Induced Innovation, Inventors, and the Energy Transition*

Eugenie Dugoua† Todd Gerarden‡
London School of Economics Cornell University

August 28, 2023

Abstract

We study how individual inventors respond to economic incentives to work on clean electricity technologies. Our empirical strategy leverages variation in natural gas prices to estimate output and entry elasticities of individual inventors. We use the estimates to quantify the medium-run effects of a natural gas price increase equivalent to the social cost of carbon. Total clean patenting would increase roughly one third. The dominant mechanism of this aggregate response is increased patenting by existing clean inventors. Our findings suggest a role for policy to increase the supply of clean inventors in order to increase clean innovation.

***JEL* Codes:** Q55, Q40, O30.
Keywords: Inventors, Energy technology, Induced innovation.

*We gratefully acknowledge funding from Cornell Center for Social Sciences, the NBER-Sloan conference on the Economics of Innovation in the Energy Sector, and LSE STICERD. We thank Pierre Azoulay, Josh Feng, Ashley Langer, Ralf Martin, and David Popp for their helpful comments, as well as the participants of the many seminars where this paper was presented and discussed. Lingxiao Cui, Suzy Guahk, Matías Navarro, Ha Pham, and Yukun Wang provided excellent research assistance.

†Department of Geography and Environment, Center for Economic Performance and Grantham Research Institute, London School of Economics. Email: e.dugoua@lse.ac.uk. Website: eugeniedugoua.com.

‡Charles H. Dyson School of Applied Economics and Management, Cornell University, Warren Hall, Ithaca, NY 14853. Email: gerarden@cornell.edu. Website: toddgerarden.com.

1 INTRODUCTION

Clean energy innovation is critical to reducing the costs of climate change mitigation and allowing society to avert the worst-case scenarios projected by climate scientists. Governments and businesses around the world have announced ambitious environmental targets, many of which hinge on assumptions about future technological improvements. While prior economic literature provides a clear rationale for subsidizing clean technology research and development (R&D) to achieve these objectives, there is limited evidence to guide the development of policy to spur clean energy innovation.

In this paper, we provide empirical evidence on the role of individual inventors in the energy transition. We focus on individual inventors because they are the crucial input into innovation and because understanding how individual inventors respond to incentives can inform the design of future policy. We use comprehensive global data on patent applications to characterize the careers of individual inventors. First, we identify inventors working on electricity generation technologies using patents' technological classifications. We then extract these inventors' patent applications and classify them as either clean, grey, or dirty electricity technologies.[1]

We document two new stylized facts about energy inventors in general, and about clean energy inventors in particular. First, we find that most inventors specialize in either clean or dirty technologies. This is consistent with models of human capital accumulation in which path dependency increases the returns to specialization, and it raises the question of whether future government policies to encourage a shift from dirty to clean technologies may be impeded by frictions that make it difficult for individual inventors to work in different fields. Second, 55% of the clean patent families in the data came from inventors who had not patented before in clean. This sizeable number highlights the crucial role of new entrants in clean innovation.

We then study how individual inventors respond to economic incentives in order to develop

1. In clean technologies, we include renewable energy technologies (i.e., solar, wind, marine, geothermal, and hydro) and nuclear energy. Grey technologies relate to energy efficiency, biomass, and waste combustion. Dirty technologies include patents related to the combustion of fossil fuels. We also collect data on any patents these inventors produce that are not related to electricity technologies.

a deeper understanding of the forces determining these stylized facts. Our primary measure of economic incentives is the price of natural gas, which is arguably the most important factor price in electricity markets. When natural gas is more expensive, clean technologies become relatively more competitive, and demand for them increases. Thus, if firms and inventors expect higher natural gas prices to persist, they have greater incentive to invent and improve clean electricity technologies.

Our empirical strategy leverages variation in natural gas prices over both countries and time to examine how inventors respond to changes in factor prices at both the intensive and the extensive margins. The residual variation in natural gas prices that we exploit stems primarily from supply shocks that are not transmitted globally due to transportation constraints.

First, we focus on active clean inventors and estimate an intensive margin output elasticity to quantify how the number of patents an inventor produces responds to natural gas prices. We use panel data methods to model how natural gas prices affect the number of clean energy patents an inventor produces, including inventor and time fixed effects to account for cross-sectional differences as well as common shocks to innovation incentives. To do so, we use information on the firms that individual inventors patent with. This leverages the role of firms, which effectively act as intermediaries that observe market signals and respond by organizing and directing inventors' research activities. Specifically, we construct measures of the incentives individual inventors are exposed to based on the incentives that the firms they patent with are exposed to. In turn, the firms' incentives are based on the markets in which they are active in innovation.

We also implement a complementary empirical strategy to mitigate concerns about the potential endogeneity of natural gas prices and the fact that inventors are likely to respond differently to transient shocks than to persistent shocks. We isolate variation from the shale gas revolution, which shifted out the supply of natural gas and generated a persistent reduction in the price of natural gas in North America relative to other regions due to natural gas transportation constraints. We leverage this variation by instrumenting for natural gas prices with an indicator variable that takes on a value of one for countries in North America after the shale gas revolution and zero otherwise.

Second, we examine the extent to which economic incentives induce new inventors to enter clean patenting. We estimate an extensive margin elasticity, which we refer to as an entry elasticity, to quantify how the number of inventors entering clean technology responds to natural gas prices. To do so, we shift our analysis to the firm level. We assemble a panel of firms patenting in clean energy and identify inventors listed on a firms' patents in a given year. Within those, we focus on inventors who are filing their first clean patent. We use inventors' patenting history to classify them as either: having never patented before; having patented outside of energy; or having patented in grey or dirty but not clean technologies. We count the number of inventors in each group and then estimate the elasticity of the number of new clean technology inventors with respect to natural gas prices for each group.

Together, these empirical strategies allow us to characterize how inventors respond along both the intensive and extensive margins and to compare the magnitudes of the responses. At the intensive margin, we find that a 10% increase in natural gas prices induces about 5% more clean families for the average clean incumbent. The direction and magnitude of the effect are consistent with prior work at the firm and technology levels. The instrumented elasticity estimates are qualitatively consistent with the non-instrumented estimates. At the extensive margin, we find that a 10% increase in natural gas prices leads to an increase in entry of up to 6% depending on the time horizon and type of entrant.

Finally, we combine these econometric estimates to study the potential effects of an increase in natural gas prices equivalent to the social cost of carbon of $51 per metric ton of carbon dioxide. We find that total clean patenting would increase roughly one third relative to baseline patenting rates in the medium run. The dominant mechanisms of this aggregate response are increased patenting by existing clean inventors and, to a lesser extent, patenting by new entrants who had not previously produced patents.

Overall, these findings show that induced innovation relies primarily on the intensive margin, that is, on already-active inventors. These results are essential to inform green innovation policy. Our findings highlight that the entry of new inventors plays a more minor role in the induced innovation response in the medium run. This highlights the need for further work to understand better what drives individuals to become clean inventors and, in particular, what specific policies could help produce more clean inventors.

This paper provides new empirical evidence to the literature on the economics of energy and environmental innovation. Prior research on directed technical change provides a theoretical justification for subsidizing clean technology R&D to redirect scientific activity toward clean technology (e.g., Acemoglu et al. 2012; Acemoglu et al. 2016; Fried 2018; Hart 2019; Lemoine, Forthcoming). Empirical analyses of patents have shown that fossil fuel price increases and environmental policies induce innovation in clean technologies at the industry and firm level (e.g., R. G. Newell et al. 1999; Popp 2002; Johnstone et al. 2010; Popp and R. Newell 2012; Noailly and Smeets 2015; Aghion et al. 2016). However, these aggregate analyses are unable to identify the extent to which induced innovation comes from changes in the use of labor, other inputs, or both. We provide new empirical evidence on how high-skilled workers respond to incentives that can be used to guide future modeling assumptions and policy design. To our knowledge, ours is the first paper in this strand of literature to examine the response of individual inventors.

This paper also contributes to the literature studying the drivers and barriers of innovation at the level of individual inventors or scientists (e.g., Romer 2000; Jones 2010; Azoulay et al. 2011; Aghion et al. 2017; Azoulay et al. 2019; Bell et al. 2019a, 2019b; Einiö et al. 2019; Agarwal and Gaule 2020; Van Reenen 2021; Akcigit et al. 2022). Deming and Noray (2020) provides evidence that STEM workers are increasingly specialized and their skills depreciate quickly. Myers (2020) also documents that the switching costs for scientists can be large in the context of biomedical research.

Our paper also builds on a growing literature that studies the impacts of the shale gas revolution. This literature exploits significant variation in U.S. natural gas prices due to the advent of horizontal drilling and hydraulic fracturing, or "fracking." Hausman and Kellogg (2015) assess the welfare and distributional implications of shale gas. Another strand of literature focuses on the implications of the change in natural gas prices on the electricity sector and environmental outcomes in the short run (e.g., Cullen and Mansur 2017; Linn and Muehlenbachs 2018; Knittel et al. 2019; Coglianese et al. 2020). Fowlie and Reguant (2022) exploit variation in the shale revolution's effects on natural gas prices across locations and industries to simulate the effects of a domestic carbon price on U.S. manufacturing.

We contribute to this literature by exploiting slightly different variation and studying different outcomes. Prior papers primarily use variation within the U.S. for estimation. By contrast, we leverage the significant change in natural gas prices in North America relative to other regions of the world to study how fuel price changes induce innovation by individual inventors.[2] This innovation could have transformational effects on environmental, electricity sector, and broader economic outcomes in the long run.

2 Stylized Facts About Energy Inventors

2.1 Data

Energy Patent Data. We use data on patent applications from the Worldwide Patent Statistical Database (PATSTAT) from the European Patent Office (2022). We extract patents related to electricity generation using patent classification codes.[3] These codes help us identify patents related to specific technologies for electricity generation, and we classify those technologies as either clean, grey, or dirty. Clean technologies include zero or low-carbon electricity generation technologies (i.e., solar, wind, marine, geothermal, hydro, and nuclear).[4] Dirty technologies include patents related to the combustion of fossil fuels (i.e., coal, oil, and natural gas).[5] In grey technologies, we group patents related to improving the efficiency of combustion processes (Lanzi et al. 2011). We

2. Acemoglu et al. (2019) present suggestive evidence of the impact of shale gas development on green innovation as motivation for a theoretical model of the long-run consequences of the shale gas revolution.

3. We use codes from the Cooperative Patent Classification and the International Patent Classification building on previous studies that have listed relevant energy codes (Johnstone et al. 2010; Lanzi et al. 2011; Dechezleprêtre et al. 2014; Popp et al. 2020).

4. A patent family is classified as "clean" if it has at least one code related to renewables or nuclear. We also consider an alternative definition of "clean" that includes some enabling technologies relevant to electricity, and excludes families that include any grey or dirty codes. Results for that definition are in the appendix.

5. Emissions intensities vary significantly across different fuels and technologies. We use the simplistic terminology "clean" and "dirty" for broad categorizations in keeping with prior work (e.g., Acemoglu et al. 2012; Aghion et al. 2016).

also classify electricity generation from biomass and waste as grey technologies since they still emit greenhouse gases despite being cleaner than traditional fossil fuels.

We aggregate patent applications at the level of patent families, which are collections of patents that are considered to cover the same technical content and therefore represent the same invention.

We date families by their priority year, which is the year when the earliest application within the family was filed.

Online Appendix Figure C.1 plots the number of clean, grey, and dirty patent families over time in our sample. The trends are similar to those documented previously by Acemoglu et al. (2019) and Popp et al. (2020), with the number of clean patent families increasing rapidly over the 2000s, followed by a decline in clean patenting since 2010. By contrast, the number of new patent families in grey and dirty technologies has been more stable over the past three decades.

Inventor Data. Our next step is to identify individual inventors to construct a panel dataset of their patenting activity over time. Intellectual property authorities such as the European Patent Office or the U.S. Patent and Trademark Office (USPTO) require that the names of all individuals who contributed to the invention be included as inventors on the application. However, patent offices do not typically use unique identifiers for individual inventors. To analyze inventors' activity over their careers, researchers must therefore use the inventor names written on patent applications to identify unique inventors. Li et al. (2014) provide such a disambiguation effort for patents filed with the USPTO between 1975 and 2010. Since we study energy inventors globally, we take similar steps to disambiguate inventors in the PATSTAT data.

Our starting point is to use the PATSTAT Standardized Name identifier, which is the result of a harmonization procedure completed prior to publication of the data. This harmonization, however, was not done on all observations: 70% of the inventors in our sample were not included. As a result, we improve on the PATSTAT identifier by standardizing inventors' names and then disambiguating inventors based on string matching.[6]

For our analysis, we focus on inventors who are listed on at least one energy patent application

6. Online Appendix A.2 explains this procedure in detail.

filed in an OECD country after 1990.[7] We define the year when the inventor becomes connected to a family as the earliest year when the inventor appears on any of the applications in the family. In the end, our sample contains a total of 873,256 energy inventors.

2.2 Descriptives

Most Energy Inventors Specialize in Clean or Dirty Technologies. Figure 1a shows the extent to which energy inventors specialize in either clean, grey, or dirty patenting based on inventors' global patent portfolios over the period spanning 1990 to 2019. To construct the graph, we classify inventors with at least one energy patent family in a given year according to their last three years of patenting.

On average throughout the period, 30% of energy inventors patent in clean energy technologies only. Inventors who patent in grey and/or dirty energy technologies are more numerous, making up 60% of energy inventors.[8] By contrast, the share of energy inventors who are active in clean energy patenting, as well as dirty and/or grey energy patenting, is much smaller, at 10%.

Figure 1a also shows how specialization has changed over time. The total number of energy inventors increased until 2012, led by a rapid rise in the number of clean inventors during the 2000s. During that period, the share of inventors working in clean energy roughly doubled, though it has plateaued since 2010 as energy patenting decreased. On the other hand, the number of inventors working on dirty and/or grey energy technologies grew more gradually over time, so that their share fell significantly over the 2000s. Finally, while the number of inventors working in both areas has increased over time, it remains small relative to the clean and dirty categories.

New Entrants are a Quantitatively Important Source of Clean Patenting To assess the contribution of different types of inventors to innovation output, we document the number of clean families produced by inventors based on their prior patenting behavior. Figure 1c summarizes the

7. We limit our geographic scope because natural gas price data is available for OECD countries only.

8. Here, for simplicity, we restrict our attention to energy-related patents since we are unable to classify non-energy patents as clean, grey, or dirty. Hence, when we say that an inventor patents only in clean, we mean that all of the energy patents the inventor produces are in clean. The inventor may also patent in other non-energy fields.

distribution of clean families over the sample period. To compute these numbers, we inversely weight patent counts by the number of inventors associated with each patent family to avoid double-counting (which we refer to as "fractional" patent counts), and then aggregate patent counts across inventors of each type.

On average, throughout the period, we find that only about half (46%) of clean families are from clean "incumbents," either inventors with prior patenting in clean technologies only (30%) or in clean as well as grey and/or dirty technologies (16%). Roughly one-third (32%) of families come from inventors who did not previously appear in the patent data. About 19% come from inventors that had previously patented in fields that we do not classify as energy. Finally, a small fraction (4%) of clean families come from inventors with prior patenting in grey and/or dirty but not clean technologies.[9]

9. We find similar distributions of incumbents versus entrants for grey and dirty families (see Online Appendix C.3).

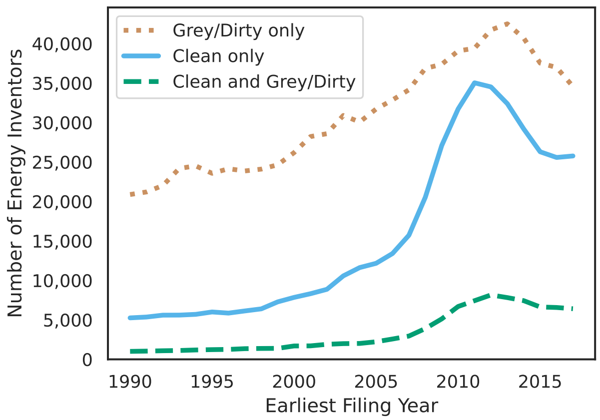

(a) Energy Inventors by Type (Levels)

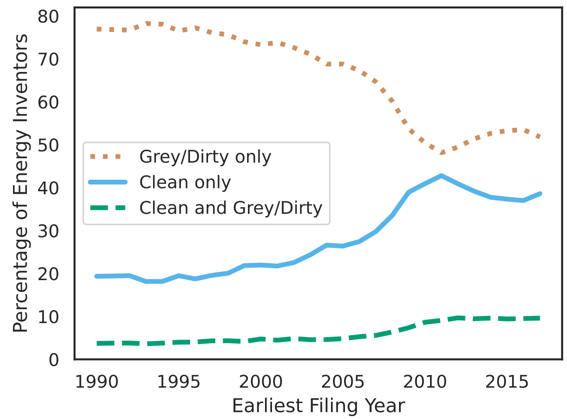

(b) Energy Inventors by Type (Shares)

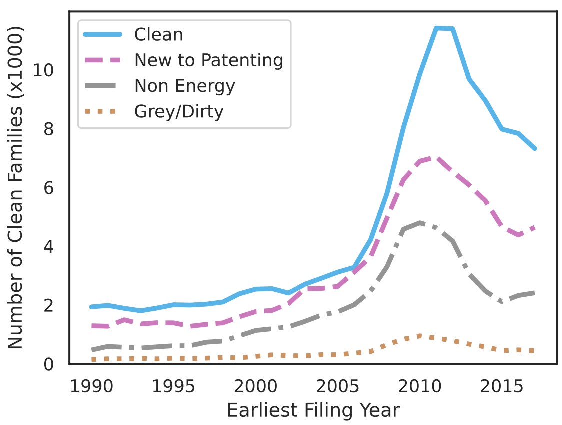

(c) Clean Families by Origin of Inventor (Levels)

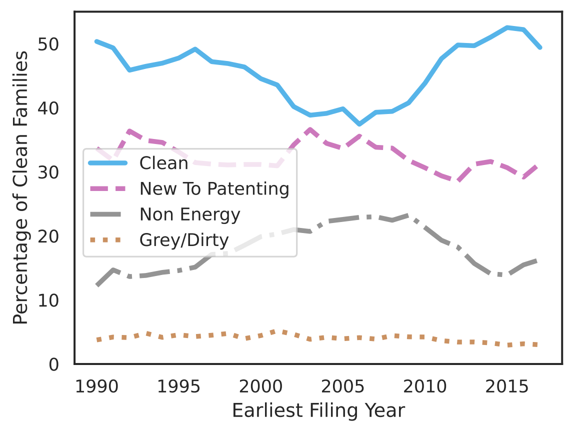

(d) Clean Families by Origin of Inventor (Shares)

FIGURE 1

Type of Energy Inventors and Clean Patent Families

Note: Figures 1a and 1b show the extent to which energy inventors specialize in either clean, grey, or dirty patenting. We focus on inventors' global patent portfolios for inventors with at least one energy patent in an OECD country after 1990. To construct the graphs, we first identify inventors with at least one energy family filed in year t, and then classify them according to their last three years of patenting activity. Figures 1c and 1d illustrate the types of inventors behind clean families over time. They plot trends over time in the levels and shares of clean families produced by inventors with previous clean patents, inventors new to patenting, inventors with previous patents outside the set of energy technologies under study, and inventors with previous grey and/or dirty patents. Families with multiple inventors are fractionally attributed to the inventors to avoid double-counting.

3 Empirical Strategy

The remainder of the paper focuses on how innovation in clean electricity generation technologies responds to changes in economic incentives, which we proxy by changes in natural gas prices. In this section, we first discuss the sources of variation in natural gas prices that we exploit. We then explain our approach to estimating clean innovation responses on both the intensive and extensive margins.

3.1 Identifying Variation

Our empirical strategy builds on a theoretical literature on induced innovation dating to Hicks (1932). Hicks hypothesized that a change in relative factor prices would spur innovation to use less of the factor which had become relatively expensive. Several studies have used patent data to test this theory in the context of energy technologies (e.g., Popp 2002; Noailly and Smeets 2015; Aghion et al. 2016). We extend this work by focusing on individual inventors of technologies used for electricity generation. We use natural gas prices as a proxy for relative factor prices in electricity generation, and therefore as an indirect proxy for the expected returns from innovation in renewable and nuclear electricity generation technologies that compete with natural gas-fired electricity generation.[10]

We use data on natural gas prices from the International Energy Agency (2020) and exploit variation across countries and time, visualized in Figure 2a.[11] The price variation across countries at a given point in time stems primarily from constraints on the transportation of natural gas. The clearest example of this is the shale gas revolution. The development of horizontal drilling and hydraulic fracturing caused prices for natural gas in North America to decline significantly in 2009.

10. While renewable and nuclear technologies primarily serve as substitutes to fossil fuel technologies, they can also be complements in some markets and time periods. This affects the interpretation but not the validity of our analysis. The Online Appendix also presents results using a broader definition of "clean" that includes enabling technologies such as smart grid and energy storage. However, the extent to which those enabling technologies are substitutes or complements to natural gas electricity generation is less clear than for clean electricity generation technologies.

11. The International Energy Agency (2020) natural gas prices are in nominal U.S. dollars per megawatt-hour. All econometric analysis in the paper includes time fixed effects, which absorb common time-varying factors including changes in the value of U.S. dollars due to inflation, so the results are invariant to using prices in real terms.

These price reductions were not seen in other regions for many years due to short-run capacity constraints on the export of natural gas.

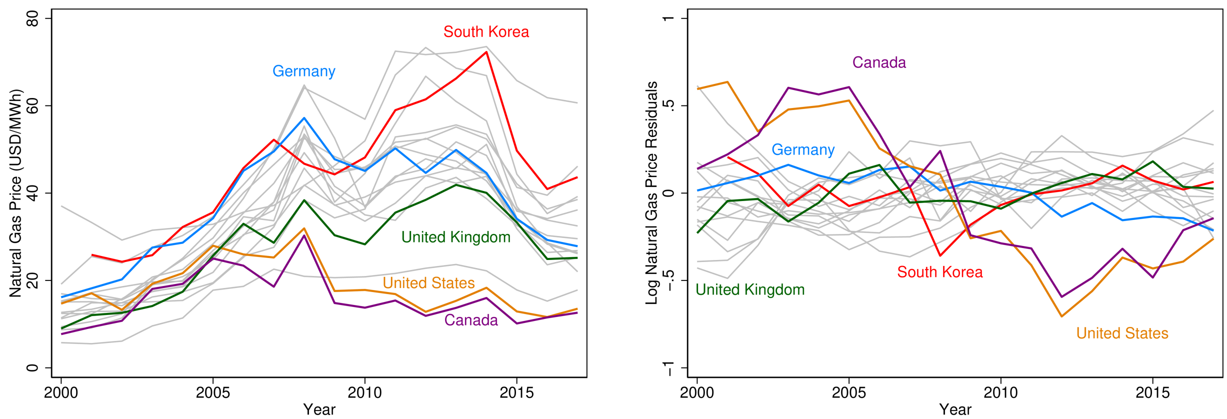

(a) Prices (b) Residuals

FIGURE 2
Natural Gas Prices and Residuals across Countries and Time

Note: Panel a plots the price of natural gas in each country over time using data from the International Energy Agency (2020). Prices are in U.S. dollars per megawatt-hour (MWh). Panel b plots residuals from a regression of the natural logarithm of the natural gas prices from Panel a on country and year fixed effects.

To mitigate concerns about the potential endogeneity of natural gas prices, we also implement an instrumental variable strategy that restricts attention to the variation in natural gas prices caused by the shale gas revolution. As described above, improvements in natural gas extraction led to significant and persistent declines in natural gas prices in North America in the middle of our sample period. The price declines were not transmitted to other regions of the world due transportation constraints.[12] These patterns are evident in the raw price data plotted in Figure 2a.

The identifying variation used in our primary empirical strategy comes from residual variation in natural gas prices after removing variation explained by country and time fixed effects. Figure 2b plots this residual variation. It is clear from this plot that a significant portion of the identifying variation stems from the change in prices before versus after the shale gas revolution for the United States and Canada.

We isolate this variation using a binary instrument that is one for the United States and Canada

12. In addition, most other countries with shale gas deposits banned hydraulic fracturing for reasons unrelated to clean energy innovation.

starting in 2009 (when the shale revolution started to take effect) and is zero in all other countries and time periods. [13] This instrument explains 51% of the residual variation in natural gas prices after accounting for country fixed effects, time fixed effects, and other control variables included in our main specification. We use a control function approach based on Lin and Wooldridge (2019) to implement the instrumental variable strategy, detailed in Appendix E.

3.2 Response at the Intensive Margin: Output Elasticity of Incumbents

To quantify the magnitude of the induced innovation response at the intensive margin, we focus on incumbent inventors that have produced at least one clean family and study to what extent they increase clean patenting in response to increases in natural gas prices. Specifically, we model patenting as a function of energy prices and inventor characteristics:

$$PAT_{it}^{C} = \exp(\beta_P \ln P_{it-1} + \beta_X X_{it-1} + \gamma_t + \eta_i) + u_{it}, \tag{1}$$

where PAT_{it}^{C} is the count of clean patent families by inventor i in year t;[14] P_{it-1} is the price of natural gas that inventor i is exposed to in year $t-1$;[15] X_{it-1} is a set of controls; γ_t and η_i denote year and inventor fixed effects; and u_{it} is an error term. In some specifications, we also include tenure fixed effects to account for how productivity evolves over the course of inventors' careers.[16] We estimate

13. Mexico is excluded from the sample for econometric analysis due to limited data on the price of natural gas.

14. We construct inventors' time-series such that the first year corresponds to the first observed patent (of any type) filed by the inventor, and the last year corresponds to the year of the last observed patent (of any type). Indeed, although we directly observe when inventors "enter" patenting, we do not know for sure when they "exit". We can, therefore, only safely input zeros for years when inventors do not file patents when these years come in between the first and last filing years of the inventor. The appendix provides a robustness check by showing results when arbitrarily truncating inventors' time-series at 50% of their observed tenure.

15. We use the previous year's prices as a proxy for individual inventors' beliefs about future prices while still allowing a lag that gives inventors time to respond to variation in price. While we do not have direct evidence on individual inventors' beliefs about natural gas prices, Anderson et al. (2013) find that U.S. consumer beliefs about gasoline prices are indistinguishable form a no-change forecast. Furthermore, Aghion et al. (2016) use a similar model for automotive firm patenting and find that lagging fuel prices by one year is reasonable compared to more flexible dynamic forms. We also estimate more flexible distributed lag models that include prices from the previous three years. This choice of lags is supported by survey evidence on inventor activities from Nagaoka and Walsh (2009), who report that the average amount of calendar time spent on research leading up to a patent application is less than two years, and that between 80% and 90% of patents involve three or fewer years of research leading up to an application.

16. We construct the tenure variable by counting the number of years since we observe an inventor's first patent.

equation 1 via Poisson pseudo maximum likelihood under the assumption that natural gas prices are conditionally weakly exogenous.

Our empirical model requires a proxy for the natural gas prices that individual inventors use to form beliefs, which we do not directly observe. For inventors who patent in conjunction with corporations, we view their incentives as primarily mediated by firms. Thus, we construct price measures for each individual that depend upon the firm(s) they are associated with are exposed to.[17] This helps us quantify the economic incentives inventors are exposed to through their employers.

We, therefore, construct inventor-specific prices in two steps. First, we compute firm-specific prices as the weighted average of country-level prices. Second, we compute inventor-specific prices as the weighted average of firm-level prices. The resulting prices are given by

$$\ln P_{it} = \sum_{j} s_{ij} \sum_{c} \frac{s_{jc} GDP_c}{\sum_c s_{jc} GDP_c} \ln P_{ct},$$

where, within the inner summation: P_{ct} is the average tax-inclusive natural gas price in country c in year t, GDP_c adjusts for differences in market size across countries, and s_{jc} captures exposure of firm j to country c. We calculate s_{jc} as firm j's share of energy patents in country c.[18] This method of constructing firm-specific prices is similar to prior analyses of induced innovation at the firm level (e.g., Noailly and Smeets 2015; Aghion et al. 2016). In the outer summation, s_{ij} is the share of inventor i's patent families that are associated with firm j.[19] We use the same weighting method to construct inventor-specific measures of the other country-level variables contained in X_{it-1}.

In the set of controls X_{it-1}, we include the GDP per capita and public spending on energy and low-carbon research, development, and demonstration (RD&D) spending by governments that inventor i is exposed to in year $t-1$. These factors are included because they are likely to

17. Patent applications always provide the names of inventors but also those of applicants (i.e., the entities detaining the intellectual property rights). More often than not, applicants are for-profit organizations. We, therefore, connect inventors to firms based on the applicants that appear on the patents. We use Orbis IP to link applicants with specific Orbis identifiers. Independent "garage" inventors who are not associated with any firms represent 16% of individual inventors in the data. For these inventors, we use the price of their home country.

18. We use the full sample to construct these shares because 65% of firms in the estimation sample do not apply for patents prior to 2000. In robustness analysis, we use pre-period patenting to mitigate concerns about the potential endogeneity of the shares. For firms without any pre-period patenting, we assign the global weighted-average price.

19. We use the full sample to construct these shares because 71% of inventors do not patent before 2000.

influence patenting, and they may be correlated with natural gas prices. We use country-year level data from the International Energy Agency (2019) on government spending on energy RD&D, and country-year level data from the World Bank (2020a, 2020b) on GDP and GDP per capita in purchasing power parity terms. We construct inventor-level controls using the same weighting strategy as for prices.

3.3 Response at the Extensive Margin: Entry Elasticity of Inventors

Next, we examine whether changes in natural gas prices induce inventors who have not previously worked on clean energy technology to begin patenting. Because we only observe inventors once they patent and do not observe their education or career history, we are unable to use within-inventor variation in natural gas prices to study extensive margin responses. Instead, we use firm-level information on patenting portfolios and the inventors they patent with. For each firm in each year, we count the number of inventors filing clean families with the firm for the first time.[20] We use these data to estimate a firm-level model analogous to the inventor-level model in equation 1:

$$E_{jt}^{k} = \exp(\beta_P^k \ln P_{jt-1} + \beta_X^k X_{jt-1} + \gamma_t^k + \eta_j^k) + u_{jt}^k, \tag{2}$$

where E_{jt}^{k} is the number of new entrant inventors of type k filing a clean family with firm j in year t.[21] We classify entrants into three types: those who previously patented in grey and/or dirty but not clean energy technologies, those who previously patented in technology areas outside of the set of energy technologies studied in this paper, and those who were not previously observed in the patent data. P_{jt-1} is the price of natural gas that firm j is exposed to in year $t-1$. We include in X_{jt-1} the GDP per capita as well as energy and low-carbon RD&D spending by governments that firm j is exposed to in year $t-1$. Firm-level prices and covariates are constructed as described in Section 3.2. Year and firm fixed effects are denoted γ_t^k and η_j^k, and u_{jt}^k is an error term. We estimate these

20. The link between PATSTAT inventors and Orbis firms is provided by Orbis IP. The coverage of the correspondence is severely limited after 2014. For this reason, we restrict our firm-level sample to years between 2000 and 2014.

21. To avoid double-counting inventors who file patents with multiple firms, we weigh the relationship between a firm and an inventor by the inverse number of firms the inventor patented with in that year.

models separately by type.

4 RESULTS

4.1 Output Elasticity Estimates

Table 1 contains estimates of the elasticity of clean patenting with respect to lagged natural gas prices. Panel A presents baseline results from models that include fixed effects and use all residual variation in natural gas prices. Panel B presents results from instrumental variable models that only use natural gas price variation from the shale gas revolution. Panel C presents results from a distributed lag model which uses all residual variation in natural gas prices in the three years prior to patenting. The columns contain alternative specifications of equation 1. All columns include year fixed effects, inventor fixed effects, and additional covariates that vary across countries and time; some columns also include tenure fixed effects. The first two columns use the simple count of clean patent families as the outcome variable. The third and fourth columns use the count of clean families weighted by the number of citations they received.[22] The last two columns use the simple count of clean patent families inversely weighted by the number of coinventors associated with each patent family (i.e., "fractional" count).[23][24]

In Panel A, all six specifications yield output elasticities around 0.5, implying that a 10% increase in natural gas prices would cause a 5% increase in clean patenting by incumbent inventors. The effect is somewhat larger when families are weighted by citations, indicating that price variation affects the production of higher-quality patents on the margin. By contrast, it is somewhat smaller when using fractional patent families, suggesting that price variation affects patenting by teams more than by individual inventors on the margin.

22. Specifically, for a family filed in year t, the weight is equal to the ratio of the number of citations the family received within three years over the number of citations that the average energy family filed in year t received.

23. For example, if an inventor produced one clean patent family in a given year in conjunction with one other inventor, the outcome would be 0.5 rather than 1. We use this approach to avoid double-counting and to facilitate comparisons to extensive margin responses in Section 4.3.

24. We also document results with other outcome variables in the appendix (e.g., granted, triadic, US-only families).

TABLE 1
Estimates of Incumbent Inventors' Elasticity of Patenting with Respect to Natural Gas Prices

	Count of Clean Patent Families					
	Simple Count		Citation-Weighted		Coinventor-Weighted	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Baseline Poisson estimates						
Prices (log, t-1)	0.548***	0.463***	0.635***	0.533***	0.468***	0.389***
	(0.037)	(0.037)	(0.047)	(0.048)	(0.047)	(0.047)
Inventors	110,454	110,454	110,454	110,454	110,454	110,454
Observations	763,630	763,630	763,630	763,630	763,630	763,630
Pseudo-R2	0.285	0.286	0.369	0.371	0.261	0.262
Panel B: Instrumental variable estimates						
Prices (log, t-1)	0.512***	0.299***	0.963***	0.703***	0.360***	0.162*
	(0.069)	(0.071)	(0.081)	(0.084)	(0.085)	(0.087)
Inventors	110,454	110,454	110,454	110,454	110,454	110,454
Observations	763,630	763,630	763,630	763,630	763,630	763,630
First-stage F-statistic	163	163	163	163	163	163
Panel C: Distributed lag estimates						
Cumulative effect (3 lags)	0.642***	0.546***	0.652***	0.565***	0.622***	0.511***
	(0.050)	(0.052)	(0.069)	(0.070)	(0.057)	(0.061)
Inventors	85,905	85,905	85,905	85,905	85,905	85,905
Observations	590,767	590,767	590,767	590,767	590,767	590,767
Pseudo-R2	0.289	0.290	0.366	0.367	0.264	0.265
Year fixed effects	X	X	X	X	X	X
Inventor fixed effects	X	X	X	X	X	X
Tenure fixed effects		X		X		X
Country-year covariates	X	X	X	X	X	X

Note: The dependent variables are the number of clean patent families, either unweighted, weighted by citations, or inversely weighted by the number of coinventors, depending on the column. Panels A, B, and C contain estimates of the same parameters using different estimation strategies. Panel A presents estimates of equation 1 estimated via Poisson pseudo-maximum likelihood. Standard errors are clustered by inventor and reported in parentheses. Panel B presents estimates of equation E.2 estimated via the control function approach described in the text, using the shale gas revolution as an instrument for natural gas prices. Standard errors are constructed via block bootstrap of the two-step control function approach, sampling inventors 250 times with replacement. The first-stage F-statistic for the instrumental variable estimates is from estimating equation E.1 at the country-year level rather than the inventor-year level, since the instrument varies at the country level and it thus provides a more conservative assessment of the instrument's strength. Panel C is analogous to Panel A except that the models include three lags of natural gas prices and all other covariates that vary across both countries and time, and the coefficients represent cumulative effects.

Panel B of Table 1 presents estimates from the instrumental variable strategy. Overall, the qualitative patterns across columns are similar to those in Panel A. The instrumented estimates of the elasticity of the count of patent families with respect to natural gas prices, in columns 1 and 2, are similar to the non-instrumented estimates in Panel A. Columns 3 and 4 present elasticities for citation-weighted patent families, which are larger in magnitude than the non-instrumented estimates.

The most likely explanation for the differences between Panels A and B is that the price variation used to identify the output elasticity is different, and that the local average treatment effect of the instrument is different from the average treatment effect.[25] The shale gas revolution generated a large decline in natural gas prices in North America that was expected to persist far into the future. This expectation of persistent price changes could have had a larger impact on the incentives for engaging in high-risk, high-reward innovation that is more likely to be cited than it had on the incentives for more incremental innovation (relative to other, potentially transient price variation).

In Panel C, we present results from a distributed lag version of the baseline Poisson model as a complementary approach to capture the long-run effects of persistent price changes. The elasticity estimates are quite consistent across columns. The cumulative effect estimates for the citation-weighted outcomes lie in between the estimates from Panels A and B. This is consistent with the transient versus persistent nature of the price variation explaining the differences between the baseline and instrumental variable estimates. Given this, and given that a large fraction of the overall variation in the data is driven by the shale revolution, we focus on the non-instrumented results for the remainder of the paper.[26]

4.2 Entry Elasticity Estimates

Table 2 contains estimates for the entry elasticity with respect to lagged natural gas prices. Each column corresponds to a different type of entrant. Panel A presents estimates from models with one lag. Panel B presents the cumulative effect from distributed lag models with three lags. In Panel A, the estimates are positive but somewhat imprecise. The entry elasticity point estimate is largest for new inventors who had not previously patented. In Panel B, the estimates for new-to-patenting and grey/dirty entrants are larger and more precisely estimated. The change in magnitude is intuitive, as

25. Other potential explanations for the differences include price endogeneity and sampling variation.

26. Appendix F also contains results for a broader definition of clean patenting that includes enabling technologies. The estimates are smaller in magnitude than the main estimates, which is as expected since enabling technologies are not direct substitutes for electricity generated from natural gas.

inventors and firms are forward-looking, and therefore likely to respond less to transient than to persistent price changes. On the other hand, we do not find clear evidence that non-energy inventors respond to price shocks.

TABLE 2
Estimates of the Elasticity of Inventor Entry with Respect to Natural Gas Prices

	Number of Clean Inventors		
	New to Patenting (1)	From Grey/Dirty (2)	From Non-Energy (3)
Panel A: Baseline Poisson estimates			
Prices (log, t-1)	0.258**	0.167*	0.044
	(0.110)	(0.096)	(0.127)
Firms	3,933	4,993	4,912
Observations	53,921	67,617	66,541
Pseudo-R2	0.692	0.605	0.643
Panel B: Distributed lag estimates			
Cumulative effect (3 lags)	0.618***	0.456***	-0.062
	(0.166)	(0.124)	(0.181)
Firms	3,779	4,703	4,642
Observations	43,733	53,109	52,559
Pseudo-R2	0.699	0.605	0.647
Year fixed effects	X	X	X
Firm fixed effects	X	X	X
Country-year covariates	X	X	X

Note: The dependent variables are the fractional number of inventors (that is, inversely weighted by the number of firms they are associated with) of each type within each firm who are new to patenting in clean patent families in that year. The sample used for estimation is a balanced panel of firms from 2000 to 2014. Panel A presents estimates of equation 2 estimated via Poisson pseudo-maximum likelihood. Standard errors are clustered by firm and reported in parentheses. Panel B is analogous to Panel A except that the models include three lags of natural gas prices and all other covariates that vary across both countries and time, and the coefficients represent cumulative effects.

These results suggest that extensive margin responses to natural gas price variation are largest for new entrants, followed by inventors who previously produced grey and/or dirty patents. The response of inventors who were previously active in other technology areas is negligible. However, these elasticities provide an incomplete picture of these extensive margin effects, because the baseline number of inventors in each group is different. The next section uses a back-of-the-envelope calculation to place these extensive margin estimates in context with one another and with the intensive margin estimates from Section 4.1.

4.3 How would Carbon Pricing Induce Innovation?

To place the intensive and extensive elasticity estimates in context, we analyze the effects of a persistent natural gas price increase equivalent to the U.S. Government's social cost of carbon of \$51 per metric ton of carbon dioxide. This corresponds to 54% of the GDP-weighted global average price of natural gas in 2014, which we use as the reference year. We model the medium-run effects of this price increase over the course of 10 years.

To calculate the aggregate impact of this change in natural gas prices, we use a first-order approximation that combines responses along the intensive and extensive margins. We use the estimated elasticities from Sections 4.1 and 4.2 along with data on baseline rates of patenting and entry to compute the contribution of each margin.[27] The extensive margin responses are computed separately by entrant type, and take into account typical patenting rates over the first 10 years after an inventor enters clean patenting. Appendix H provides a formal description of our approach and more details on its implementation.

Table 3 summarizes the results. In the medium run, intensive margin responses by incumbent inventors are the largest source of induced patenting. Within the extensive margin responses, entry to patenting by new inventors is quantitatively most important. Extensive margin responses by inventors who had previously produced patents related to grey and dirty technologies are next most important. Finally, entry by inventors who had previously worked on other technologies outside of the energy sector contributes a relatively small, though imprecisely estimated, amount. In total, this represents an increase of 26% relative to a scenario in which the baseline rate of clean patenting from 2014 had been constant over the 10-year period.

To assess the sensitivity of these results, we present analogous estimates using distributed lag models as well as using an unbalanced panel of firms for the extensive margin analysis in Online Appendix Table H.1.[28] While the absolute magnitudes of patenting activity depend on the

27. To avoid double-counting, inputs are based on "fractional" numbers: we use elasticities estimated using the count of clean families inversely weighted by the number of coinventors (for the intensive margin) and the number of inventors inversely weighted by the number of firms they are associated with (for the extensive margin).

28. In addition, Online Appendix Table H.2 presents results using 2010 as the base year. Online Appendix Table H.3 presents analogous results for a broader definition of clean patenting that includes enabling technologies.

specification, the relative importance of each margin does not: in all cases, the largest sources of induced patenting activity are increased patenting by incumbent inventors, followed by entry of new inventors without prior patents.

TABLE 3
Predicted Impacts of Carbon Pricing on Clean Patenting

Source	**Patents**	**Share (%)**
Intensive margin response		
Incumbent inventors	36,710	80.7
	(4,438)	(6.2)
Extensive margin response		
Entry from grey/dirty	1,615	3.6
	(927)	(2.0)
Entry from non-energy	540	1.2
	(1,557)	(3.4)
Entry to patenting	6,616	14.5
	(2,807)	(5.5)
Total	45,481	100.0
	(5,555)	.

Note: Predicted changes in the number of clean patent families due to a persistent 54% increase in natural gas prices over the course of 10 years, relative to a base year of 2014. The total change in patenting represents an increase of 26% relative to baseline patenting rates. Output and entry elasticities are estimated using a single lag of natural gas prices as in Panel A of Tables 1 and 2. Inputs for the extensive margin analysis are derived from a balanced panel of firms from 2000 through 2014 as in Table 2. Standard errors are constructed using the delta method.

Some caveats and discussion are in order. First, this analysis is an approximation. While the price change we study is on the same order of magnitude as the country-level natural gas price variation observed in the raw data, our first-order approximation does not account for higher-order effects of natural gas prices on innovation by incumbents and entry by new inventors. If the supply of patents or inventors are highly convex, our predictions may overstate the magnitude of induced innovation.

Second, our analysis focuses on the effects of a change in natural gas prices. In reality, carbon pricing would also increase the price of other emitting sources of electricity generation such as coal. Furthermore, economy-wide carbon pricing could lead to increased demand for electricity from other sectors, such as electric vehicle charging from the transportation sector, which would also affect the returns to clean innovation. Both of these effects are beyond the scope of our analysis.

Third, our analysis does not account for differences in the quality of different patents. The results in Table 3 are simple counts of patent families. Given the relative magnitudes of the estimates in Table 1, the effects are likely to be larger for alternative measures such as citation-weighted patents that attempt to proxy for the quality of innovations.

5 CONCLUSION

We draw three sets of conclusions from the empirical evidence in this paper. First, inventors are highly specialized: most inventors patenting in electricity generation technologies work exclusively on either clean, grey, or dirty technology. Consequently, about half of clean patents are produced by inventors with prior clean patents. Furthermore, clean patenting output by these inventors responds to changes in natural gas prices.

Second, the other half of clean patents come from new entrants, highlighting their important role in clean innovation. But perhaps surprisingly, we find that entry by these inventors does not respond strongly to variation in natural gas prices.

Third and finally, our analysis of carbon pricing shows that induced innovation is driven primarily by intensive margin increases in the patenting output of incumbent inventors. Extensive margin entry of new inventors plays a more minor role. This contrasts with the roughly equal split of patenting between the two groups from the descriptive analysis. These differences underscore the importance of designing policy based on which innovation activities are most likely to respond on the margin rather than on average.

These findings raise the question of whether government policies to encourage a shift from dirty to clean technologies may be impeded by frictions that make it difficult for individual inventors to work in different fields. In particular, our finding that induced innovation relies primarily on the intensive margin highlights the need for further work to understand better what drives individuals to become clean inventors and what specific policies could help produce more clean inventors.

References

Acemoglu, Daron, Philippe Aghion, Lint Barrage, and David Hémous. 2019. *Climate Change, Directed Innovation, and Energy Transition: The Long-run Consequences of the Shale Gas Revolution.* Working Paper.

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous. 2012. "The Environment and Directed Technical Change." *American Economic Review* 102 (1): 131–166.

Acemoglu, Daron, Ufuk Akcigit, Douglas Hanley, and William Kerr. 2016. "Transition to Clean Technology." *Journal of Political Economy* 124 (1): 52–104.

Agarwal, Ruchir, and Patrick Gaule. 2020. "Invisible Geniuses: Could the Knowledge Frontier Advance Faster?" *American Economic Review: Insights* 2 (4): 409–424.

Aghion, Philippe, Ufuk Akcigit, Ari Hyytinen, and Otto Toivanen. 2017. *The Social Origins of Inventors.* Working Paper 24110. National Bureau of Economic Research, December.

Aghion, Philippe, Antoine Dechezleprêtre, David Hémous, Ralf Martin, and John Van Reenen. 2016. "Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry." *Journal of Political Economy* 124 (1): 1–51.

Akcigit, Ufuk, John Grigsby, Tom Nicholas, and Stefanie Stantcheva. 2022. "Taxation and Innovation in the Twentieth Century." *The Quarterly Journal of Economics* 137 (1): 329–385.

Anderson, Soren T., Ryan Kellogg, and James M. Sallee. 2013. "What Do Consumers Believe about Future Gasoline Prices?" *Journal of Environmental Economics and Management* 66 (3): 383–403.

Azoulay, Pierre, Christian Fons-Rosen, and Joshua S. Graff Zivin. 2019. "Does Science Advance One Funeral at a Time?" *American Economic Review* 109 (8): 2889–2920.

Azoulay, Pierre, Joshua S. Graff Zivin, and Gustavo Manso. 2011. "Incentives and Creativity: Evidence from the Academic Life Sciences." *The RAND Journal of Economics* 42 (3): 527–554.

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen. 2019a. "Joseph Schumpeter Lecture, EEA Annual Congress 2017: Do Tax Cuts Produce More Einsteins? The Impacts of Financial Incentives Versus Exposure to Innovation on the Supply of Inventors." *Journal of the European Economic Association* 17 (3): 651–677.

———. 2019b. "Who Becomes an Inventor in America? The Importance of Exposure to Innovation." *The Quarterly Journal of Economics* 134 (2): 647–713.

Coglianese, John, Todd D. Gerarden, and James H. Stock. 2020. "The Effects of Fuel Prices, Environmental Regulations, and Other Factors on U.S. Coal Production, 2008-2016." *The Energy Journal* 41 (1).

Cullen, Joseph A., and Erin T. Mansur. 2017. "Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution." *American Economic Journal: Economic Policy* 9 (3): 106–133.

Dechezleprêtre, Antoine, Ralf Martin, and Myra Mohnen. 2014. *Knowledge Spillovers from Clean and Dirty Technologies.* CEP Discussion Paper 1300.

Deming, David J., and Kadeem Noray. 2020. "Earnings Dynamics, Changing Job Skills, and STEM Careers." *The Quarterly Journal of Economics* 135 (4): 1965–2005.

Einiö, Elias, Josh Feng, and Xavier Jaravel. 2019. *Social Push and the Direction of Innovation.* Working Paper. May.

European Patent Office. 2022. *PATSTAT Global 2022 - Single Edition (Spring).* https://www.epo.org/searching-for-patents/business/patstat.html.

Fowlie, Meredith L., and Mar Reguant. 2022. "Mitigating Emissions Leakage in Incomplete Carbon Markets." *Journal of the Association of Environmental and Resource Economists* 9 (2): 307–343.

Fried, Stephie. 2018. "Climate Policy and Innovation: A Quantitative Macroeconomic Analysis." *American Economic Journal: Macroeconomics* 10 (1): 90–118.

Hart, Rob. 2019. "To Everything There Is a Season: Carbon Pricing, Research Subsidies, and the Transition to Fossil-Free Energy." *Journal of the Association of Environmental and Resource Economists* 6 (2): 349–389.

Hausman, Catherine, and Ryan Kellogg. 2015. "Welfare and Distributional Implications of Shale Gas." *Brookings Papers on Economic Activity,* 71–125.

Hicks, John R. 1932. *The Theory of Wages.* London: Macmillan.

International Energy Agency. 2019. *International Energy Agency Energy Technology Research and Development Database, 1974-2017.* https://doi.org/10.5257/iea/et/2018-09.

———. 2020. *International Energy Agency Energy Prices and Taxes, 1960-2019.* https://doi.org/10.5257/iea/ept/2019Q4.

Johnstone, Nick, Ivan Haščič, and David Popp. 2010. "Renewable Energy Policies and Technological Innovation: Evidence Based on Patent Counts." *Environmental and Resource Economics* 45 (1): 133–155.

Jones, Benjamin F. 2010. "Age and Great Invention." *The Review of Economics and Statistics* 92 (1): 1–14.

Knittel, Christopher R., Konstantinos Metaxoglou, and André Trindade. 2019. "Environmental Implications of Market Structure: Shale Gas and Electricity Markets." *International Journal of Industrial Organization* 63:511–550.

Lanzi, Elisa, Elena Verdolini, and Ivan Haščič. 2011. "Efficiency-Improving Fossil Fuel Technologies for Electricity Generation: Data Selection and Trends." *Energy Policy,* Asian Energy Security, 39 (11): 7000–7014.

Lemoine, Derek. Forthcoming. "Innovation-Led Transitions in Energy Supply." *American Economic Journal: Macroeconomics.*

Li, Guan-Cheng, Ronald Lai, Alexander D'Amour, David M. Doolin, Ye Sun, Vetle I. Torvik, Amy Z. Yu, and Lee Fleming. 2014. "Disambiguation and Co-Authorship Networks of the U.S. Patent Inventor Database (1975–2010)." *Research Policy* 43 (6): 941–955.

Lin, Wei, and Jeffrey M. Wooldridge. 2019. "Chapter 2 - Testing and Correcting for Endogeneity in Nonlinear Unobserved Effects Models." In *Panel Data Econometrics,* edited by Mike Tsionas, 21–43. Academic Press, January.

Linn, Joshua, and Lucija Muehlenbachs. 2018. "The Heterogeneous Impacts of Low Natural Gas Prices on Consumers and the Environment." *Journal of Environmental Economics and Management* 89:1–28.

Myers, Kyle. 2020. "The Elasticity of Science." *American Economic Journal: Applied Economics* 12 (4): 103–34.

Nagaoka, Sadao, and John P. Walsh. 2009. *The R&D Process in the US and Japan: Major Findings from the RIETI-Georgia Tech Inventor Survey.* RIETI Discussion Paper Series 09-E-010. Research Institute of Economy, Trade and Industry.

Newell, Richard G., Adam B. Jaffe, and Robert N. Stavins. 1999. "The Induced Innovation Hypothesis and Energy-Saving Technological Change." *The Quarterly Journal of Economics* 114 (3): 941–975.

Noailly, Joëlle, and Roger Smeets. 2015. "Directing Technical Change from Fossil-Fuel to Renewable Energy Innovation: An Application Using Firm-Level Patent Data." *Journal of Environmental Economics and Management* 72:15–37.

Popp, David. 2002. "Induced Innovation and Energy Prices." *The American Economic Review* 92 (1): 160–180.

Popp, David, and Richard Newell. 2012. "Where Does Energy R&D Come from? Examining Crowding out from Energy R&D." *Energy Economics* 34 (4): 980–991.

Popp, David, Jacquelyn Pless, Ivan Haščič, and Nick Johnstone. 2020. *Innovation and Entrepreneurship in the Energy Sector.* Working Paper, Working Paper Series 27145. National Bureau of Economic Research, May.

Romer, Paul M. 2000. "Should the Government Subsidize Supply or Demand in the Market for Scientists and Engineers?" *Innovation Policy and the Economy* 1:221–252.

Van Reenen, John. 2021. "Innovation and Human Capital Policy." In *Innovation and Public Policy,* 61–84. Chicago: University of Chicago Press.

World Bank. 2020a. *GDP (Constant 2017 US$).* Accessed October 14, 2020. https://data.worldbank.org/NY.GDP.MKTP.PP.KD.

———. 2020b. *GDP Per Capita (Constant 2017 US$).* Accessed October 14, 2020. https://data.worldbank.org/NY.GDP.PCAP.PP.KD.