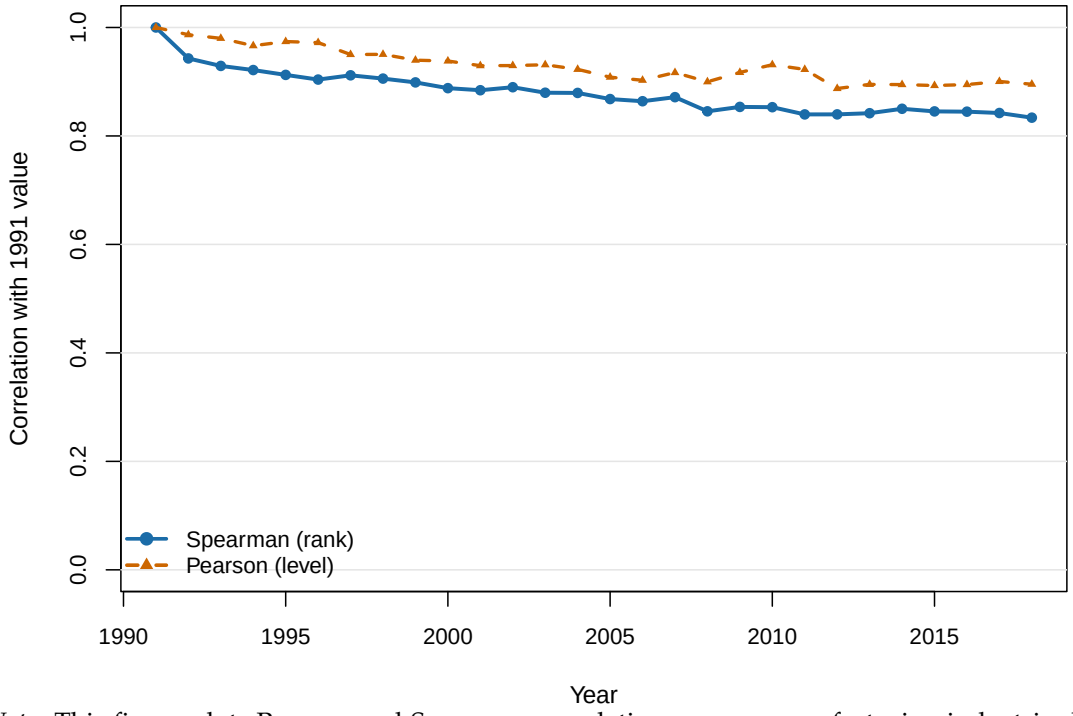

Note: This figure plots Pearson and Spearman correlations across manufacturing industries between energy intensity in each indicated year and in 1991. Energy intensity is energy cost as a share of sales.

Figure A6: Persistence of industry energy input cost shares

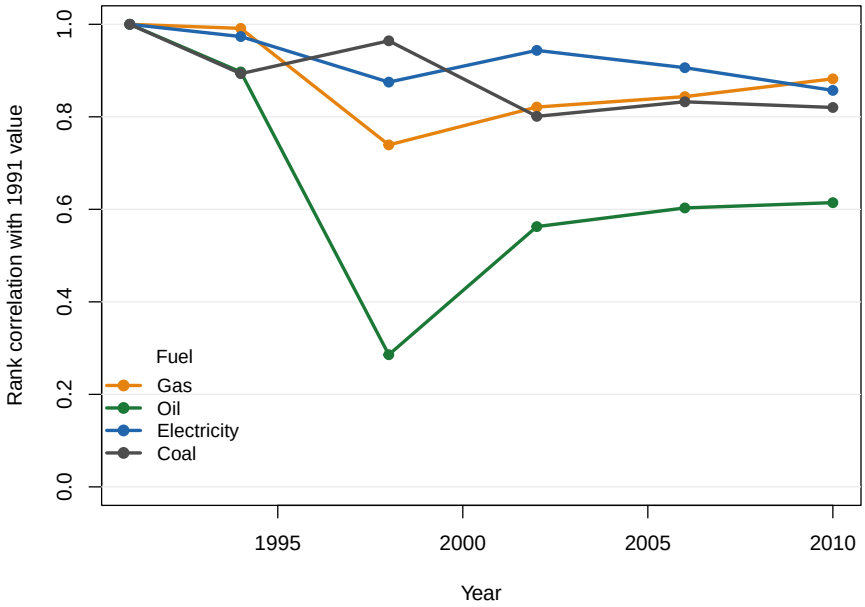

Note: This figure plots, for each fuel and MECS survey year, the Spearman correlation across manufacturing industries between the fuel's share of energy expenditures and its 1991 value. Fuel cost shares combine MECS fuel quantities with contemporaneous relative fuel prices.

Figure A7: Persistence of firm-level exposure by energy input

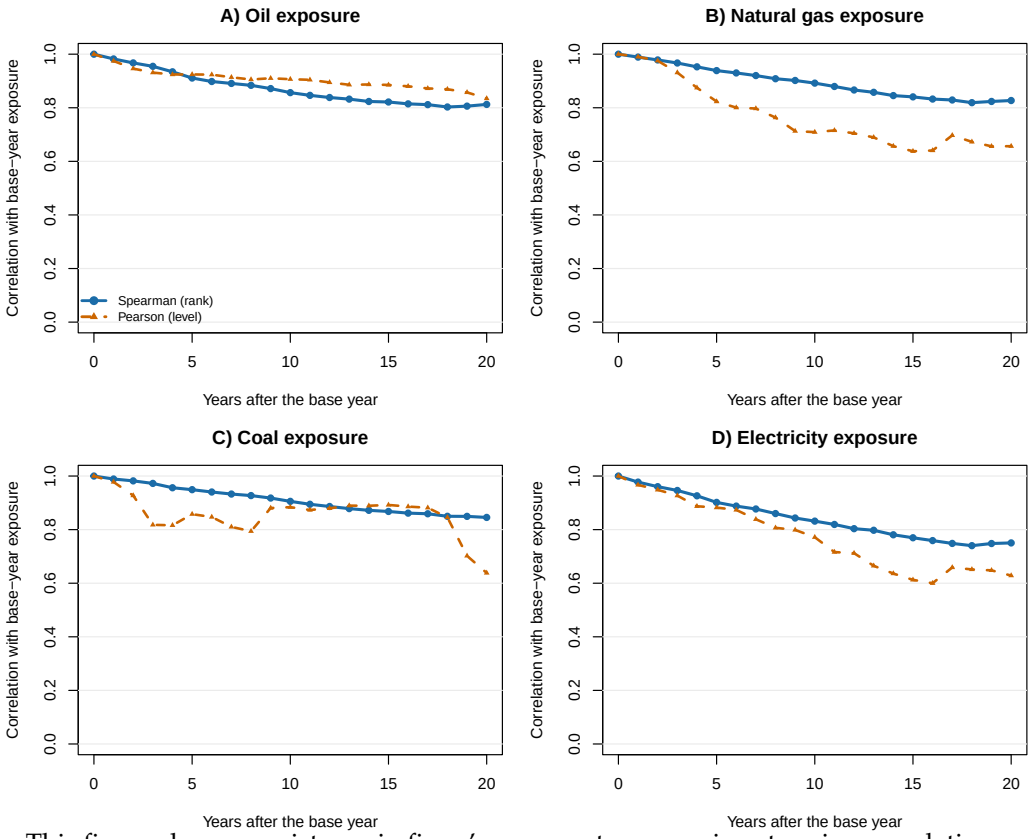

Note: This figure shows persistence in firms' exposure to energy inputs using correlations across firms between exposure to each energy input in the indicated year and in the firm's first observed year. Panel A plots oil exposure, Panel B natural gas exposure, Panel C coal exposure, and Panel D electricity exposure. Exposure to an energy input is its predicted cost as a share of firm sales. Each correlation uses all estimation-sample firms observed at the indicated horizon.

Figure A8: Persistence of the shock across base years

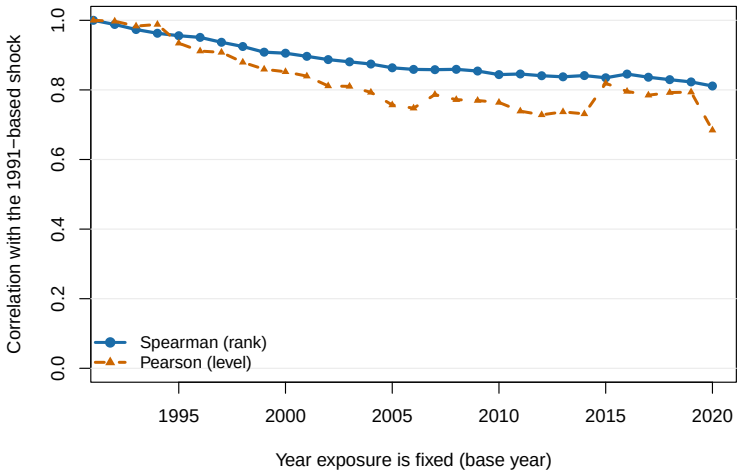

Note: This figure reports the cross-firm correlation between the Energy Cost Shock constructed with firm exposure fixed in each indicated base year and the baseline shock constructed with exposure fixed in 1991.

Figure A9: Distribution of breadth

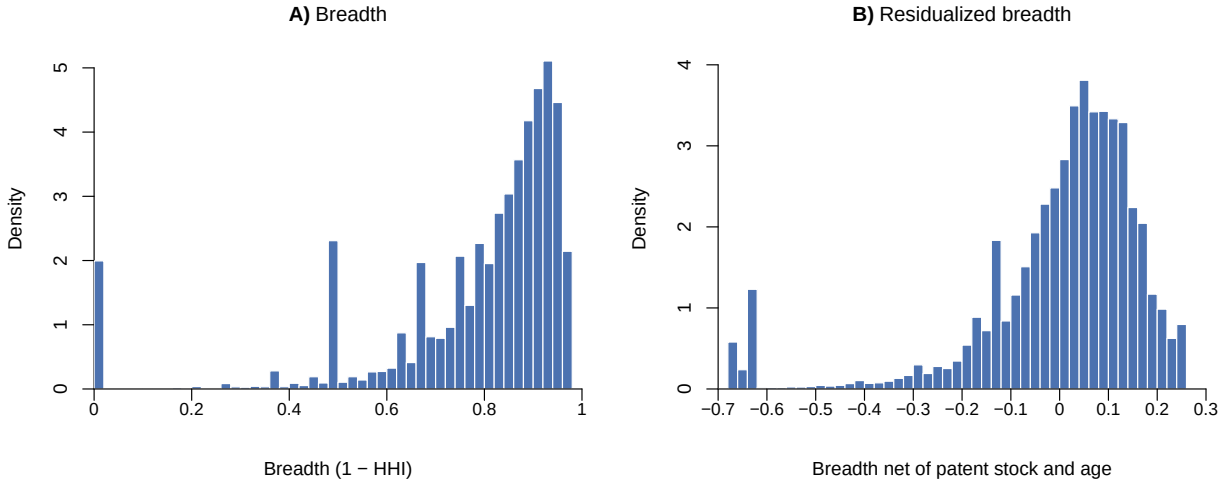

Note: Panel A plots the distribution of breadth, defined as one minus the Herfindahl concentration of an inventor's prior patenting across CPC main groups, excluding Y02 classifications. Panel B plots breadth residualized on lagged patent stock, career age, and career age squared. The sample includes inventor-years in the estimation sample.

Figure A10: Distribution of firm energy cost exposure

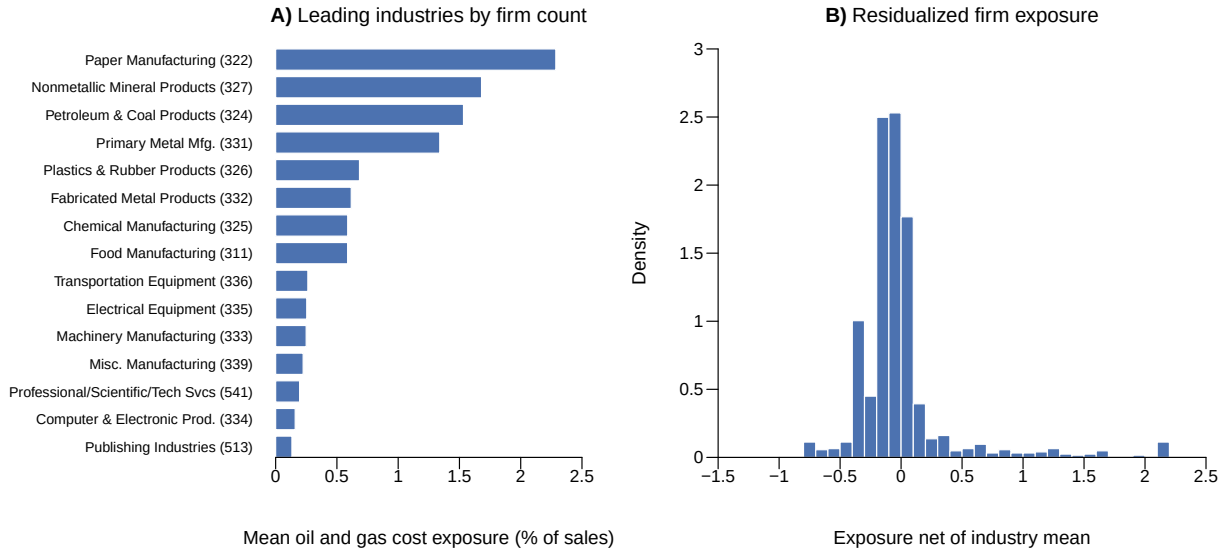

Note: Panel A reports mean predetermined oil and natural gas cost exposure, as a share of sales, for the 15 NAICS3 industries with the most sample firms. These industries account for 94.2% of the sample firms. Panel B plots the residualized distribution of firm exposure after removing its NAICS3 industry mean. Exposure is constructed from the firm's operating-segment composition, energy intensity, and oil and natural gas cost shares in the corresponding industries.

Figure A11: Distribution of the energy-cost shock

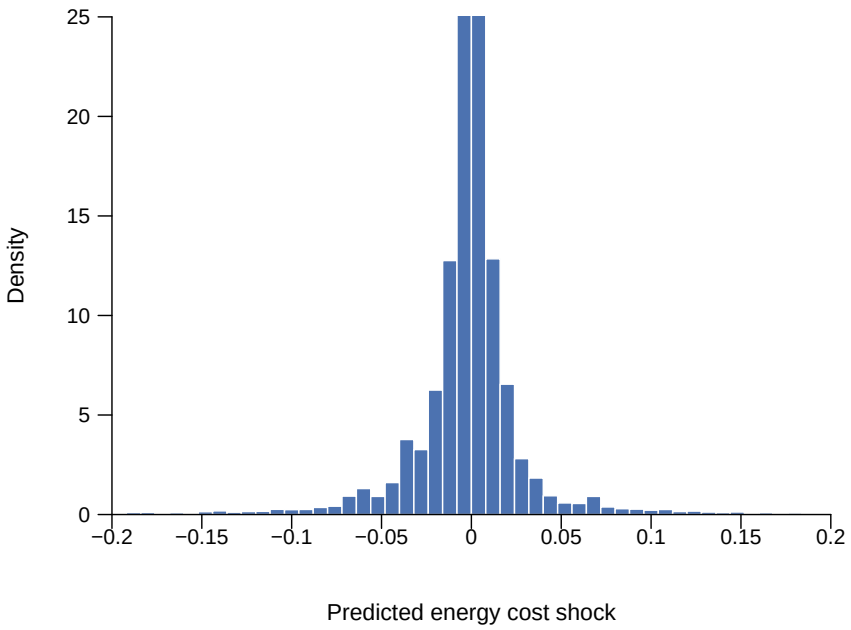

Note: This figure plots the distribution of the firm-year Energy Cost Shock, defined in Section 2.4. The shock sums, across oil and natural gas, each firm's predetermined cost exposure to the energy input multiplied by the predicted change in its price induced by oil supply news shocks. The horizontal axis is restricted to $[-0.2, 0.2]$, retaining 98.70% of firm-years. The sample includes firm-years in the estimation sample.

Tables

Table A1: Robustness checks: alternative green-patent definitions

			API-selected fields		Dugoua and Gerarden (2025)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Definition	Baseline	Energy-related Y02	Broad	Narrow	Clean electricity	Clean or grey electricity	Process green
Panel A)	Dependent variable: Green patent count						
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0784***	0.0913***	0.0823***	0.0598**	0.1517***	0.1697***	0.0484**
	(0.0284)	(0.0258)	(0.0252)	(0.0303)	(0.0584)	(0.0524)	(0.0245)
	[+2.0%]	[+2.3%]	[+2.1%]	[+1.5%]	[+3.9%]	[+4.4%]	[+1.2%]
Breadth_{t-1}	0.6456***	0.5824***	0.6149***	0.6716***	0.3444***	0.3805***	0.6821***
	(0.0674)	(0.0777)	(0.1015)	(0.1178)	(0.1300)	(0.1134)	(0.0790)
Observations	583,525	433,354	274,605	202,219	116,176	138,156	431,140
Panel B)	Dependent variable: Breakthrough share						
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0356***	-0.0395**	-0.0450**	-0.0475*	-0.0319**	-0.0256**	-0.0342**
	(0.0136)	(0.0179)	(0.0205)	(0.0253)	(0.0137)	(0.0109)	(0.0172)
	[−7.1%]	[−7.6%]	[−8.3%]	[−9.8%]	[−4.3%]	[−3.7%]	[−6.6%]
Breadth_{t-1}	0.0145	0.0119	0.0035	0.0126	0.0019	-0.0009	0.0213
	(0.0168)	(0.0183)	(0.0257)	(0.0302)	(0.0326)	(0.0301)	(0.0190)
Outcome mean	0.126	0.131	0.137	0.122	0.187	0.173	0.130
Adjusted R^2	0.415	0.448	0.493	0.579	0.503	0.502	0.425
Observations	65,541	41,767	26,967	16,614	11,252	13,678	38,010
Inventor FE	Y	Y	Y	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y

Note: This table presents estimates of equation (2) using alternative definitions of green innovation. Panel A reports Poisson pseudo-maximum-likelihood (PPML) estimates with the count of the inventor's green patents as the outcome. Panel B reports OLS estimates with the share of the inventor's green patents that are breakthroughs as the outcome, where a breakthrough is a patent in the top decile of the importance measure (Section 2.2). Column (1) uses the baseline definition, any patent assigned a Y02 climate classification. Column (2) restricts green patents to the energy-related Y02B, Y02D, Y02E, and Y02P subclasses. Columns (3) and (4) use broad and narrow sets of Y02 main groups selected with assistance from a large language model. We provided the model with the official titles and descriptions of the Y02 main groups and asked it to classify whether the technologies in each group could reduce a firm's own oil or natural-gas use. The narrow definition retains groups classified as directly related, while the broad definition additionally includes groups classified as potentially related. Columns (5) and (6) use the clean-electricity and clean-or-grey-electricity definitions of Dugoua and Gerarden (2025), respectively. Column (7) restricts the baseline definition to patents whose titles Tham et al. (2026) classify as describing a process, alone or together with a product. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to have mean zero and a standard deviation of one. The bracketed row reports the implied percent change in the outcome for a one-standard-deviation increase in the shock and an inventor one standard deviation above mean breadth, expressed as a percent of the outcome mean in Panel B. The "Outcome mean" row is the unconditional mean of the Panel B outcome. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A2: Timing of the firm-level response to energy-cost shocks

	Placebo	Contemporaneous	Baseline	Post-shock		
	(1)	(2)	(3)	(4)	(5)	(6)
	$t-2$	$t-1$	t	$t+1$	$t+2$	$t+3$
Shock_{t-1}	0.0095	0.0369**	0.0375**	0.0246	0.0236	0.0144
	(0.0146)	(0.0147)	(0.0187)	(0.0178)	(0.0167)	(0.0157)
Observations	20,993	20,993	20,993	19,233	17,618	16,147
Firm FE	Y	Y	Y	Y	Y	Y
Industry×year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y

Note: This table presents PPML estimates of the firm's green patent count, tracing the response across outcome horizons $t-2$ to $t+3$ relative to the baseline shock year, $t-1$, on the exact estimation sample of the firm-level results table. Shock is the lagged energy-cost shock, as defined in Section 2.4, standardized to zero mean and unit standard deviation. In columns (1) to (3), the event shock occurs, respectively, one year after, in the same year as, and one year before the outcome, while financial controls are dated one year before each outcome. In columns (4) to (6), the event shock and financial controls remain fixed at $t-1$, and the specifications additionally control for every shock realized between $t-1$ and the outcome. The financial controls are log size, liquidity, leverage, profitability, and R&D intensity. All columns include firm and industry-by-year fixed effects. Standard errors (in parentheses) are clustered at the firm level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A3: Alternative aggregate shocks

	Dependent variable: Green patent count				
	Baseline Energy Shock	Geopolitical Risk Shock	Monetary Policy Shock	Central Bank Information Shock	Government Spending News Shock
	(1)	(2)	(3)	(4)	(5)
Shock_{t-1}	0.0377**	-0.0158	-0.0014	-0.0344	0.0232
	(0.0188)	(0.0137)	(0.0153)	(0.0267)	(0.0195)
	[+3.8%]	[−1.6%]	[−0.1%]	[−3.4%]	[+2.3%]
Firm FE	Y	Y	Y	Y	Y
Industry×year FE	Y	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y	Y
Observations	21,012	21,012	21,012	21,012	17,348

Note: This table presents Poisson pseudo-maximum-likelihood (PPML) estimates of the firm's green patent count on five alternative shocks, each interacted with the share of the firm's 1991 sales spent on natural gas and oil, lagged one year and standardized to have mean zero and a standard deviation of one. Column (1) uses the baseline energy-cost shock, built from the oil supply news shocks of Känzig (2021). Column (2) is the annual change in the log mean Geopolitical Risk index of Caldara and Iacoviello (2022). Columns (3) and (4) are the annual sums of the pure monetary policy shock and the central bank information shock of Jarociński and Karadi (2020), respectively. Column (5) is the annual defense-news shock of Ramey and Zubairy (2018), scaled by nominal GDP; the series ends in 2015, so column (5) excludes firm-years after 2016. All columns include firm and industry (NAICS3) by year fixed effects and lagged firm controls: log size, liquidity, leverage, profitability, and R&D intensity. Standard errors (in parentheses) are clustered by firm. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A4: Robustness checks: inventor-team composition

	Green patents		Other patents	
	(1)	(2)	(3)	(4)
Co-inventors	At least one new	No new	At least one new	No new
Panel A)	Baseline PPML estimation samples			
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.1104***	0.0104	-0.0026	0.0099
	(0.0363)	(0.0533)	(0.0262)	(0.0354)
	[+2.8%]	[+0.3%]	[−0.1%]	[+0.2%]
Breadth_{t-1}	0.6663***	1.7262***	0.2526***	0.7216***
	(0.0737)	(0.1071)	(0.0437)	(0.0951)
Outcome mean	0.209	0.178	0.819	0.836
Observations	479,501	341,430	557,103	487,334
Panel B)	Common starting sample within patent type			
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0989**	-0.0062	-0.0086	0.0039
	(0.0464)	(0.0742)	(0.0322)	(0.0389)
	[+2.5%]	[−0.2%]	[−0.2%]	[+0.1%]
Breadth_{t-1}	0.8382***	2.2875***	0.3425***	0.7673***
	(0.0989)	(0.1473)	(0.0412)	(0.0995)
Outcome mean	0.328	0.322	0.963	1.029
Observations	241,936	242,734	461,189	461,319
Inventor FE	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y

Note: This table presents PPML estimates of equation (2) using the inventor panel, with patent counts split by team composition as the outcome. Columns (1) and (2) split the inventor's green patents by team composition, and columns (3) and (4) do the same for non-green patents. A co-inventor is new if the inventor and the co-inventor had not appeared together on any earlier patent, at any assignee. Columns (1) and (3) use the count of patents with two to fifteen inventors on which at least one co-inventor is new. Columns (2) and (4) use the count of patents without a new co-inventor, which includes team patents on which every co-inventor had previously patented with the inventor and solo patents. Panel A estimates each outcome on its available PPML sample. The number of observations differs across columns because PPML drops inventors and firm-years whose outcome is zero in every sampled year. Panel B begins with the inventor-years that enter both estimation samples for a given patent type. PPML may subsequently drop a small number of additional separated observations, so the final sample sizes can differ slightly within each pair. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4, standardized to be zero-mean and have a standard deviation of 1. The bracketed rows report the implied percentage change in the outcome per one-standard-deviation increase in the shock for an inventor one standard deviation above mean breadth. Table 5 reports the corresponding OLS estimates using indicators for the same outcomes. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A5: Energy-cost shocks and green innovation quality: additional measures

	Breakthrough innovation				Scientific importance		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Outcome	Share	Count	Log count	Breakthrough dummy	Log importance	Highly cited share	Log economic value
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0356***	-0.3147***	-0.0261**	-0.0386***	-0.0101**	-0.0139*	-0.0040
	(0.0136)	(0.1168)	(0.0109)	(0.0144)	(0.0044)	(0.0082)	(0.0100)
	[−7.1%]	[−7.6%]	[−0.7%]	[−6.5%]	[−0.3%]	[−2.0%]	[−0.1%]
Breadth	0.0145	0.4593***	0.0622***	0.0422**	0.0042	0.0167	-0.0226
	(0.0168)	(0.1020)	(0.0163)	(0.0193)	(0.0071)	(0.0175)	(0.0148)
Inventor FE	Y	Y	Y	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y
Outcome mean	0.126	0.197	0.119	0.149	0.219	0.177	2.083
Observations	65,541	19,371	65,541	65,541	65,541	126,149	130,656
Adjusted R^2	0.415	-	0.386	0.380	0.829	0.352	0.889

Note: This table presents estimates of equation (3) using alternative measures of the quality of green innovation. Column (1) uses the share of the inventor's green patents in the top decile of the importance measure as the outcome, estimated by OLS. Column (2) uses the count of green breakthrough patents, estimated by Poisson pseudo-maximum-likelihood (PPML). Columns (3) and (4) use the log of one plus that count and an indicator for filing any green breakthrough patent, respectively, estimated by OLS. Column (5) uses the mean log forward-to-backward similarity ratio, the continuous measure of importance from Kelly et al. (2021), estimated by OLS. Column (6) uses the share of the inventor's green patents in the top decile of five-year forward citations, estimated by OLS. Column (7) uses the mean log economic value of the inventor's green patents, from Kogan et al. (2017), estimated by OLS. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. The bracketed row reports the implied change in the outcome per one standard deviation increase in the shock, for an inventor one standard deviation above mean breadth: as a percentage of the outcome mean in columns (1), (4), and (6), as an approximate percent change in the log outcomes of columns (3), (5), and (7), and as a percent change in the count in column (2). The "Outcome mean" row is the unconditional mean of each outcome. Columns (1) to (5) include inventor-years with a scored green patent through 2016; column (6) includes inventor-years with any green patent; and column (7) includes inventor-years for which economic value is defined. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A6: Alternative explanations: predicted breakthrough share from inventor matching

	Green breakthrough share		
	(1)	(2)	(3)
Dependent variable	Actual	Predicted using green patents	Predicted using all patent types
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0300**	0.0031	0.0030
	(0.0143)	(0.0021)	(0.0022)
	[−6.2%]	[+0.6%]	[+0.6%]
Breadth_{t-1}	0.0020	0.0217***	0.0194***
	(0.0179)	(0.0057)	(0.0059)
Observations	59,389	59,389	59,389
Adjusted R^2	0.436	0.713	0.711
Outcome mean	0.122	0.127	0.128
Inventor FE	Y	Y	Y
Firm×year FE	Y	Y	Y
Controls	Y	Y	Y

Note: This table compares the actual breakthrough share of green patents with the share predicted from the predetermined characteristics of their inventor teams. Appendix Section D describes the prediction exercise and the team characteristics. Column (1) uses the actual breakthrough share. Column (2) predicts each green patent using coefficients estimated from green patents. Column (3) predicts the same green patents using coefficients estimated from all green and non-green patents. Each first stage includes firm-by-year fixed effects, which are excluded from the predicted breakthrough share. All three second-stage regressions use the same inventor-firm-year observations. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). The bracketed row reports the implied change as a percentage of the outcome mean per one-standard-deviation increase in the shock for an inventor whose breadth is one standard deviation higher. Treating the first-stage predictions as fixed, the difference between columns (1) and (2) has $p = 0.028$, and the difference between columns (1) and (3) has $p = 0.027$. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A7: Alternative explanations: inventor-team composition and the scientific value of green innovation

	Green breakthrough share		Other breakthrough share	
	(1)	(2)	(3)	(4)
New co-inventor	Yes	No	Yes	No
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0371***	-0.0176	-0.0046	-0.0033
	(0.0136)	(0.0208)	(0.0048)	(0.0066)
	[−6.4%]	[−4.2%]	[−1.0%]	[−0.9%]
Breadth_{t-1}	0.0312	0.0242	0.0209***	0.0498***
	(0.0248)	(0.0328)	(0.0073)	(0.0141)
Observations	29,364	27,060	142,232	115,986
Adjusted R^2	0.326	0.498	0.204	0.265
Outcome mean	0.146	0.107	0.115	0.092
Inventor FE	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y

Note: This table presents OLS estimates of equation (3) using the inventor panel, the breakthrough-share analogue of the inventor-team composition results in Table 5. The outcome is the share of the inventor's green patents that are breakthroughs in columns (1) and (2), and the share of other patents that are breakthroughs in columns (3) and (4); a breakthrough is a patent in the top decile of the importance measure (Section 2.2). A co-inventor is new if the inventor and the co-inventor had not appeared together on any earlier patent, at any assignee. Columns (1) and (3) compute the breakthrough share over patents with two to fifteen inventors on which at least one co-inventor is new. Columns (2) and (4) use the inventor's remaining patents, including patents with only established co-inventors and solo patents. Each column uses inventor-years with an observed breakthrough share for the corresponding patent category. Appendix Table A4 reports the corresponding patent-count estimates. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4, standardized to have mean zero and a standard deviation of one. The bracketed row reports the change in the outcome as a percent of its mean per one-standard-deviation increase in the shock, for an inventor one standard deviation above mean breadth. The Outcome mean row is the unconditional mean of each outcome. A stacked test allowing the fixed effects and controls to differ across outcomes yields $p = 0.48$ for equality of the interaction coefficients in columns (1) and (2), and $p = 0.83$ for equality in columns (3) and (4). All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A8: CPC hierarchy: illustrative technology classifications

Technology family	Illustrative main groups within the subclass	Illustrative subgroups within one main group
Iron smelting (C21B)	C21B 5/00 Pig iron in blast furnaces; C21B 7/00 Blast furnaces	C21B 5/003 Coal injection; C21B 5/006 Automatic process control
Air conditioning (F24F)	F24F 1/00 Room units; F24F 3/00 Central-station systems	F24F 1/02 Self-contained room units; F24F 1/022 Compressor cycles
Internal-combustion engines (F02B)	F02B 5/00 Positive-ignition engines; F02B 7/00 Compression-ignition engines	F02B 5/02 Methods of operating
Photographic apparatus (G03B)	G03B 19/00 Cameras; G03B 21/00 Projectors	G03B 19/02 Still-picture cameras; G03B 19/18 Motion-picture cameras

Note: The second column gives distinct main groups within each subclass. The third gives subgroups within the main group named first in the second column. CPC titles are abbreviated for readability.

Table A9: Proximity of green and non-green patents to green knowledge

	Green			Non-green			
Industry	N	Mean	SD	N	Mean	SD	t
Basic chemicals	26,851	0.584	0.263	237,264	0.491	0.273	55.023***
Other chemicals	9,012	0.656	0.266	208,565	0.493	0.277	56.851***
Oil refining	3,892	0.616	0.288	17,465	0.474	0.262	28.212***
Electricity	49,928	0.735	0.183	154,882	0.424	0.267	292.274***
Metals	6,216	0.666	0.272	86,669	0.488	0.276	49.844***
Fabricated metals	36,870	0.664	0.268	242,091	0.475	0.263	126.134***
Glass	2,854	0.589	0.284	17,651	0.486	0.273	18.227***
Cement	7,472	0.723	0.242	68,077	0.476	0.270	82.861***
Motor vehicles	10,376	0.878	0.180	172,313	0.477	0.268	213.752***
Other transport	17,109	0.737	0.236	131,384	0.469	0.268	137.110***
Total	**574,279**	**0.714**	**0.243**	**7,304,054**	**0.483**	**0.276**	**684.363****

Note: This table compares the proximity of green and non-green patents to the green knowledge domain within ten energy-intensive manufacturing industries defined over four-digit NAICS: Basic chemicals (3251), Other chemicals (3252, 3253, 3255, 3256, 3259), Oil refining (3241), Electricity (electrical power equipment, 3353), Metals (3311–3315), Fabricated metals (3321–3327, 3329), Glass (3272), Cement (3273–3274), Motor vehicles (3361–3363), and Other transport (3364–3366, 3369). Proximity is the patent-level within-industry-year percentile of green-knowledge adjacency (`adj_all_pct_inclg`); eligible patents have a defined proximity percentile (industry-year cell size > 1) and an assigned four-digit NAICS. The green reference vector for year t uses green patents filed in $t-5$ to $t-1$, so a focal green patent never enters the reference that scores it. The Total row pools all eligible patents across every industry (not only the ten shown). t is a two-sided Welch difference-in-means statistic (green minus non-green). $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A10: Proximity predicts receiving a citation from a green patent within five years

	(1) All patents	(2) Non-green patents	(3) Green patents
Proximity	0.1556***	0.0891***	0.1368***
	(0.0254)	(0.0135)	(0.0272)
Observations	6,173,939	5,750,725	423,008
Outcome mean	0.0867	0.0617	0.4263
SD(proximity)	0.284	0.279	0.239
Effect of 1 SD increase (pp)	4.41	2.49	3.27
Effect / outcome mean	50.9%	40.3%	7.7%
NAICS4×year FE	Y	Y	Y

Note: This table presents patent-level linear probability estimates with the outcome variable set to one if the patent is cited by at least one green patent filed in the subsequent five years. Proximity is the patent's within-NAICS4-by-filing-year percentile of adjacency to green knowledge. Column (1) includes all patents, column (2) non-green patents, and column (3) green patents. All columns include NAICS4-by-filing-year fixed effects. Standard errors are clustered by NAICS4 industry. The effect rows report the estimated effect of a one-standard-deviation increase in proximity in percentage points and relative to the outcome mean.

Table A11: Robustness checks: alternative breadth definitions

		Field definition			Distance and fields	
	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	Coarser fields (subclass)	Finer fields (subgroup)	Cited prior art	Distance-weighted (Rao–Stirling)	Relative to field peers
Panel A)	Dependent variable: Green patent count					
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0784***	0.0423**	0.0904***	0.1374***	0.0683***	0.0147**
	(0.0284)	(0.0199)	(0.0313)	(0.0475)	(0.0245)	(0.0069)
	[+2.0%]	[+1.2%]	[+1.9%]	[+1.7%]	[+1.6%]	[+1.5%]
Breadth_{t-1}	0.6456***	0.4368***	0.7396***	0.5704***	0.6159***	0.0925***
	(0.0674)	(0.0686)	(0.0810)	(0.1205)	(0.0894)	(0.0140)
Observations	583,525	583,525	583,525	579,077	583,525	583,525
Panel B)	Dependent variable: Breakthrough share					
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0356***	-0.0232**	-0.0386**	-0.0675*	-0.0274**	-0.0067**
	(0.0136)	(0.0108)	(0.0157)	(0.0380)	(0.0130)	(0.0030)
	[−7.1%]	[−5.2%]	[−6.3%]	[−6.6%]	[−5.0%]	[−5.3%]
Breadth_{t-1}	0.0145	0.0059	0.0224	0.0397*	0.0048	0.0033
	(0.0168)	(0.0121)	(0.0158)	(0.0217)	(0.0145)	(0.0035)
Outcome mean	0.126	0.126	0.126	0.126	0.126	0.126
Observations	65,541	65,541	65,541	65,035	65,541	65,541
Inventor FE	Y	Y	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y	Y	Y

Note: This table presents estimates of equation (2) using alternative definitions of breadth. Panel A reports Poisson pseudo-maximum-likelihood (PPML) estimates with the count of the inventor's green patents as the outcome. Panel B reports OLS estimates with the share of the inventor's green patents that are breakthroughs as the outcome, where a breakthrough is a patent in the top decile of the importance measure (Section 2.2). Column (1) uses the baseline measure, one minus the HHI of the inventor's prior patenting across non-Y02 CPC main groups. Columns (2) and (3) aggregate the inventor's prior patenting to the coarser CPC-subclass level and the finer CPC-subgroup level, respectively. Column (4) constructs breadth from the CPC main groups of the patents cited by the inventor's prior patents. Column (5) uses a Rao–Stirling index that weights each pair of CPC main groups by their technological distance, measured as one minus the cosine similarity between their pre-sample co-classification profiles. Column (6) standardizes the baseline breadth measure within the inventor's primary CPC class and year, where the primary class accounts for the largest share of the inventor's accumulated prior patenting. Every measure excludes Y02 codes. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to have mean zero and a standard deviation of one. The bracketed row reports the implied percent change in the outcome for a one-standard-deviation increase in the shock and a one-standard-deviation increase in the corresponding breadth measure, expressed as a percent of the outcome mean in Panel B. The "Outcome mean" row is the unconditional mean of the Panel B outcome. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A12: Robustness checks: alternative depreciation rates for breadth

	Annual depreciation rate			
	(1) Baseline (0%)	(2) 10%	(3) 15%	(4) 20%
Panel A)	Dependent variable: Green patent count			
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0784***	0.0788***	0.0820***	0.0856***
	(0.0284)	(0.0290)	(0.0285)	(0.0279)
	[+2.0%]	[+2.0%]	[+2.1%]	[+2.2%]
Breadth_{t-1}	0.6456***	0.5632***	0.5130***	0.4722***
	(0.0674)	(0.0657)	(0.0644)	(0.0632)
Observations	583,525	583,525	583,525	583,525
Panel B)	Dependent variable: Breakthrough share			
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0356***	-0.0347***	-0.0336**	-0.0323**
	(0.0136)	(0.0133)	(0.0132)	(0.0130)
	[−7.1%]	[−7.0%]	[−6.8%]	[−6.6%]
Breadth_{t-1}	0.0145	0.0121	0.0122	0.0130
	(0.0168)	(0.0156)	(0.0148)	(0.0140)
Outcome mean	0.126	0.126	0.126	0.126
Adjusted R^2	0.415	0.415	0.415	0.415
Observations	65,541	65,541	65,541	65,541
Inventor FE	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y

Note: This table presents estimates of equation (2) using the inventor panel, rebuilding the non-Y02 CPC main-group breadth stock at alternative depreciation rates. Panel A reports Poisson pseudo-maximum-likelihood (PPML) estimates with the count of the inventor's green patents as the outcome. Panel B reports OLS estimates with the share of green patents in the top decile of the importance measure as the outcome (Section 2.2), which we define as a breakthrough patent. Column (1) is the baseline measure, which carries every past field at full weight, $\delta = 0$ (Section 2.3); columns (2) to (4) rebuild the stock as a perpetual inventory with depreciation rates $\delta = 0.10$, $\delta = 0.15$, and $\delta = 0.20$. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. The bracketed row reports the implied percent change in the outcome per one standard deviation increase in the shock, for an inventor one standard deviation above mean breadth, expressed as a percent of the outcome mean in Panel B. The "Outcome mean" row is the unconditional mean of the Panel B outcome. All columns control for the share of the inventor's prior patents that are green and the log of one plus the inventor's lagged patent stock, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A13: Robustness checks: multi-inventor patents

				Team controls			
	(1) Baseline	(2) Fractional weighting	(3) Small teams (≤ 3)	(4) Breadth	(5) Green share	(6) Prior quality	(7) All
Panel A)	Dependent variable: Green patent count						
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0784***	0.0783***	0.0407***	0.0672***	0.0835***	0.0798***	0.0676***
	(0.0284)	(0.0243)	(0.0154)	(0.0172)	(0.0213)	(0.0234)	(0.0151)
	[+2.0%]	[+2.0%]	[+1.1%]	[+1.7%]	[+2.1%]	[+2.0%]	[+1.7%]
Team breadth × Shock	–	–	–	0.0302	–	–	0.0404
				(0.0465)			(0.0431)
Team green share × Shock	–	–	–	–	-0.0071	–	-0.0095
					(0.0152)		(0.0151)
Team prior quality × Shock	–	–	–	–	–	0.0728*	0.0726*
						(0.0438)	(0.0415)
Observations	583,525	583,525	496,174	345,870	345,870	344,620	344,620
Panel B)	Dependent variable: Breakthrough share						
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0356***	-0.0304**	-0.0281**	-0.0179**	-0.0339**	-0.0330**	-0.0193***
	(0.0136)	(0.0127)	(0.0120)	(0.0077)	(0.0136)	(0.0142)	(0.0068)
	[−7.1%]	[−6.2%]	[−6.1%]	[−3.6%]	[−6.8%]	[−6.6%]	[−3.8%]
Team breadth × Shock	–	–	–	-0.0440*	–	–	-0.0447*
				(0.0251)			(0.0240)
Team green share × Shock	–	–	–	–	-0.0290***	–	-0.0284***
					(0.0063)		(0.0066)
Team prior quality × Shock	–	–	–	–	–	-0.0125	-0.0116
						(0.0257)	(0.0267)
Outcome mean	0.124	0.124	0.123	0.123	0.123	0.123	0.123
Adjusted R^2	0.415	0.415	0.398	0.426	0.426	0.426	0.427
Observations	65,541	65,541	55,977	62,448	62,448	62,448	62,448
Team size control	N	N	N	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y	Y	Y	Y

Note: This table presents estimates of equation (2) using the inventor panel. Panel A reports Poisson pseudo-maximum-likelihood (PPML) estimates with the count of the inventor's green patents as the outcome. Panel B reports OLS estimates with the share of the inventor's green patents that are breakthroughs as the outcome, where a breakthrough is a patent in the top decile of the importance measure (Section 2.2). Column (1) uses the baseline specification. Column (2) reconstructs the breadth measure after weighting each patent by the inverse of its team size. Column (3) reconstructs the breadth measure from patents with at most three inventors. Columns (4) through (7) also control for the inventor-year average team size. Column (4) controls for the average breadth of the co-inventors on patents filed by the inventor in year t, and its interaction with the shock. Column (5) controls for the average green share of these co-inventors, and its interaction with the shock. Column (6) controls for the average prior breakthrough share of these co-inventors through $t-1$ and its interaction with the shock. Column (7) controls for all three co-inventor measures and their interactions with the shock. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. The bracketed row reports the implied percent change in the outcome per one standard deviation increase in the shock, for an inventor one standard deviation above mean breadth, expressed as a percent of the outcome mean in Panel B. The "Outcome mean" row is the unconditional mean of the Panel B outcome. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A14: Robustness checks: alternative firm energy-cost exposure measures

	Baseline	Largest segment	Baseline, SIC sample	Firm SIC exposure
	(1)	(2)	(3)	(4)
Panel A)	Dependent variable: Green patent count			
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0784***	0.0821***	0.0790***	0.0570**
	(0.0284)	(0.0247)	(0.0286)	(0.0289)
	[+2.0%]	[+2.1%]	[+2.0%]	[+1.4%]
Breadth_{t-1}	0.6456***	0.6458***	0.6531***	0.6534***
	(0.0674)	(0.0673)	(0.0697)	(0.0695)
Observations	583,525	583,525	549,708	549,708
Panel B)	Dependent variable: Breakthrough share			
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0356***	-0.0371**	-0.0356***	-0.0327**
	(0.0136)	(0.0172)	(0.0137)	(0.0165)
	[−7.1%]	[−7.4%]	[−7.1%]	[−6.5%]
Breadth_{t-1}	0.0145	0.0143	0.0139	0.0150
	(0.0168)	(0.0168)	(0.0175)	(0.0175)
Outcome mean	0.126	0.126	0.126	0.126
Observations	65,541	65,541	62,413	62,413
Inventor FE	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y

Note: This table presents estimates of equation (2) using alternative constructions of firms' energy-cost exposure. Panel A reports Poisson pseudo-maximum-likelihood (PPML) estimates with the count of the inventor's green patents as the outcome. Panel B reports OLS estimates with the share of the inventor's green patents that are breakthroughs as the outcome, where a breakthrough is a patent in the top decile of the importance measure (Section 2.2). Column (1) uses the baseline exposure, which combines the industry-level energy intensity and fuel composition of the firm's manufacturing segments using their shares of segment sales in the firm's first observed segment year. Column (2) instead assigns the firm the industry-level energy intensity and fuel composition of its largest manufacturing segment in that year. Column (3) reproduces the baseline specification on the sample for which the SIC-based exposure in column (4) is available. Column (4) assigns the firm the industry-level energy intensity and fuel composition associated with its Compustat SIC measured in or closest to its baseline segment year. Each exposure measure is fixed before the outcome year and is interacted with the predicted oil and natural-gas price changes described in Section 2.4. Shock is standardized to have mean zero and a standard deviation of one. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). The bracketed row reports the implied percent change in the outcome for a one-standard-deviation increase in the shock and an inventor one standard deviation above mean breadth, expressed as a percent of the outcome mean in Panel B. The "Outcome mean" row is the unconditional mean of the Panel B outcome. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A15: Robustness checks: alternative energy-cost shock definitions

	Baseline	Aggregation horizon		Energy input	Sample restriction
	(1)	(2)	(3)	(4)	(5)
Definition	12 months	6 months	18 months	Natural gas	Manufacturing share $> 50\%$
Panel A)	Dependent variable: Green patent count				
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0784***	0.0870***	0.0777**	0.0676**	0.0876***
	(0.0284)	(0.0268)	(0.0322)	(0.0277)	(0.0311)
	[+2.0%]	[+2.2%]	[+2.0%]	[+1.7%]	[+2.2%]
Breadth_{t-1}	0.6456***	0.6455***	0.6451***	0.6464***	0.6358***
	(0.0674)	(0.0675)	(0.0675)	(0.0675)	(0.0711)
Observations	583,525	583,525	583,525	583,525	544,305
Panel B)	Dependent variable: Breakthrough share				
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0356***	-0.0292**	-0.0348***	-0.0329***	-0.0362**
	(0.0136)	(0.0118)	(0.0108)	(0.0120)	(0.0149)
	[−7.1%]	[−5.8%]	[−6.9%]	[−6.6%]	[−7.4%]
Breadth_{t-1}	0.0145	0.0141	0.0148	0.0141	0.0191
	(0.0168)	(0.0168)	(0.0168)	(0.0168)	(0.0167)
Outcome mean	0.1264	0.1264	0.1264	0.1264	0.1229
Adjusted R^2	0.415	0.415	0.415	0.415	0.401
Observations	65,541	65,541	65,541	65,541	61,901
Inventor FE	Y	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y

Note: This table presents estimates of equation (2) using the inventor panel. Panel A reports Poisson pseudo-maximum-likelihood (PPML) estimates with the count of the inventor's green patents as the outcome. Panel B reports OLS estimates with the share of the inventor's green patents that are breakthroughs as the outcome, where a breakthrough is a patent in the top decile of the importance measure (Section 2.2). Column (1) uses the baseline energy-cost shock. Columns (2) and (3) construct the predicted energy-input price changes over six and eighteen months, respectively, rather than the baseline twelve months. Column (4) uses the natural-gas component of the baseline shock alone. Column (5) restricts the baseline specification to firms with a manufacturing sales share above 50 percent. Each shock is lagged one year and standardized to be zero-mean and have a standard deviation of 1. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). The bracketed row reports the implied percent change in the outcome per one standard deviation increase in the shock, for an inventor one standard deviation above mean breadth, expressed as a percent of the outcome mean in Panel B. The "Outcome mean" row is the unconditional mean of the Panel B outcome. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A16: Robustness checks: non-corporate patenting

	(1)	(2)	(3)	(4)
		Excl. inventors with university/gov. patents		
	Baseline	During corporate career	Ever	Résumé-matched sample
Panel A)	Dependent variable: Green patent count			
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0784***	0.0834***	0.0862***	0.1097***
	(0.0284)	(0.0280)	(0.0288)	(0.0424)
	[+2.0%]	[+2.1%]	[+2.2%]	[+2.8%]
Breadth_{t-1}	0.6456***	0.6915***	0.7203***	0.6260***
	(0.0674)	(0.0630)	(0.0643)	(0.0845)
Observations	583,525	527,055	510,716	269,212
Panel B)	Dependent variable: Breakthrough share			
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0356***	-0.0378***	-0.0361***	-0.0500**
	(0.0136)	(0.0144)	(0.0134)	(0.0207)
	[−7.1%]	[−7.5%]	[−7.1%]	[−10.0%]
Breadth_{t-1}	0.0145	0.0180	0.0173	0.0250
	(0.0168)	(0.0169)	(0.0172)	(0.0217)
Outcome mean	0.126	0.127	0.128	0.124
Adjusted R^2	0.415	0.410	0.412	0.413
Observations	65,541	58,789	57,024	26,725
Inventor FE	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y

Note: This table presents estimates of equation (2) using the inventor panel. Panel A reports Poisson pseudo-maximum-likelihood (PPML) estimates with the count of the inventor's green patents as the outcome. Panel B reports OLS estimates with the share of the inventor's green patents that are breakthroughs as the outcome, where a breakthrough is a patent in the top decile of the importance measure (Section 2.2). Column (1) reports the baseline estimates. Columns (2) and (3) exclude inventors with a patent assigned to a university or government entity, identified from PatentsView assignee types and organization names. Column (2) excludes inventors with such a patent filed between their first and last corporate patents, and column (3) excludes inventors with such a patent at any point in their career. Column (4) restricts the sample to inventors matched to Revelio résumé data, as in column (2) of Table 9. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. The bracketed row reports the implied percent change in the outcome per one standard deviation increase in the shock, for an inventor one standard deviation above mean breadth, expressed as a percent of the outcome mean in Panel B. The "Outcome mean" row is the unconditional mean of the Panel B outcome. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A17: Robustness checks: alternative quantity outcomes

	Green patents			Other patents		
	(1)	(2)	(3)	(4)	(5)	(6)
Outcome	Dummy	Log count	Share	Dummy	Log count	Share
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0099***	0.0088***	0.0146***	-0.0056	-0.0086*	-0.0146***
	(0.0030)	(0.0027)	(0.0054)	(0.0038)	(0.0046)	(0.0054)
	[+1.1%]	[+0.2%]	[+1.5%]	[−0.3%]	[−0.2%]	[−0.5%]
Breadth_{t-1}	0.0365***	0.0205**	0.0805***	0.0184	0.0209*	-0.0805***
	(0.0078)	(0.0089)	(0.0101)	(0.0136)	(0.0122)	(0.0101)
Inventor FE	Y	Y	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Outcome mean	0.237	0.208	0.244	0.535	0.614	0.756
Adjusted R^2	0.180	0.267	0.301	0.254	0.444	0.301
Observations	584,187	584,187	371,602	584,187	584,187	371,602

Note: This table presents OLS estimates of equation (2) using alternative patent count outcomes. Columns (1) and (4) use indicators for filing any green and non-green patent, respectively, as the outcome. Columns (2) and (5) use the log of one plus the corresponding patent count. Columns (3) and (6) use the corresponding share of the inventor's patents. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. The bracketed row reports the implied change in the outcome per one standard deviation increase in the shock, for an inventor one standard deviation above mean breadth, as a percentage of the outcome mean. The "Outcome mean" row is the unconditional mean of each outcome. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A18: Robustness checks: alternative fixed effects and firm controls

	Green		Other	
	(1)	(2)	(3)	(4)
Panel A)	Dependent variables: Patent counts			
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	0.0371	0.0532**	-0.0040	-0.0083
	(0.0267)	(0.0248)	(0.0171)	(0.0205)
	[+0.9%]	[+1.4%]	[−0.1%]	[−0.2%]
Breadth_{t-1}	-0.1385**	0.6268***	-0.1027	0.1506***
	(0.0629)	(0.0723)	(0.0702)	(0.0502)
Observations	576,190	574,317	576,190	574,496
Panel B)	Dependent variables: Breakthrough shares			
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0228**	-0.0328**	-0.0032	-0.0018
	(0.0116)	(0.0132)	(0.0029)	(0.0034)
	[−4.6%]	[−6.6%]	[−0.8%]	[−0.4%]
Breadth_{t-1}	0.0702***	0.0012	0.0580***	0.0283***
	(0.0172)	(0.0175)	(0.0132)	(0.0071)
Outcome mean	0.126	0.126	0.104	0.104
Adjusted R^2	0.131	0.351	0.063	0.219
Observations	65,015	65,015	222,685	222,685
Firm FE	Y	Y	Y	Y
Industry×year FE	Y	Y	Y	Y
Inventor FE	N	Y	N	Y
Controls	Y	Y	Y	Y

Note: This table presents estimates of equation (2) using the inventor panel, replacing firm-by-year fixed effects with firm and four-digit-NAICS-by-year fixed effects, and including lagged firm financial controls. Panel A reports Poisson pseudo-maximum-likelihood (PPML) estimates with the count of green patents in columns (1) and (2), and other patents in columns (3) and (4), as the outcome. Panel B reports OLS estimates with the corresponding breakthrough share as the outcome, where a breakthrough is a patent in the top decile of the importance measure (Section 2.2). Columns (2) and (4) add inventor fixed effects. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. The bracketed row reports the implied percent change in the outcome per one standard deviation increase in the shock, for an inventor one standard deviation above mean breadth, expressed as a percent of the outcome mean in Panel B. The "Outcome mean" row is the unconditional mean of each Panel B outcome. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, career age squared, and lagged firm size, liquidity, leverage, profitability, and R&D intensity. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A19: Robustness checks: alternative breakthrough thresholds

	Comparison group (top decile)				Breakthrough threshold (NAICS4×year)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Definition	Baseline	NAICS3 ×year	NAICS4×year green patents	NAICS3×year green patents	Top 5%	Top 15%	Top 20%
$\text{Breadth}_{t-1} \times \text{Shock}_{t-1}$	-0.0356***	-0.0320**	-0.0403**	-0.0328**	-0.0290**	-0.0392**	-0.0379**
	(0.0136)	(0.0145)	(0.0160)	(0.0133)	(0.0130)	(0.0164)	(0.0159)
	[−7.1%]	[−7.2%]	[−9.9%]	[−8.0%]	[−10.9%]	[−5.4%]	[−4.1%]
Breadth_{t-1}	0.0145	-0.0011	-0.0005	-0.0062	-0.0050	0.0203	0.0283
	(0.0168)	(0.0128)	(0.0133)	(0.0111)	(0.0096)	(0.0208)	(0.0215)
Inventor FE	Y	Y	Y	Y	Y	Y	Y
Firm×year FE	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y
Outcome mean	0.126	0.113	0.103	0.103	0.067	0.183	0.235
Adjusted R^2	0.415	0.436	0.389	0.414	0.366	0.448	0.467
Observations	65,541	65,445	64,311	65,088	65,541	65,541	65,445

Note: This table presents OLS estimates of equation (3) using alternative definitions of breakthrough innovation. The outcome is the share of the inventor's green patents classified as breakthroughs. Columns (1) through (4) vary the reference group used to define the top decile: all patents within NAICS4-by-year in column (1), all patents within NAICS3-by-year in column (2), green patents within NAICS4-by-year in column (3), and green patents within NAICS3-by-year in column (4). Columns (5) through (7) retain the baseline NAICS4-by-year reference group and define breakthroughs as patents in the top 5, 15, and 20 percent, respectively. Breadth is the lagged dispersion of an inventor's prior patenting across fields, defined as one minus the HHI index over the CPC main groups of prior patents excluding Y02 climate classification codes (Section 2.3). Shock is the lagged energy-cost shock, as defined in Section 2.4. The shock is standardized to be zero-mean and have a standard deviation of 1. The bracketed row reports the implied change in the outcome per one standard deviation increase in the shock, for an inventor one standard deviation above mean breadth, as a percentage of the outcome mean. The "Outcome mean" row is the unconditional mean of each outcome. The sample includes inventor-years with a scored green patent through 2016. All columns control for the share of the inventor's prior patents that are green, the log of one plus the inventor's lagged patent stock, career age, and career age squared, and include inventor and firm-by-year fixed effects. Standard errors (in parentheses) are clustered two-way by firm and inventor. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.