

# Induced innovation under competitive pressure

Leon Bremer

WORK IN PROGRESS – Version of February 27, 2025

## Abstract

Because directed technical change has the potential to delink economic activity from environmental degradation, understanding the drivers of environment-related innovation is valuable to policy makers. This paper empirically investigates two drivers of clean innovation, energy prices and market competition, grounded in the induced innovation and inverted-U literature, respectively. Using a self-constructed dataset, I find evidence of the induced innovation hypothesis, but only when studying energy price changes, not when studying energy price levels. I do find some evidence for an inverted-U relationship between competition and innovation, although the relationship is statistically uncertain and disappears when the quadratic relationship is relaxed. When combining the two hypotheses, the results are inconclusive. I do not find consistent and significant evidence that market competition affects the induced innovation effect. While the induced innovation findings are robust against different econometric specifications and estimation methods, they vary depending on whether energy prices are first-differenced or not. The literature would benefit from foundations for clarifying which specification is preferred.

## 1 Introduction

As new technologies are invented that contribute to disconnecting pollution from economic activity, we are better able to sustain economic activity while addressing problems with pollution and climate change. These so-called clean innovations thereby contribute to a societal good. Policy makers aiming to promote such innovations would benefit from understanding the factors influencing firms' innovation decisions. This understanding could for instance be exploited when devising environmental policies.

This paper adds to the literature by investigating the combination of two related strands of literature. On the one hand, the literature on induced innovation discusses how higher input prices drive firms towards the invention and adoption of more efficient technologies to especially economize on the more expensive inputs. On the other hand, scholars have been interested in the effect of market competition on innovation. Do competitive pressures promote or hinder innovation? This paper combines the two fields, investigating the relationship between energy input prices and

environment-related innovation under different market competition intensities. I will test the hypotheses separately and simultaneously.

I construct a unique data set of mostly European firms. Methods are proposed and implemented to overcome two main challenges. First, I exploit the work from Bremer (2023) which applies a machine learning algorithm to merge the patent data to firm financial data based on firm names. The algorithm scores names on similarity, thereby allowing for some degree of different spelling and misspellings in the data. Second, the location of operations is often obscure for larger firms. Firms could operate multiple plants, and the head offices do not need to be located in the same country as the production or innovation facilities. I address this ambiguity by studying innovation behavior on the firm group level. This avoids within-group specialization from introducing measurement error. The resulting sample covers a variety of industries and European countries.

Innovation is measured by successful patent applications. To understand the direction of innovation, patent classification codes are used to categorize patents into clean and dirty innovations. To estimate the effects of energy prices and competition on patenting, count models are fitted using regression. I closely follow the empirical literature in the field to obtain several alternative estimation results. I also highlight the choice about which variation to exploit for the induced innovation hypothesis. Should one consider energy price levels, a measure consisting of mostly cross-sectional variation? Or should one consider energy price changes, a measure isolating the time variation?

I find that the findings for induced innovation from energy prices depends on the type of variation that is exploited. When considering energy price changes, measured as growth rates, I find positive and statistically significant results, in line with the induced innovation hypothesis. This result is consistent across estimation methods, like Poisson, negative binomial or instrumental variable specifications. On the other hand, the regression specifications considering energy price levels result in insignificant coefficient estimates of the energy price on clean patenting. When studying the effects of energy prices on dirty and overall innovation, the results are mixed and mostly statistically insignificant.

These results show the importance of the identification strategy. When using energy price levels, one implicitly focuses more on cross-sectional differences. On the one hand, this could be desirable to answer the induced innovation hypothesis, as sectors facing higher energy prices might continually be induced to patent clean innovations. On the other hand, one might want to focus on the time variation in energy prices, as that is more likely the exogenous shock leading to a change in innovation strategy. Firms likely know about the different fuels used in particular processes and what their prices are. But they likely do not know how fuel prices change. The latter argues in favor of focusing on the time variation in energy prices by using energy price changes.

Estimating the quadratic relationship between competition and innovation yields findings of an inverted-U relationship. However, these findings are statistically uncertain. Moreover, the fitted relationship predicts a peak of the inverted-U curve at relatively low levels of competition.

When combining the induced innovation hypothesis with the hypothesized inverted-U relationship between competition and innovation, the results are inconclusive. To test the combined hypotheses, the energy price term is interacted with the competition terms in the regressions, and subsequently the marginal effects of energy prices on innovation are estimated for different levels of market competition. When considering energy price levels there is again no evidence for induced clean innovation, but energy prices do negatively affect dirty innovation outcomes for low and medium levels of market competition. When studying energy price changes, the only statistically significant finding is that of induced clean innovation at low levels of market competition.

These findings also draw two important conclusions regarding methodologies in the induced innovation literature. First, the induced innovation hypothesis tests yield different answers depending on the form in which energy prices enter the estimation. When energy price growth rates are considered, I find strong evidence for induced clean innovation from energy prices. When energy price levels are considered, I find no such evidence. Second, these findings are highly robust to the estimation method used. I employ Poisson, negative binomial and two instrumental variable approaches but the induced innovation conclusions are not challenged.

Furthermore, I show that the induced innovation conclusions change somewhat when analyzing innovation on the firm level instead of the firm group level. Where the conclusion for energy price changes is unchallenged, the conclusion from the energy price levels specification is different. There is again no clean induced innovation, but there is a statistically significant reduction in dirty innovation in response to energy price levels. The clean induced innovation estimate is also closer to zero, while the estimate for dirty innovation is now larger in absolute terms. As many studies in the induced innovation literature are based on firm-level analyses using energy price levels, this is an interesting finding.

The literature also regularly studies Triadic patents only. These are innovations that are patented at all three major patent offices around the world, namely in Europe, Japan and the US. The main drawback is that these patents are uncommon and lead to significantly fewer observations. This likely explains why I find no statistically significant effects for the two hypotheses that I test for. However, the estimates for the energy price level specifications are in line with the induced innovation hypothesis. This again shows that methodological considerations are important for the induced innovation literature.

Whereas all firms are expected to experience energy-saving incentives from increases in energy input prices, firms in energy-intensive industries likely experience

the strongest incentives. Testing this hypothesis yields evidence for a story of budget constraints and directed technical change. I find that firms in energy-intensive industries reduce their overall innovation output in response to energy input price increases. These reductions are driven by dirty innovation, but not clean innovation. On the other hand, firms operating in industries with low energy intensity do not significantly reduce overall or dirty patenting in response to energy price increases, but they do increase clean patenting.

The remainder of this paper is structured as follows. [Section 2](#) discusses the relevant literature. Variable construction, the data and its sources are discussed in [Section 3](#). [Section 4](#) develops and discusses the econometric model and empirical findings are presented in [Section 5](#). [Section 6](#) concludes.

## 2 Related literature

There is a broad range of literature that relates to the interplay of innovation, energy prices, competition and the environment. I split this section in three parts. First, I discuss relevant literature on induced innovation. Second, the literature on the role of market competition in innovation is discussed. Third, knowledge stocks and spillovers are discussed, each important drivers of patenting outcomes. A discussion of statistical methods used in the literature can be found in [Section 4.2](#) of this paper when I cover methodologies.

While I discuss the literature in separate sections, studies across sections are intimately related. Several studies also cover topics from multiple sections. I will discuss them in the section where it is most relevant for this paper, but I acknowledge that one could come to a different grouping.

### 2.1 Induced innovation

Following Hicks (1932), the induced innovation hypothesis states that firms innovate to save on inputs that become relatively more expensive. In general, firms can achieve certain input savings while maintaining the same output level in two ways. One response could be to substitute some of the expensive inputs for relatively cheaper inputs. The difficulty and profitability of substitution depend on the substitution elasticities and the relative prices. Another response is to change the production methods, by investing in new technologies or by developing new technologies through innovation activities. These new technologies are expected to rely less on the more expensive input. The latter is the channel of interest of the induced innovation hypothesis.

Studying Hick's induced innovation hypothesis, Newell et al. (1999) estimate the energy efficiency of consumer durables in response to energy prices. They use Sears catalogs to collect information on consumer durables and their attributes. They mainly find that energy prices did induce energy efficiency of consumer durables after

product energy labelling was introduced. This shows that energy prices do induce changes in products on offer, likely through the innovation channel. It also shows that this effect requires consumers to have access to information, as shown by the effect of energy labelling. When consumers cannot compare energy efficiency between products, the producers of these goods put in less effort to improve energy efficiency. Arguably the issue of information is less likely to play a role for firms trying to reduce their energy bill. Firms likely have better information about their production processes and the energy efficiency of their equipment, because the firms themselves are paying the energy bill.

Empirical literature has also found evidence of the induced innovation hypothesis when studying responses to changes in energy input prices. Often patent data is used to study innovation. In that category, Popp (2002) tests the induced innovation hypothesis using citation-weighted patent counts. He uses a distributed lag model with adaptive expectations for the energy prices. In such a model firms form expectations about future energy prices based on a weighted average of past prices (Popp et al., 2010). Popp (2002) finds that energy prices positively affect innovation in energy technologies. Also the knowledge stock positively affects innovation. He further finds that not controlling for the knowledge stock inverts the relationship between innovation and the energy prices.

Linn (2008) tests whether energy input prices have an effect on innovation in the US manufacturing sector. Using plant-level data he finds that the demand for energy responds stronger to energy price changes for entrants compared to incumbents. To conclude that the heterogeneous demand response is due to energy-efficient innovation, one must assume that entrants are capable of acquiring newer technologies. Although statistically significant, the effect of energy prices on energy intensity is small.

A large literature also studies the effects of policies on innovation outcomes. Some studies use a combination of energy prices and policies, by studying tax-inclusive energy prices, or by considering both energy prices and energy or environmental policies. One could also argue that price-based energy policies have the same effect as energy price changes, and hence induced innovation from energy price changes could serve as a proxy for price policies like carbon taxes, as argued by Marin and Vona (2019) and discussed by Jaffe et al. (2003). Yet others take a theoretical approach and study which policies are most effective in promoting innovation and welfare (Fischer et al., 2003).

Aghion et al. (2016) estimate the effect of tax-inclusive fuel prices on clean, gray and dirty patenting in the auto industry. Patents are categorized in these clean, gray and dirty categories according to their IPC codes. These codes refer to highly specific technological fields and are attached to patent applications. The authors find that higher fuel prices stimulate clean innovations and reduce dirty innovations. Note how for car manufacturers the fuel price is not an input in the production process, but rather a demand-side factor. As consumers face these fuel prices, they benefit more or

less from a car's efficiency depending on the prevalent fuel prices. Furthermore, they find path dependence both from the firm's patent stocks as well as from the industry's patent stocks available in the country of operation.

In a similar study, Rozendaal and Vollebergh ([forthcoming](#)) find that fuel standards induce clean innovation in the auto industry. They find no evidence of policies negatively affecting dirty patenting. They also do not find evidence for fuel prices inducing clean innovation. Lastly, they find strong evidence of path dependence in the firm's own patent stocks, but not for spillovers from patent stocks of other firms in the country of operation.

Johnstone et al. ([2010](#)) use patent data to test the effect of different types of public policy on innovation. Using a country-level panel data set, they study innovation in renewable energy technologies. They find that market-based policies are successful if the environment-related technologies are nearly as competitive as fossil fuel technologies. Technologies that are much less competitive than fossil fuel technologies require more targeted policies in order to develop. This finding is confirmed by Nicolli and Vona ([2016](#)), who find that mature policies benefit from quota systems and that newer technologies benefit more from subsidies and R&D support. On the contrary, Nesta et al. ([2014](#)) find no direct evidence of renewable energy policies promoting clean energy innovations. However, when they interact their policy term with an index for market competition, they find that renewable energy policies do induce clean innovation in the more competitive markets.

Calel and Dechezleprêtre ([2016](#)) use a matched difference-in-differences methodology to estimate the effect the EU Emissions Trading System (EU ETS) has on patenting of low-carbon technologies. Comparing plants under the EU ETS to a matched selection of unregulated plants allows the authors to compare similar firms and identify the effect of the EU ETS. They find that the EU ETS increased patenting of low-carbon technologies in Europe with one percent. While statistically significant, this policy-induced innovation effect is not large.

While many of the induced innovation effects are found to be small, Shapiro and Walker ([2018](#)) show there is a large potential for technical change to improve production processes. They investigate why air pollution fell so dramatically in the US manufacturing sector from 1990 to 2008, even though output increased. By decomposing industrial emissions, they isolate which component has been the main driver of emissions reductions. They find that emissions dropped mostly due to within-product changes in emissions intensity and not due to decreasing output or changes in the composition of output. This finding shows that the channel of innovation and adoption of cleaner production methods can be sizable. They further employ a numerical model to estimate the environmental stringency faced by the firms. They show that the increase in environmental stringency can explain most part of the emissions reductions in US manufacturing.

Policy-induced innovation is also heavily discussed in the literature on the Porter

hypothesis (see Porter & Van der Linde, 1995b). The Porter hypothesis states that regulation can be beneficial for the regulated firms. The constraints from regulation force firms to innovate, increasing production efficiency and in turn making firms more competitive. Met with skepticism, follow-up studies have focused on three sub-hypotheses, specified by Jaffe and Palmer (1997). (1) A *narrow* hypothesis claims that outcome-based regulation is more effective than command-and-control regulation, (2) a *weak* hypothesis describes the positive effect of regulation on innovation, and (3) a *strong* hypothesis claims that environmental regulation makes a firm more profitable or competitive. Lanoie et al. (2011) empirically estimate these three forms and find evidence for a weak hypothesis, but not for a strong hypothesis. They find ambiguous results when testing the narrow hypothesis. For a further overview of the literature on the Porter hypothesis, one could inform Ambec et al. (2013).

Many studies on induced innovation use patent data. It has arguably become the main measure of innovation in this literature. While patent statistics have several benefits, they also have drawbacks. The main benefits are that patent data offers a way to quantify innovation outcomes and provide this for many innovators over many years (Griliches, 1990; Haščič & Migotto, 2015). Patent data is available for many countries across the globe and they contain long time series. However, one of the main drawbacks is comparability (Dechezleprêtre, Ménière, & Mohnen, 2017; Haščič & Migotto, 2015; OECD, 2009). Comparability issues occur due to different patenting rates and patent quality across technologies, patent offices, sectors or countries. Solutions are proposed to correct patents for their quality by either weighting patent counts by their number of citations, as in Popp (2002), or by selecting only patents applied for at multiple patent offices, as proposed by Dernis and Khan (2004) who suggest to only consider *Triadic* patents, which are applied for at all three major patent offices in the US, Europe and Japan. Newer text-based methods use patent texts to distill information about newness and additionality of the patent's innovation (Kelly et al., 2021). Hall et al. (2005) empirically find that patent citations do positively affect market value of firms, indicating that more frequently cited patents are more valuable. Abrams et al. (2013), however, find that the relationship between patent citations and firm revenue follows an inverted-U pattern. Although such corrections might aid comparability, they could come at the expense of restricting the number of observed innovations. As not all innovations are patented, one does not gain a complete overview of innovations by studying patent data. Restricting the sample through quality filters exacerbates this issue.

An overview of developments in the induced innovation literature is provided by Sarr and Noailly (2017). And a comprehensive literature review of induced innovation is conducted by Grubb et al. (2021). They also present an overview of the estimated elasticities between clean patenting and energy prices in the empirical literature. This collection of estimates does not show a definitive consensus on the topic, with elasticity estimates covering a wide range and even some negative values.



## 2.2 Competition and innovation

The inverted-U relationship between competition and innovation has been modeled and estimated by Aghion et al. (2005). They estimate a quadratic relationship between competition and innovation on the industry level using data from the UK. They fit this panel data with a Poisson regression and find evidence for an inverted-U relationship. This means that too little or too much competition results in less innovation. In a theoretical model the authors attempt to explain the inverted-U relationship. They argue that there are two opposing forces at play, (1) an escape-competition effect and (2) the Schumpeterian effect of innovation. The factor that determines which effect is dominant is the share of sectors characterized by neck-and-neck competition in relation to the share of sectors characterized by competition between laggards and leaders. The escape-competition effect is experienced by neck-and-neck firms and its incentives to innovate increase with competition. The Schumpeterian effect is experienced by laggards and its incentives to innovate decrease with competition. The authors argue that the endogenous share of neck-and-neck sectors decreases with competition, thereby producing the inverted-U. At low levels of competition there is a large share of neck-and-neck firms, making the low, but upward-sloping incentives of the escape-competition effect dominant. For higher levels of competition the downward-sloping incentives of the Schumpeterian effect become ever more dominant. Therefore the effect of competition on innovation is first upward-sloping and later downward-sloping.

The discussion of the relationship between market competition and innovation goes back to the work of Schumpeter (1950). The Schumpeterian argument, as used by Aghion et al. (2005) above, would be that laggards have fewer incentives to innovate in more competitive markets. The possible rents that can be captured from an innovation are lower in more competitive markets, and hence the incentives to innovate and creatively destroy the existing technologies are lower.

Scherer (1967) was the first to estimate a non-linear relationship between competition and innovation. While not the main aim of the paper, the author noticed how the relationship between competition, measured as concentration, and the number of scientists and engineers employed differs between industry groups. Discussing these results he suggests that part of this heterogeneity could be explained by the level of market competition. Adding a squared competition term to the regressions yields an inverted-U relationship, although the coefficient estimates are statistically uncertain.

The literature later on reverted back to the simpler linear relationship between competition and innovation. Nickell (1996) investigates the relationship between competition and firm performance. He finds evidence for a positive effect of competition on total factor productivity (TFP) growth. This finding is robust against alternative measures of competition and to using an instrumented regression. Competition is either measured from a survey-based indicator of the number of product market competitors or a measure of the profit share. Nickell et al. (1997) confirm this finding using the same profit share measure, which they refer to as *rents*. They further add



interactions between the competition term and either financial pressure or shareholder pressure. Both these last two terms are positively related with productivity growth, and they substitute for competitive pressures.

Blundell et al. (1999) empirically estimate the effect of market share and competition on innovation, using the fixed effects count model developed by Blundell et al. (1995). Where the first term captures the market share of the firm itself, the second term captures the concentration in the industry. They find positive effects of market share on innovation and negative effects of concentration on innovation. Innovation is measured as a simple count of innovations, as collected by a research center. When alternatively using patent counts, the market share results reproduce, but the concentration results become statistically insignificant.

More recently, Bao and Chen (2018) investigate incumbents' responses to the threat of entry of foreign firms over the period of 2000-2007. In their study innovation is measured by patent applications, and the threat of foreign entry is deduced from news articles. They select news about foreign direct investments (FDI), which captures the expected arrival of foreign competition in the market. They find that, amongst other responses, incumbent firms respond to such threats by innovating. The effect is statistically significant, but small. FDI threats result in firms increasing patenting with 1.4%.

These findings are in line with Aghion et al. (2009) who also study innovation incentives to incumbent firms from firm entry. They estimate the effect of foreign entry on patenting outcomes of publicly-traded UK firms over the period 1987-1993. To deal with endogeneity they instrument foreign entry with UK policy reforms. A patent count model is fitted with a zero-inflated Poisson regression. They find that foreign entry promotes patenting. Moreover, they find that this effect reduces with the distance to frontier of the industry the firm is active in. They also control for competition, measured by average industry profitability, and find an inverted-U relationship with patenting.

The literature on competition and innovation mostly does not take into account the direction of innovation. To the best of my knowledge, Aghion et al. (2023) are the first to do this. They ask whether market competition promotes clean innovations. The direction of innovation is measured as the difference in the logs of clean and dirty patent counts. The study focuses on the auto industry. The findings confirm that environmental attitudes in a country cause the direction of innovation to become cleaner. Additionally, they find that market competition, measured by markups, strengthens this effect. So in more competitive markets, the environmental attitudes of people have a greater effect on steering innovation in the clean direction.

Vives (2008) investigates the relationship between competition and innovation theoretically. Developing a microeconomic model with both Bertrand and Cournot competition, he finds that competition generally has positive effects on innovation. The main point, however, is that results depend on how one measures competition

and innovation. The main recommendation for empirical researchers is therefore to choose the competition measure carefully and to control for underlying determinants of competition. For a further discussion of competition measures, I refer to OECD (2021).

Recent studies have also documented relevant insights in market dynamics and the firm's ability to pass through changes in input prices. De Loecker et al. (2020) document an increase in markups, market shares and profits since the early 1980s. The authors show that the increase in average markups is mostly driven by increased market share of firms that already had higher markups. They show that the markup distribution has become more skewed in 2014 compared to 1980. The authors therefore also point to the complexity of deriving statements about the entire market from simple statistics like the average. Furthermore, Ganapati et al. (2020) investigate a firm's ability to pass-through energy input price changes to the consumer. Using plant-level census data for the US manufacturing sector and using shift-share instruments they find that the energy input price pass-through rate is about 70%. They also note that this rate differs between industries, arguing that firms with more market power are able to pass through a larger share of the input price changes. The authors argue that their findings have important implications for policy makers, as likely environmental price policies will not lead to a 100% pass through to consumers.

## 2.3 Knowledge stocks

The concepts of innovation and knowledge are intimately related in economics. Where innovations refer to new ideas, knowledge refers to the accumulation of ideas. When studying innovation, it is difficult to ignore the role of knowledge stocks. Many of the studies discussed above take into account how knowledge stocks affect innovation outcomes. They thereby acknowledge or test the dependence of innovations on existing knowledge stocks. The consensus from these tests is that existing knowledge helps further innovation, the *standing on shoulders* argument.

Most of these studies also distinguish different types of knowledge stocks. For example, some might split knowledge stocks across technologies. The consensus is that innovations in a particular technological field, for instance renewable energy, benefits more from the knowledge stocks in that same technological field rather than in other fields. Another split that is made in most of the above studies is the own knowledge stock versus the knowledge stocks of others. For example, firm-level analyses will consider the firm's own knowledge stocks and knowledge stocks of others that the firm could potentially tap into. Think of competitors' knowledge in the same industry or country. Additionally, some studies delve into the channels through which others' knowledge becomes available to a firm.

Abramovitz (1986) discusses determinants of differences in productivity levels between countries, and especially the dynamics of convergence and divergence. Focusing on convergence, the *catch-up* hypothesis predicts that laggard countries could benefit

from the existing knowledge in leading countries, leading to a higher growth rate and subsequent convergence. However, the author states that convergence is not a given and depends on a country's ability to adopt new technologies. This *social capability* can be determined by many factors, amongst which are cultural and social values. This study highlights the complexity of knowledge flows. Even if knowledge exists in one place, it does not mean it is available to all.

Grossman and Helpman (1991) argue that the knowledge available to a firm also consists of foreign knowledge that spills over. They model these spillovers as having a positive correlation with how well a firm is connected with its foreign counterparts. They argue that those connections are established through international trade. To test this channel of knowledge flows, Coe and Helpman (1995) empirically test the driving forces of countries' total factor productivity (TFP). Using country-level panel data covering the 1970s and 80s, they find that TFP is not only positively affected by domestic R&D capital, but also by foreign R&D capital. Their findings further show that this foreign knowledge spills over through the international trade channel. In a response to Coe and Helpman (1995), Keller (1998) shows that weighting the foreign R&D capital by random trade shares provides an even stronger effect on TFP than the trade-related shares considered by Coe and Helpman (1995). While domestic TFP still relates to foreign R&D capital, it might therefore not be through the international trade channel.

In an attempt to quantify the relevance of different knowledge sources, Jacobs et al. (2002) estimate the effect domestic and foreign R&D have on domestic TFP. The domestic knowledge stock is split by own-sector R&D and other sectors' R&D efforts. To determine the foreign knowledge stock and the other sectors' knowledge stock that is available to a particular sector, the authors use input-output tables and trade data to weigh the external knowledge. This assumes that knowledge spills over through the trade channels. They derive knowledge from R&D expenditures. They use sector-level panel data for the Netherlands to empirically estimate the effects of these three R&D knowledge stocks on TFP. The findings show that a one percent increase in the sector's own knowledge stock increases TFP with 0.37%, while this is 0.15% and 0.03% for domestic other sectors' knowledge and foreign knowledge, respectively.

A significant literature also uses patent data to study knowledge flows and spillovers. Not only can patent data be used to measure innovation outcomes and knowledge stocks, patent citations are frequently used to quantify knowledge flows. As patents can cite other patented work, a citation might indicate the use of previous knowledge in a new innovation. For example, Jaffe et al. (1993) use patent citations to study cross-border knowledge flows. Using inventors' locations of the cited and citing patent, the authors study the geographical dimension of knowledge flows. They consider three geographical levels, namely the country, state and metropolitan level. They further consider only patents applied for in the US in either 1975 or 1980. The

findings show that about two-thirds of citations are from within the same country, the US, some 10% comes from the same state, and some 4-9% comes from within the same metropolitan area. To determine whether knowledge spillovers are localized, the authors compare these rates to control rates that are based on expected locations of citation links under the assumption of no localization. This exercise shows that the citations within metropolitan area are the most overrepresented, followed by the state level. US patent citations in 1975 and 1980 were therefore highly localized.

Some studies estimate differences in the spillover effects between clean and dirty technologies. Noailly and Shestalova (2017) find that renewable energy technologies have larger spillovers than fossil fuel technologies. They use patent citations as an indicator for knowledge spillovers between and within renewable energy and fossil fuel patents. Also Dechezleprêtre, Martin, and Mohnen (2017) find that innovations in clean technologies show higher levels of spillovers than innovations in dirty technologies and innovations in other emerging technologies. Using patent citation data they estimate spillovers of clean innovations to be 40% higher than for the other innovations. Verdolini and Galeotti (2011) use patent citations to construct available foreign knowledge stocks. They show that both domestic as well as foreign knowledge stocks promote energy-related patent applications. They also show that energy prices foster these patent applications.

Although environment-related R&D can have benefits for the environment, it might also have opportunity costs as it might crowd out other R&D. Using patent data from Delphion and NBER, and financial data from Compustat, Popp and Newell (2012) do not find evidence for such a crowding out effect, neither on the sector level, nor on the firm level. In line with Noailly and Shestalova (2017) they find that patents for alternative energy are cited more often, suggesting that they are more valuable than other patents.

### 3 Data

This section presents the data used in this paper, alongside the variable constructions needed for the analysis. The sample is constructed from several data sources. The four main data sources are discussed in separate sections here. These are the patent data, firm financial data, energy price data, and competition data. Further data sources are mentioned together in a separate section. Variable constructions are discussed throughout these sections. A separate section discusses the aggregation of firm statistics to the firm group level. The last two sections present the final samples and some descriptives of these samples.

### 3.1 Patenting

Patent data is collected from PATSTAT’s 2018 Autumn edition. The PATSTAT data is rich in the number of patent applications, the years covered, the countries covered, and the details per patent application. For example, data on the applicant, technological field and patent office are available.

To provide a measure of innovation output on the firm level, I rely on patent application counts. Throughout, only granted patent applications are considered. To avoid double counting of inventions, one can count patent families instead of individual patents. Patent families group patents referring to the same invention. These can be patent applications at different patent offices or subsequent patent applications to rejected applications. The analysis in this research will be on the family level, unless specified otherwise. I use the DOCDB patent family available in PATSTAT.

Further, patents can be filtered for quality. Applicants might seek broader protection for higher quality ideas. The information on the patent office at which a patent is applied for, offers information on the patent’s quality. Such logic is broadly used in the related literature (see, for example, Aghion et al., 2016; Nicolli & Vona, 2016; Rozendaal & Vollebergh, forthcoming). One could for instance select only patents at the three major patent offices, the European Patent Office (EPO), United States Patent and Trademark Office (USPTO) and the Japan Patent Office (JPO). Even more stringent quality filters would only consider patent families that contain patent applications applied for at multiple patent offices. The Biadic patent family would have patents at both the EPO and USPTO, while the Triadic patent family would have patents at all three major patent offices. While the quality of patents likely increases with these stringent filters, the number of observed inventions decreases. This introduces more zeros in the dependent variable, leading to a more restrictive overview of innovation activities. The methodological approach to the zeros issue will be discussed in Section 4. Alternatively, scholars have used citation-weighting to account for patent quality (see, for example, Popp, 2002). The main drawback of citation weights is that citations take time to accumulate, requiring the researcher to trim the data sample. In order to keep longer time series and to have a more complete overview of innovation activities, I opt for counting patents at any of the three main patent offices.


The patent application’s details on technology classifications are used to specify several classes of relevant technologies. I make use of the international patent classification (IPC) and the cooperative patent classification (CPC) systems to identify the specific technologies. I produce four classes of patent technologies, namely clean, gray, dirty and other. To categorize patents as clean, I utilize the class list of the OECD (2016) that identifies environment-related technologies. Gray and dirty innovation directions are identified through the IPC and CPC classifications listed by Jee and Srivastav (2024), who leverage the works by Aghion et al. (2016), Dechezleprêtre et al. (2019), and IEA (2021). While the dirty energy patents search strategy by the

IEA (2021) is suggested to be accompanied by key word searches, I take the more conservative approach to label all patents with the IEA’s suggested CPC codes as dirty. Added benefit is that the labelling exercise is simplified by skipping the key word search. This shortcut also does not require the patent text data. Any patents that are not labelled as clean, gray or dirty are categorized as other technologies. It can occur that patents contain technology classes that categorize them in more than one of the four classes. In such cases ties are broken in favor of the gray category in order not to inflate either the clean and dirty categories. And in case patents are categorized as both clean and dirty, dirty prevails. Examples of technologies are provided in Table 16 in Appendix C.6.

An alternative strategy to identify clean innovations is to use the Y02 technology classes from the CPC classification. The Y02 class specifically describes technologies for mitigation or adaptation to climate change. This strategy is also widely used in the literature. Due to its simple search strategy, its implementation is quick. Throughout, the clean definition based on the OECD’s search strategy is preferred, as it takes into account a broader set of technology classes, and because it is more in line with the gray and dirty search strategies. Figure 1 presents the aggregated yearly patent counts over time. Three interesting observations can be made. First, there is an increasing trend in all categories. There is a dip in 2014, likely due to the patent publication lag and delays in the inclusion of patents into the PATSTAT database. Patents are only made public once they have received a publication decision by the patent office. As the patent application process can be lengthy, the more recent years in a patent database will suffer from truncation due to the publication lag. Figure 12 in Appendix C.1 provides the distribution of this publication lag. Second, one notices a stronger increase in clean patenting over the other patent classes, which also becomes clear from Figure 13 in Appendix C. Third, the vast majority of patent applications are in the other category. Dirty and clean patent counts have similar sizes to each other, especially in the more recent years. Gray patents are relatively uncommon.

To summarize the above, patent counts can be described by  $PAT_{jt}^{zq}$  for firm  $j$  in year  $t$ , where  $z$  indicates the type of technology and  $q$  indicates the quality. Throughout this paper  $z$  will take the values clean, gray, dirty, other or all. And quality  $q$  consists of several choices discussed above, namely whether only applications at specific offices are considered. For any patent count only granted applications are considered, and counts are on the family level to avoid double counting of innovations. As a patent family can consist of multiple individual patent applications, the technology  $z$  and quality  $q$  are determined by inspecting all patents within the family. If any of the patents in the patent family are in technology  $z$  or have quality  $q$ , the patent family is considered to belong to those classes. Where I do allow for patents to have multiple qualities, as they are never used simultaneously, I do not allow for a patent family to belong to multiple technologies  $z$ . Ties are therefore broken according to the rule discussed earlier, with precedence for gray, than dirty, than clean.





**Figure 1:** CLEAN, GRAY, DIRTY AND OTHER PATENT COUNTS OVER TIME.

*Note:* Patent counts are on the family level. Counts are aggregated within technology class and within year. Only granted patents are considered. And only patents are considered that are applied for by firms covered in this paper. Patents at all patent offices are considered. The other patent category is plotted on the right axis.

Note that a patent application counts towards the patent output of firm  $j$  if firm  $j$  is an applicant to that patent. While this definition could lead to double counting, if a patent has multiple applicants, this is mostly irrelevant for the hypotheses being tested in this paper. The issue of double counting is also limited. Of all the granted patent applications used in this paper, more than 94% have only one applicant (see [Table 13](#)).

Knowledge stocks are another statistic derived from the patent data. They take the above patent counts and accumulate them over time using the perpetual inventory method as in Cockburn and Griliches (1988), such that

$$PS_{jt}^{zq} = PAT_{jt}^{zq} + (1 - \delta)PS_{jt-1}^{zq} \quad (1)$$

where  $PS_{jt}^{zq}$  is the patent stock. The parameter  $\delta$  is the annual depreciation rate. The rationale of a depreciation rate for knowledge is that knowledge becomes obsolete over time. Unless otherwise specified, I will set  $\delta = 0.1$ . Different depreciation rates are used in the patenting literature. While a depreciation rate of 10% is common, see, for example, Verdolini and Galeotti (2011), used rates go up to 15 or 20%, as in Noailly and Smeets (2016) and Aghion et al. (2016), respectively. On the other hand, De Rassenfosse and Jaffe (2018) suggest a depreciation rate as low as 2-7%. For the calculation of the patent stocks, I only consider patents for which the earliest application year is on or after 1975. This is because the other data used in this paper starts only decades later.

Other patent stocks can be derived similarly. Country-level or sector-level patent stocks can be calculated using the same formula as in [Equation 1](#) when replacing the patent count to a country- or sector-level patent count.  $PAT_{ct}$  becomes the number of patent applications by firms located in country  $c$ . Analogously  $PAT_{it}$  is the number of patent applications in sector  $i$ . Also these patent counts are split across the classes



clean, gray, dirty, other and all. Due to the risk of double counting one should note that  $PAT_{ct} \leq \sum_{j \in J_c} PAT_{jt}$ , where  $J_c$  is the set of firms in country  $c$ . This is carefully taken into account when counting patent stocks on the sector and country level. Also for these patent counts only granted patents are considered.

Note that none of the information to identify the firm  $j$ , the country  $c$  and the sector  $i$  are available in PATSTAT. These data come from the Amadeus data that is merged onto the patent data, as discussed further on in this section.

### 3.2 Firm financials and ownership

Firm financial data from Amadeus Financials is used to obtain information on a firm's number of employees, revenue, costs, assets, debt, cash flow, and other accounting data, as well as information on the location of the firm and the sector the firm is active in. The Amadeus Subsidiaries data offers information on ownership links. The information on subsidiary firms is less specific, but it does contain information on the number of employees, ownership shares, sector and location.

The firms in the Amadeus Financials database are only located in Europe. The subsidiaries can be located in other parts of the world. However, due to the location of firms in the Amadeus Financials database, the subsidiaries are also predominantly located in Europe. In fact, nearly 84% of subsidiaries are located in Europe. The firms in these databases are active in similar industries. Studying the NACE industry codes on the two-digit level shows that the most common industries in both databases are wholesale and retail trade, and manufacturing. For the full description of industries in the Amadeus databases I refer to [Table 15](#) in [Appendix C.5](#).

The Amadeus Subsidiaries database is used to construct firm groups. To do so, any firm that owns another firm, but that is itself not owned by another firm, is considered a parent firm. All its subsidiaries and their subsidiaries are then considered part of that parent firm's group. One shortcoming of the Subsidiaries database is that they contain observations at one moment in time.

Several steps are taken to clean the two databases. For the Amadeus Financials data the main objective of cleaning is to harmonize the reported figures. Not all firms report in the same currency and firms might have different reporting periods. The Subsidiaries database does not directly provide a clear overview of firm group structures. Instead it list ownership shares between two firms. A small algorithm is used to unravel the ownership structure. The exact cleaning steps for both databases can be found in [Appendix B](#).

### 3.3 Energy prices

For the energy prices I leverage the work by Sato et al. (2019) who construct average energy prices for different sectors that also vary by country and year. These energy prices are constructed from fuel mixes and the prices of the respective energy types.

This also allows for the construction of shift-share instruments. For these instruments the energy type shares are held constant, such that the instrument only varies with price changes of the different energy types.

The energy price data is harmonized to allow for comparisons across energy types, across currencies and over time. Energy prices are reported in 2010 USD per ton of oil equivalents. Energy prices are inclusive of any taxes. I do not correct energy prices for purchasing power, because energy is globally traded.

### 3.4 Competition

There are several competition measures proposed in the literature (for an overview see OECD, 2021). I use data from the CompNet database that offers competition statistics with variation across countries, industries and year. A host of different competition statistics are provided, but in this paper I will use markups. The markups are derived from fitting a translog production function with OLS. For the exact derivation of the variables I refer to the CompNet user guide (CompNet, 2022).

The markup data from CompNet contains country-industry-year variation. Contributors to the CompNet data have access to each country’s microdata, on the firm level. Those data are used to produce the statistics on the industry-year level for each country. Contributors share these statistics with CompNet. By using the same methods across all national data offices, the CompNet data is comparable across countries.

The CompNet data contains several moments of the competition statistics. In the main specification I use the mean statistic of the markup. This is the mean within country, sector and year. The markup statistic is reported as a multiplier. For example, the interpretation of a markup of 1.2 is a 20% markup.

As economic theory states that larger markups are a sign of lower market competition, it helps to transform the markup variable. I therefore define the competition measure in this paper as

$$COMP_{jt} = -\log(MU_{jt}) \quad (2)$$

where  $MU$  is the markup and  $j$  and  $t$  again denote the firm group and year. The log approximately transforms the multiplier into a share and the negation provides a more intuitive interpretation of the competition measure. A higher value of  $COMP$  can be interpreted as a higher level of market competition. Note that when studying the squared competition variable, the negation has to be restored.

### 3.5 Other data sources

Data on country-level GDP and population levels are obtained from the Worldbank. GDP statistics are harmonized to compare them across countries and over time. GDP is therefore denoted in 2017 USD and corrected for purchasing power parity (PPP).

### 3.6 Firm group aggregation

The ownership structure is of interest, as firm groups often behave as one strategic unit. Imagine a firm group consisting of three firms, a holding firm with headquarters in London, a manufacturing firm located in Berlin, and a research and development firm in Paris. When energy prices in France's manufacturing industry increase, the total energy costs of the firm are little affected, as most energy is used in the manufacturing plant in Berlin. But if energy prices in Germany's manufacturing industry increase, the firm group might increase R&D efforts at the Paris firm, leading to patent applications. Using the location of the firm group's parent firm in London will completely miss any of these dynamics. Therefore, throughout the analysis firm groups will be considered as the main unit of analysis. The locations and industries of all the firms in the firm group will be used to construct patent counts, measures of competition and measures of relevant energy prices.

To construct firm groups, a decision has to be made regarding minimum required ownership shares. I opt for filtering ownership links that have shares of 50% or higher. Minority stakes are thereby excluded, making it more likely that the resulting firm group makes strategic decisions as one unit.

Note that the majority of above mentioned statistics are reported or calculated on the firm level. In order to analyze firm group behavior these statistics need to be aggregated.

For patent counts and patent stocks the aggregations are straightforward. Any patent applied for by the firms in the firm group are considered to be the innovation output of the firm group. Double counting within the firm group is avoided by counting unique patent applications within the firm group.

Some statistics are summed over the firms in the firm group, like the GDP and the population of the countries that a firm group is active in. If the firms of a firm group are active in Germany and France, the GDP and population of these two countries are added up and attached to this firm group. These statistics are also used to calculate the GDP per capita that the firm group is exposed to. Similarly, in order to study spillovers of knowledge, patent stocks on the country level are summed as well. This means that the above firm group is exposed to the combined patent stocks of Germany and France. Averaging these patent stocks across countries would not appreciate the exposure to more knowledge when located in multiple locations.

Other statistics cannot be summed, as they measure an intensity that needs to be comparable across firm groups. This goes for the energy prices and competition statistics. For these variables averages are taken over the firms in a firm group. As simple averages likely assume inappropriate weights, a weighted average is used in the main specifications. The weights are based on the number of employees at the individual unconsolidated firms within the firm group. This brings two challenges, namely there is an additional layer of information needed that could be missing, and the number of employees is endogenous to processes studied in this paper. To tackle

both issues, I resort to a simple solution. I take the average number of employees over the years 2012-2015 for each individual firm. These years contain the fewest missing data, and a static variable is less likely to have endogeneity issues. For the energy price data, about 4.3% of observations is lost when weighting by employees, while about 3.5% of observations is lost for markup data.

Missing data requires caution when constructing firm group-level statistics. Besides the inconvenience of introducing missing values, missing data introduces undesirable composition effects. Take the simple example of a firm group consisting of two firms, of which firm A has data for the years 2000 and 2001 and firm B has data only for 2000. Aggregating within the firm group by simply taking the sum or the average of the variable across firms within a year will result in a compositional difference over time. For the year 2000 the statistic is the aggregate of firm A and B, while for 2001 the aggregate is only based on firm A. As the missing data is likely uninformative, one does not want to introduce such compositional effects in aggregated statistics. Instead, I remove firms from groups that have many missing observations relative to the other firms in the group. Consecutively, I remove the time periods for which not all firms in the group have observations.<sup>1</sup> This assures there are no composition effects in the statistics.

Removing individual firms with few observations from the firm group allows to retain longer time series for the firm groups. It also assures that groups consisting of many firms are not underrepresented. Large groups are more likely to have missing data somewhere in the group. Underrepresenting large firm groups likely affects results, as large groups are expected to be active patent applicants.

### 3.7 Sample and descriptives

The main sample consists of the patenting behavior of firm groups and their exposure to energy prices, competition, and other variables. An alternative sample consists of firm-level statistics. If not specified, the sample on the group-level is used.

The main sample runs from 1995 up to and including 2014 and is unbalanced. Both time boundaries are imposed. 1995 is the first year for which energy data is available. Competition data is available only from 1999 onwards. While there is data on patenting until 2018, the last four years are incomplete due to publication lags. Following Rozendaal (2024) I trim the patent data four years before the last year in the patent data, effectively using data until 2014. The sample further consists of only firms that are linked between the PATSTAT and Amadeus databases. The firm (group) identifiers come from Amadeus and the patenting outcomes come from PATSTAT. This also means that firms that never patent are not represented in this

---

<sup>1</sup>Specifically, three steps are taken. First, any year in a group for which fewer than 65% of firms have observations is dropped. Second, a firm is dropped from the group if it has fewer than 50% of the observations of the group's average. Third, any year for which not all remaining firms in the group have an observation is dropped.

sample, because they will not be found in the PATSTAT data.

Linking the patent data from PATSTAT to the Amadeus firm data is the most challenging part of the sample construction. As these databases have no common identifiers, I link these databases on firm names. This is not straightforward due to alternative and inconsistent spellings across and within databases. Especially the applicant names in the PATSTAT data are poorly spelled. Spelling mistakes are not uncommon and firm names that seem to refer to the same firm have alternative spellings or abbreviations. Patent applicants also do not have consistent identifiers within the PATSTAT database.

I therefore utilize text algorithms to link firm names in a fuzzy matter. This allows for slightly different spellings, increasing the number of links between the databases, and resulting in a more complete view of patenting behavior. The algorithm is described in detail in Bremer (2023), where it is also benchmarked against a simpler exact name match. The fuzzy algorithm vastly outperforms an exact match, by increasing the number of matched firms in the Amadeus Financials database with 116%, and by matching 25.5% of patent applications instead of 3.6% when using an exact match.

While I refer to Bremer (2023) for any details, I quickly mention the matching process here. The algorithm takes four steps. In step 1 the text data, meaning the firm names, is thoroughly cleaned and harmonized. Step 2 turns the text data into numerical vectors based on the text characteristics of the respective firm name and the set of all firm names. It then calculates the geometric angle between each vectorized firm name using the cosine similarity. Step 3 determines which firm names should be considered true matches. To achieve this, a sample of matches is manually labelled as either a correct match or an incorrect match. A probit regression model is fitted on this labelled data, and successively used to out-of-sample predict the matching status of the remaining unlabelled links. Step 4 performs disambiguation. As some patent applicants are linked to multiple firms, this step breaks this duplicity and determines to which firm the applicant should be linked.


Summary statistics of the main variables are presented in Table 1. Notice that the patent statistics have no missing values. Any firm group that does not link to patent applicants that has patent applications in a particular year is assumed to have not patented that year. Further, one should notice that the patent statistics are highly skewed. There are many zeros in the patenting process, especially when filtering for higher quality patents by restricting patent counts to patent applications at the major patent offices. The energy prices and markups are (weighted) averages within the firm group. The weighted variables use the number of employees at the firms within the firm group as the weight. These variables remain denominated in the same way as the simple averages. GDP per capita and spillovers are based on sums within the firm group. As GDP per capita is a ratio, GDP and population are summed within the firm group before calculating the ratio.

**Table 1: SUMMARY STATISTICS.**

	Mean	SD	Min	Median	Max	Obs
Patents clean (all)	0.109	2.46	0.00	0.00	473.00	356,620
Patents dirty (all)	0.126	4.21	0.00	0.00	828.00	356,620
Patents all (all)	2.172	37.68	0.00	0.00	6,202.00	356,620
Patents clean (main)	0.073	1.72	0.00	0.00	274.00	356,620
Patents dirty (main)	0.075	2.71	0.00	0.00	584.00	356,620
Patents all (main)	1.415	24.17	0.00	0.00	2,014.00	356,620
Patents clean (triadic)	0.001	0.05	0.00	0.00	11.00	356,620
Patents dirty (triadic)	0.001	0.04	0.00	0.00	7.00	356,620
Patents all (triadic)	0.022	0.49	0.00	0.00	76.00	356,620
Patent stocks clean (main)	0.438	9.12	0.00	0.00	1,312.60	356,620
Patent stocks dirty (main)	0.474	15.38	0.00	0.00	2,416.80	356,620
Patent stocks all (main)	9.969	172.25	0.00	0.00	14,307.68	356,620
Spillovers country clean (main)	4,836.954	6,893.22	0.00	2,152.53	57,940.63	356,620
Spillovers country dirty (main)	6,245.139	8,729.15	0.00	3,082.89	63,251.56	356,620
Spillovers country all (main)	112,162.596	157,904.95	0.00	49,457.15	1,147,779.68	356,620
Energy prices (2010USD/toe)	775.371	370.85	89.71	693.21	2,805.88	206,644
Energy prices (weighted)	793.958	387.74	89.71	710.86	2,805.88	199,304
Markup	1.293	0.28	0.54	1.23	5.33	129,153
Markup (weighted)	1.265	0.29	0.54	1.19	5.33	125,447
GDP per capita (PPP 2017\$)	36,189.591	12,853.49	2,020.92	40,252.47	120,647.82	356,618

*Note:* All variables are aggregated on the firm group level. All patent statistics are on the DOCDB family level. The term in brackets for the patent variables indicates the office at which patents are applied for, where the main offices are any of USPTO, EPO and JPO, and triadic refers to patent families applied for at all three main offices. Weighted variables refer to the weighted aggregation on the firm group level. They remain denoted in the same denomination as the original variable.

A description of the variation in the energy price data is presented by [Figure 2](#). It plots the energy prices of different firm groups over time. The nine firm groups that are at the 10th to the 90th decile in 1995 are plotted. This figure shows the variation over time, as well as between firm groups. Over the span of these twenty years energy prices have mostly increased, especially since the early 2000s. The distance between the outermost deciles has grown over time with the increase of the energy prices. Lastly, I notice that there are some rank changes over time. For example, the 30th percentile ranks eighth out of nine in 2014. Although in general ranks do not differ drastically over time.




**Figure 2:** ENERGY PRICES OVER TIME.

*Note:* Each line tracks the energy prices of a firm group over time. Nine firm groups are tracked, one for each decile at the start in 1995. As the energy data is widely available already in 1995, I only select firm groups with data for the entire sample of 1995-2014. Energy prices are in 2010USD per ton of oil equivalent. The energy prices are weighted by the individual firm's number of employees.

The same graph is plotted in [Figure 3](#) for the energy price changes. These changes are defined as the log of the first difference of the energy prices and therefore approximate growth rates. This plot shows no clear trend. Furthermore, often the changes co-move across the deciles. There are also clear deviations from these co-movements, for example in 1998 and 2000 for the firm group at the 70th percentile. Compared to the energy prices in levels, the energy price growth rates contain less cross-sectional variation and more time variation.

Similarly, [Figure 4](#) plots the competition deciles over time. Again the firm groups at each decile are tracked over time. Deciles are determined in 2002 due to imperfect data availability before that date. There is no clear trend in the competition data, and there is only limited variation in competition over time. Competition mostly varies across firm groups, while time variation is limited. Some deciles cross each others





**Figure 3:** ENERGY PRICES CHANGES OVER TIME.


*Note:* Each line tracks the log of the first-differenced energy price of a firm group over time. Nine firm groups are tracked, one for each decile at the start in 1996. As the energy data is widely available already in 1995, I only select firm groups with data for the entire sample of 1996-2014 (one lag is lost to construct the first-difference). The energy prices are weighted by the individual firm's number of employees.

paths, but many stay around their original position in the distribution.

A stylized illustration of the inverted-U hypothesis is presented by [Figure 5](#). An inverted-U is apparent for both overall innovation as well as clean innovation innovation. However, the pattern is less pronounced for clean innovation. The maximum for clean innovation can be found close to the upper boundary of competition in this figure. The figure for overall innovation presents a similar story as the one in [Aghion et al. \(2005\)](#).

Similarly, one could illustrate the relationship between energy prices and innovation. [Figure 6](#) presents this relationship for clean and dirty patenting separately. The figure plots the relationships separately for different market competition levels. Firm groups are allocated to a competition tercile based on their average markups. The figure shows that there is a positive relationship between energy prices and both clean and dirty patenting in this stylized setting. Looking at the clean innovation figure, one observes that the relationship is strongest for firm groups operating in markets with medium or high competition. For dirty patenting the relationship is strongest for the middle competition tercile. The other two terciles also have wider confidence intervals. Comparing between clean and dirty patenting I conclude that the relationship is stronger for clean patenting, both economically as well as statistically. Only for the middle tercile the relationship is similar, although the confidence interval is wider for dirty patenting.

Although the stylized relationships in [Figure 5](#) and [6](#) hint towards the existence



**Figure 4:** COMPETITION OVER TIME.

*Note:* Each line tracks the market competition at a firm group over time. The competition statistic is  $-\log(\text{Markup})$ , with a markup of 1.1 referring to a 10% markup. Nine firm groups are tracked, one for each decile at the year 2002. As the competition data is not available for all years, the year 2002 is selected, as it has good coverage from there onwards. I only select firm groups with data for the entire sample of 2002-2014. The markups are weighted by the individual firm's number of employees.


of an inverted-U relationship between competition and innovation, and the existence of induced innovation, these figures provide no evidence of these relationships. In order to provide evidence for these hypotheses, more formal and thorough tests are required. The next section presents the methodology for these tests.

While the Amadeus databases cover many different industries, the availability of energy prices and competition statistics reduce the coverage somewhat. The firm group sample mostly contains firms in the manufacturing, construction and wholesale sectors. An overview of the industries represented by the Amadeus data and the firm group sample is presented in [Table 15](#) in [Appendix C.5](#).

### 3.8 The firm-level sample

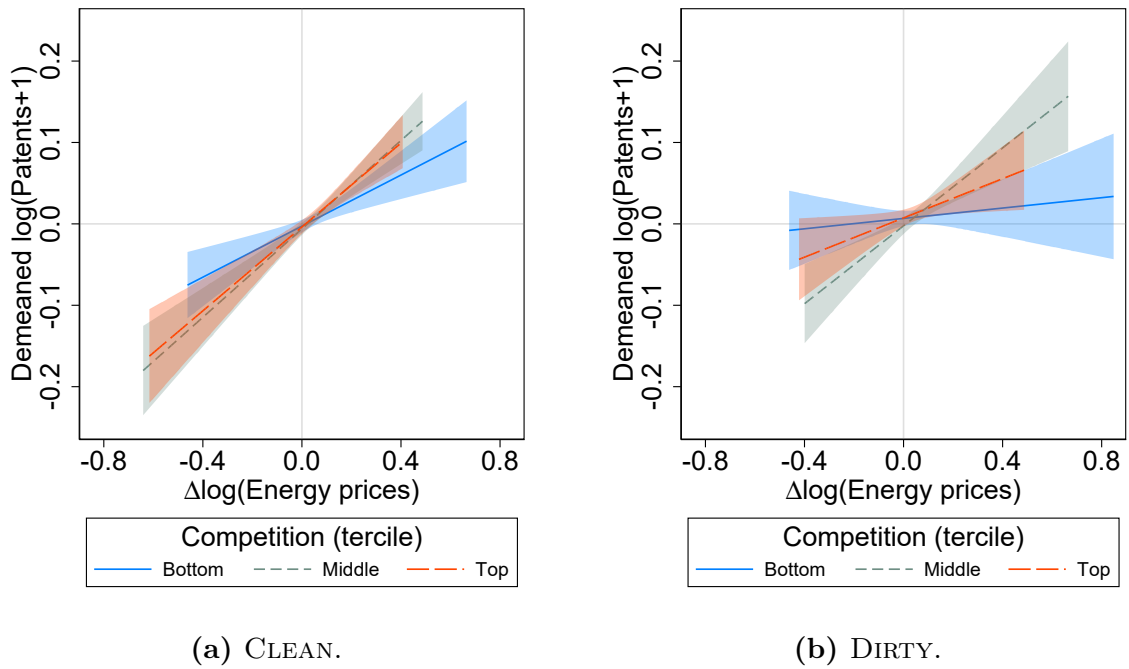
While the firm group is the preferred unit of analysis, one might be interested in studying innovation responses on the individual firm level. Comparing outcomes across these different unit levels might be insightful, especially since the related literature often studies individual firm innovation behavior. I therefore also construct a firm-level sample.

There are two main differences between the firm-level sample and the group-level sample. First, there is no need to aggregate variables across firms. Only the patenting outcomes of the individual firm are considered, and the energy prices, competition and



**Figure 5:** INVERTED-U BETWEEN COMPETITION AND INNOVATION.

*Note:* Each point in this plot represents one firm group’s combination of average competition and total patenting. The log of total patents is taken, thereby consequently removing any zeros. The competition statistic is  $-\log(\overline{markup})$ , with a markup of 1.1 referring to a 10% markup, and the bar denoting the average over time. The average markup data is trimmed to leave only the 10th to the 90th percentile. The separate figures refer to all and clean patenting, respectively. Only patents applied for at any of the main patent offices are counted. The line presents a quadratic fit through the data points.



**Figure 6:** RELATION BETWEEN ENERGY PRICES AND INNOVATION.

*Note:* Each plot presents the linear fit between patenting and energy prices. The shaded areas present 95% confidence intervals. Patents are measured as  $\log(pat + 1)$  and are demeaned on the firm level. Energy prices are in first differences over time and lagged one year. Each plot presents the relationship for the subset of firm groups that on average experiences low, medium or high market competition. The separate figures consider clean and dirty patents, respectively. Only patents applied for at any of the main patent offices are counted. And only firm groups are considered that ever patented in the respective technology class.

GDP per capita variables do not need to be aggregated within the firm group.

Second, the matching between the PATSTAT and Amadeus database is slightly different. While details are provided in Bremer (2023), I here address the most important difference. As the PATSTAT and Amadeus databases are linked with a fuzzy firm name matching algorithm, there is inherent uncertainty around each match. It is well possible that a patent applicant in the PATSTAT database gets matched with more than one firm in the Amadeus database. As the firms in the Amadeus database are unique, at most one match can be the true match. To break the ties, the predicted likelihood,  $\hat{p}$  of the match is used, together with the firm characteristics from the Amadeus database. First, only matches that are within a small distance, namely 0.01, from the highest  $\hat{p}$  of that patent applicant's links are considered. The reason not to select the match with the highest  $\hat{p}$  directly, is to accommodate for numerical imprecision. Second, only firms with unconsolidated data are considered, as for this sample the statistics of the individual firm are of interest. Third, the matches are sorted decreasingly on indicators of firms size. The first of each patent applicant is then considered the true match.

While matching to the firm group level could also lead to duplicates in the proposed matches, the issue is smaller because often proposed matches have duplicates within the firm group. Any within-firm group duplicates do not need to be disambiguated. The chance of introducing mistakes through this disambiguation exercise is larger for the firm-level sample than for the firm group sample.

For the summary statistics of the sample of individual firms, I refer to [Table 14](#) in [Appendix C](#). The main differences with the firm group sample are that the individual firm sample has more observations, as there are more firms than firm groups, and that the patent counts are mostly smaller for the individual firms.<sup>2</sup>

## 4 Methodology

This section discusses the empirical strategy taken in this paper. First, it presents the three hypotheses of interest. Second, it discusses estimation techniques used in related literature. Third, the empirical strategy is presented that should provide answers to the three hypotheses.

### 4.1 Hypotheses

This paper tests the co-occurrence of two hypotheses, namely the induced innovation hypothesis and the hypothesis that competition and innovation have an inverted-U

---

<sup>2</sup>It can occur that the maximum patent counts of the individual firm sample are greater than the maximum of the firm group sample. This results if that particular individual firm does not make it into the firm group sample. This is possible, for example, when the particular individual firm is part of a firm group, but is only owned for a small part. I have ignored small ownership shares in the firm group construction.

relationship. The hypotheses will be tested separately and jointly in a count model regression in which patent counts on the firm-year level are explained by energy prices, competition, and their interactions.

The induced innovation hypothesis states that changes in input prices change the incentives of firms to reduce the usage of that input in their production (see Hicks, 1932). When an input's price increases relative to other inputs' prices, firms want to economize on it. When they decrease, that incentive is no longer there. Applying the hypothesis' logic to energy input prices, one expects firms facing higher energy prices to push in the direction of energy savings. Such savings can come from fuel switching or efficiency. In any case, there is an incentive for directed innovation. The direction of innovation is towards the energy-saving and clean technologies.

Empirically, an inverted-U relationship between market competition and innovation is observed (see Aghion et al., 2005). With regards to market competition and directed innovation, there is a less clear theoretical underpinning for the expected relationship. Whether an inverted-U pattern replicates in the clean innovation domain is to be tested.

With energy input prices and market competition expected to play a role in the innovation decision, the third hypothesis tests how they interact. How does market competition affect the induced innovation channel? This question has no clear underpinning in the existing economic literature, but theoretical works on both hypotheses can be used to start thinking about their interaction. If competition turns out to affect the induced innovation effect, this would be valuable information for policy makers as well. In that case policy makers should take into account how price policies affect clean innovation decisions differently depending on the level of market competition a firm operates in.

In the model by Aghion et al. (2005) there are three relationships that together form the inverted-U between competition and innovation. First, there is a positive relationship between competition and innovation for neck-and-neck firms. This relationship is referred to as the *escape competition effect*. Second, a negative relationship between competition and innovation is described for firms operating in *unlevelled* markets, or markets with *leaders and laggards*. This relationship is referred to as the *Schumpeterian effect*. Third, the endogenous share of sectors that are neck-and-neck, compared to being unlevelled, decreases with competition. Together these three relationships create the inverted-U between competition and innovation. At low levels of competition the increasing escape competition effect is dominant, and at higher levels of competition the decreasing Schumpeterian effect becomes more dominant.

The induced innovation literature describes a positive relationship between input prices and innovations that economize on that given input (Hicks, 1932). Others go further and state that not only directed innovations are induced, but also other innovations (Porter & Van der Linde, 1995a). However, they are talking about environmental policies and not input prices. But if one believes that energy price increases

mimic environmental price policies, energy prices could induce non-clean innovation as well. Having a positive relationship between energy input prices and innovation, the question is whether the intensity, or slope, of this relationship changes with competition.

Returning to the framework of Aghion et al. (2005) one could rephrase that question to whether neck-and-neck competitors and laggards experience different induced innovation incentives. Let us speculate on that question. For example, Porter and Van der Linde (1995a) state that firms that have a culture of change or have managers that have incentives to embrace change are more likely to innovate in response to legislation. It is however the question whether neck-and-neck firms or laggards are more likely to have such cultures or management. On the other hand, some interactions might be clearer. As innovation is costly, only firms in less competitive markets might have the financial means to respond to input price changes with innovation. This channel would hint at induced innovation incentives being strongest in less competitive markets.

With these notes on the interaction between the two hypotheses, I will take the question to the data. With no clear theoretical prior, empirics might shed more light on the matter.

## 4.2 Estimation techniques

The innovation literature has used patent statistics intensively. The literature therefore has also evolved different estimation techniques. Most often the dependent variable of interest is a patent count. As these counts are non-negative and indivisible, linear regression is often deemed inappropriate. Most other patent studies therefore rely on count models.

The most used count model is the Poisson regression, which is estimated using maximum likelihood or a pseudo Poisson maximum likelihood (PPML) as implemented by Correia et al. (2020) (for example Dugoua & Gerarden, 2023). The Poisson model fits the count characteristic of the dependent variable well. The drawback of the Poisson distribution is its equidispersion, as the conditional mean and the conditional variance are the same. As patent counts are often overdispersed, the negative binomial distribution might be preferred (as in Noailly & Shestalova, 2017). Negative binomial regression allows the conditional variance to be greater than the conditional mean of the outcome variable.

As patent counts are characterized by its many zeros, one more extension is sometimes used. Zero-inflated models estimate the model in two steps (see, for example, Noailly & Smeets, 2015). First, the zeros are explained by a set of independent variables in a binary outcomes model, like probit or logit. Second, a count model, like Poisson or negative binomial, is fitted on the non-zero outcomes. This technique thereby explicitly addresses the excess zeros. The downside of zero-inflated models is that its interpretation is also twofold. The model results in parameter estimates for



the binary outcome model as well as parameter estimates for the count model, conditional on the outcome being non-zero. Often these two outcomes are interpreted as the extensive and intensive margin, respectively. The interpretation becomes that the patenting process consists of two decisions, namely (1) does the firm patent in that particular year, and (2) if so, how many patents does it apply for. While zero-inflated models allow for more freedom in fitting the patent outcomes, the question is whether the innovation decision of firms follows these two sequential decisions.

Special attention is paid to the unit fixed effects in estimating count models. The unit fixed effects could be a nuisance to the independent variables in the regression. This is especially the case when the regression specification includes patent stocks. They are likely correlated with the firm fixed effects. Having somewhat short time dimensions in the panel data, makes the non-linear count model estimation vulnerable to the incidental parameter problem. The correlation between the fixed effects and the independent variables could result in inconsistent estimates of the coefficients to the independent variables.

Hausman et al. (1984) suggest to replace the fixed effects with the sum of patent counts over the sample period. While this removes the nuisance of the unit fixed effects, it requires strict exogeneity of the dependent variables. Alternatively, Blundell et al. (1995) suggest to replace the unit fixed effects with pre-sample average patent outcomes, as well as a dummy for when this pre-sample average is non-zero. For both estimators, the estimation can proceed as normal using count models like Poisson or negative binomial.

Aghion et al. (2016) employ both the methods by Hausman et al. (1984) and Blundell et al. (1995, 1999), but additionally develop an alternative that does not rely on long pre-sample data series. Although patent data do have long time series, the question is whether pre-sample averages are proper proxies for unit fixed effects. If the patenting process in the sample period is different from the process in the pre-sample period, the pre-sample averages might not be valid fixed effects. Aghion et al. (2016) point out that this could especially be the case for clean patent applications, which increased substantially in the later years. Their alternative specification therefore does not rely on pre-sample patenting. Instead their method relies on a control function to estimate the unit fixed effects. GMM is then used to simultaneously fit the main regression model as well as the control function equation. An important take-away from their methodological exercises is that the results are qualitatively the same across methods. This even holds when employing OLS with (regular) fixed effects where the dependent variable is the log difference between clean and dirty patents, or specifically  $\log(PAT^{clean} + 1) - \log(PAT^{dirty} + 1)$ . For a more detailed discussion of these different estimators, I refer to their appendix.

As with many studies, there might be endogeneity concerns. These can be addressed with the right instruments. As the regression model is not linear, instrumentation cannot be achieved with two-stage least squares (2SLS). Instead, one could

instrument an endogenous regressor using generalized method of moments (GMM). As the instruments need to be orthogonal to the error term, one could use that to write up the moment conditions.

Alternatively, a control function approach could be taken, as in Dugoua and Gerarden (2023) and explained by Wooldridge (2015). This control function methodology is similar to 2SLS, as it consists of two stages. The first stage is identical to 2SLS. But instead of using the fitted values of the endogenous regressor from the first stage in the second stage, the fitted error term from the first stage enters the second stage, besides the endogenous regressor. The intuition is that the first-stage fitted error term contains the endogenous variation of the endogenous variable. By including it in the second stage it soaks up the endogenous variation and the remaining variation in the endogenous variable should be exogenous. Dugoua and Gerarden (2023) use this control function method in combination with a second-stage Poisson regression which they fit using pseudo-Poisson maximum likelihood. Bootstrapping is used to produce standard errors. The authors do use regular unit fixed effects in their setup and do not control for patent stocks.

For the induced innovation hypothesis energy prices are used in this paper. As these energy prices could be endogenous, they can be instrumented using a *shift-share* or *Bartik* instrument (Bartik, 1991). A shift-share instrument can be constructed for variables that vary due to both composition effects as well as time variation, like an energy price composite consisting of prices of multiple fuels. As the time variation of individual fuel sources is likely exogenous, but the composition of fuels is endogenous, a shift-share instrument could be appropriate. The shift-share instrument keeps the composition constant while allowing for the time variation of the individual components.

Goldsmith-Pinkham et al. (2020) discusses Bartik instruments and how they can instrument endogenous regressors. Noticeable is that the examples consider regressors in growth rates instead of in levels. In the induced innovation literature there is no consensus about the form of the energy prices in the regression. Often studies consider energy price levels (as in Aghion et al., 2016; Dugoua & Gerarden, 2023; Marin & Vona, 2021; Popp, 2002), and some studies consider both levels and first-differences (see, for example, Marin & Vona, 2019). Whereas energy price levels are more common in the patent literature, growth rates are by far most common in the Bartik instruments literature.

Energy price levels and growth rates arguably measure rather different things. Growth rates ignore levels and they thereby cancel much of the cross-sectional variation in energy prices. Instead variation in growth rates likely closer represent shocks, something the Bartik instrument literature tries to isolate and use.

In the context of the induced innovation hypothesis the question is which type of variation one should exploit. If one expects that consistently different energy prices drive innovation, energy price levels make sense. Such identification is more vulnerable

to omitted variables, as both the innovation outcomes as well as the energy prices can be driven by country or sector specific characteristics. If instead one seeks to utilize the arguably more exogenous time variation, ignoring level differences, energy price growth rates are more of interest. Unexpected variation in growth rates are likely a result of exogenous shocks to energy prices.

The empirical patent literature also resorts to lagged regressors in order to strengthen identification. Where contemporaneous regressors are more likely to suffer from reversed causality and omitted variable bias, imposing delayed effects reduces this risk. Additionally, patent applications are the result of a somewhat lengthy innovation process. This process easily takes a year. While a lag structure of one year is common in the patent literature, some studies experiment with longer or multiple lag structures (Dugoua & Gerarden, 2023; Marin & Vona, 2021).

### 4.3 Empirical strategy

Across the hypotheses the dependent variable is the count of successful patent applications. Count models, like Poisson and negative binomial regressions, as discussed in the previous section are therefore appropriate. The main model of focus is the Poisson regression model. In the findings section other models will be fitted as well, like the negative binomial regression model.

In order to test the hypotheses, the patent outcome needs to be explained by energy prices and competition. To uncover the relationships of interest, the following regression equation will be fitted

$$\begin{aligned}
PAT_{jt}^{zq} = \exp & \left( ENPR_{jt-1}\beta_E \right. \\
& + COMP_{jt-1}\beta_C + COMP_{jt-1}^2\beta_{C2} \\
& + ENPR_{jt-1} \times (COMP_{jt-1}\beta_{EC} + COMP_{jt-1}^2\beta_{EC2}) \\
& + \sum_{m \in \mathcal{Z}} PS_{jt-1}^{mq}\beta_P + SPILL_{ct-1}^{mq}\beta_S \\
& + \sum_{m \in \mathcal{Z}} \mathbb{1}_{\{PS_{jt-1}^{mq}=0\}}\kappa_0 + \overline{PAT}_{j0}^{mq}\kappa_1 + \mathbb{1}_{\{\overline{PAT}_{j0}^{mq}=0\}}\kappa_2 \\
& \left. + X_{jt-1}\gamma + \delta_t \right) + \varepsilon_{jt}
\end{aligned} \tag{3}$$

where  $PAT$  are successful patent application counts, and the subscripts  $j$  and  $t$  denote the firm and year, respectively, and the superscripts  $z$  and  $q$  denote the technology class of a patent, for example clean or dirty technologies, and the quality of the patent, for example patent applications at the main patent offices only, respectively.  $ENPR$  are energy prices, and  $COMP$  is a measure of market competition. Furthermore,  $PS$  are the firm's own patent stocks and  $SPILL$  are the patent stocks within the countries the firm is located in, thereby measuring knowledge spillovers. The set  $\mathcal{Z}$  contains the relevant technology classes of the patents, for example clean and dirty technology classes, such that patent stocks of different technology classes enter the

regression separately. This allows testing of knowledge spillovers across technologies.  $X$  is a set of control variables, mostly consisting of only GDP per capita.  $\overline{PAT}_{j0}$  are pre-sample average patent applications, such that  $\kappa_1$  and  $\kappa_2$  capture pre-sample innovation behavior, a proxy for firm fixed effects.  $\kappa_0$  allows firms with zero patent stocks to have a different innovation outcome.  $\delta$  are year fixed effects, and  $\varepsilon$  is the error term. Note that all regressors, except for  $\delta$ , are lagged with one year. By default I will only include firms in the regression that during the time of the sample have at least one patent in the category of the dependent variable. For example, when considering clean patenting at the main patent offices, only firm groups that during the sample period have patented clean innovations at one of the main patent offices are considered. This selection reduces the excess zeros problem and ignores firm groups that have no correlation between the dependent and independent variables.

Note how the regression in [Equation 3](#) forms a log-linear relationship between the dependent and independent variables. This stems from modelling the mean of the Poisson distribution as the exponent of all regressors. The desirable feature of this log-linear relationship is the interpretation of the coefficient estimates, especially when one log-transforms the independent variables. The resulting log-log relationship between dependent and independent variables then results in an elasticity interpretation of the coefficient estimates. All of the independent variables in [Equation 3](#) are log-transformed, except for the binary variables. In order to avoid the introduction of missing observations, the patent stocks,  $PS$  and  $SPILL$ , and the pre-sample average patent outcome,  $\overline{PAT}$ , are incremented with one before the log is taken.

The competition variable requires special attention, as it is measured from markups. The definition of the competition variable  $COMP$  is based on the markup variable  $MU$ , such that  $COMP = -\log(MU)$ . By negating the markup variable the interpretation becomes easier, as one can think of the variable as a measure for competition. One difficulty comes from the squared competition term, as squaring  $COMP$  would cancel out the negative relationship with the markup. Hence the squared competition term is defined as  $COMP^2 = -\log(MU)^2$ .<sup>3</sup>

[Equation 3](#) shows the full regression model used to test the combination of the induced innovation and inverted-U hypotheses. When testing the hypotheses separately, several lines will be left out. To test the induced innovation hypothesis, lines two and three will be disregarded. When testing the inverted-U relationship between competition and innovation, lines one and three will be discarded. The induced innovation hypothesis predicts that  $\beta_E$  is positive when considering  $PAT^{clean}$  as the dependent variable. The inverted-U hypothesis predicts that  $\beta_C$  is positive and that  $\beta_{C2}$  is negative. Additionally the inverted-U hypothesis requires  $\beta_C$  and  $\beta_{C2}$  to be such that the maximum in the relationship between competition and innovation is

---

<sup>3</sup>To think about this in an easier way, one can think of the relationship between markups and innovation. This is hypothesized to have a U relationship. When negating both markup terms, an inverted-U relationship is tested for.

reached at a reasonable level of competition. Assuming  $\beta_C > 0$  and  $\beta_{C2} < 0$ , the maximum is found at  $COMP^{max} = -\frac{\beta_C}{2\beta_{C2}}$ . For the combination of the two hypotheses, the full model is estimated. Due to the interaction terms there are five parameters of interest making the interpretation difficult. To aid interpretation, I will calculate the marginal effects of the energy price on the patent outcome for different values of the competition measure. This provides induced innovation results for different levels of market competition. And as the marginal effects can be interpreted as elasticities, the interpretation is straightforward.

Throughout this paper the energy price variable  $ENPR$  will either be in levels or in first-differences. Given that  $ENPR$  is the log-transformation of the energy price faced by the firm, the first-difference is a close approximation of the growth rate of this energy price. When using first-differences the energy price variable will be denoted  $\Delta ENPR$ . The regression in Equation 3 is further adjusted by excluding year fixed effects when considering  $\Delta ENPR$ , as they likely intervene with the time variation in the energy price growth rates that this specification is set out to exploit. As the growth rates put more emphasis on time variation and less on cross-sectional variation, it is undesirable to have time fixed effects to soak up any of this variation. As discussed in Section 4.2 energy price changes might provide more exogenous variation, aiding the identification strategy.

To further bolster the identification strategy, a shift-share instrument will be used for the energy price. As discussed in Section 3, the energy price data from Sato et al. (2019) comes with a shift-share instrument. The energy price variable is a weighted variable of different prices of energy types, with the weights the actual share of each energy type in the energy mix. The shift-share instrument keeps the composition of energy types, hence the weights, constant at the 1995 level and only allows variation from individual energy types' prices. The instrument therefore rules out any variation due to fuel switching, a decision that could be endogenous to the innovation decision. When using energy price levels, the instrument is also in levels, and when studying first-differenced energy price, the instrument is accordingly transformed to a first difference.

I utilize two instrumental variable approaches, following the discussion in Section 4.2, namely GMM and the control function approach by Dugoua and Gerarden (2023). The GMM approach specifies moment conditions based on the orthogonality of the instrument with respect to the error term from the Poisson regression. The control function approach follows a two-step procedure. For this approach I deviate from Equation 3 and closely follow the specification in Dugoua and Gerarden (2023). In the first step the endogenous energy price variable is regressed on the instrument and the other regressors, or precisely

$$ENPR_{jt-1} = ENPRSS_{jt-1}\beta_{ES} + X_{jt-1} + \tau_{t-1} + \nu_{jt-t} \quad (4)$$

with  $ENPRSS$  the shift-share instrument of the energy price, and  $X$  control variables that also enter the second stage regression. In the main specification  $X$  is identical

to  $X$  in [Equation 3](#) and consists of only GDP per capita. Time fixed effects,  $\tau_t$ , are included only when energy price levels are considered, but not when energy price changes are considered.

The second step consists of the Poisson regression

$$PAT_{jt}^{zq} = \exp \left( ENPR_{jt-1} + \hat{\nu}_{jt-t} + X_{jt-1} + \omega_{t-1} + \varepsilon_{jt} \right) \quad (5)$$

where the variables are the same as in [Equation 3](#),  $\hat{\nu}$  is the fitted error term from the first stage, and  $\omega$  are year fixed effects. Again, year fixed are only included when considering energy price levels and not when considering first differences.

Furthermore, I will analyze the direction of innovation for the induced innovation hypothesis, following one of the robustness specifications of Aghion et al. (2016). The direction of innovation is defined as the log of the ratio of patent outcomes across different technology classes. Such a regression does not have a count outcome variable and hence ordinary least squares (OLS) can be used to fit it. The regression specification becomes

$$\log \left( \frac{PAT_{jt}^{clean,q} + 1}{PAT_{jt}^{dirty,q} + 1} \right) = ENPR_{jt-1}\beta_E + \sum_{m \in \mathcal{Z}} PS_{jt-1}^{mq}\beta_P + SPILL_{ct-1}^{mq}\beta_S \quad (6)$$

$$+ X_{jt-1}\gamma + \eta_j + \delta_t + \varepsilon_{jt}$$

where the variables follow the conventional notation. Variables again enter this regression in logs. As this regression can be estimated with the OLS within estimator, the incidental parameter problem is less likely to cause issues and unit fixed effects,  $\eta_j$ , can be included. There is no need for pre-sample proxies of these fixed effects. Aghion et al. (2016) specify two different regressions, namely one with only firm fixed effects and one that also contains country-year fixed effects.

## 5 Findings

This section presents the main findings of this paper. It separately presents the results to the different hypotheses. First, the results to the induced innovation hypothesis are presented. Second, the results for the inverted-U relationship between competition and innovation are presented. Third, this section presents results for the interaction between the two hypotheses. Fourth, robustness tests are performed and discussed. Lastly, a discussion wraps up this section.

Throughout this section the unit of analysis is the firm group. Firm-level results are discussed in the robustness section.

### 5.1 Induced innovation

When fitting [Equation 3](#) using a Poisson regression without the competition terms, the induced innovation hypothesis is tested. [Table 2](#) presents the results for the regressions

based on energy price levels. The different regressions consider different technology classes and different quality filters on the patent applications. For each regression the sample is restricted to firm groups that at least once patented in the category of the dependent variable over the duration of the sample. For example, when studying clean patenting at the main patent offices, only firm groups are studied that during the 1995-2014 have at least once patented in this category. Looking at the coefficient estimates to the energy price variable, there is no evidence that energy prices induce clean innovation, as shown in the first two columns. Both coefficients are negative, which goes counter to the induced innovation hypothesis. However, only the energy price coefficient in the specification with patent applications at all patent offices is marginally statistically significant. Interestingly, dirty patenting at all patent offices does respond negatively to energy price changes. This coefficient estimate is greater in absolute terms than the energy price coefficient to the clean innovation outcome. The estimate is also statistically significant at the 1% level. However, this effect does not hold for the applications at the main patent offices, as the negative coefficient is not statistically significantly different from zero. And while not statistically significant, in the overall patenting regressions in columns 5 and 6, the energy price coefficient has a positive sign. As these estimates follow from a log-log relationship, the interpretation of the energy price coefficient in column 1, for example, is that a one percent increase in energy prices results in a 0.18% decrease in clean patent applications at all patent offices.

Further, the results clearly show the positive relationship between the firm group's own patent stocks and its successful patent applications. There is also clear path dependence, where clean patent stocks help further clean innovations, and where dirty patent stocks help further dirty innovations. Noticeable is also how dirty patent stocks aid each type of innovation, even clean innovation. This relationship does not hold the other way around, as clean patent stocks have no relationship with dirty patent applications. Interesting is also to compare these coefficients with the coefficient of the stock of all patents. Dirty patent stocks are not more valuable than the stock of all patents when it comes to clean innovation.

Spillovers within countries are statistically significant in a few regressions only. And perhaps surprisingly, these spillovers do not always affect the patenting within the technology types. Instead, clean patent stocks aid dirty patenting, and dirty patent stocks harm clean patenting outcomes. When it comes to all types of innovations, in columns 5 and 6, the country spillovers seem to play no role. These findings do not paint a picture of technology-related spillovers on the country level. Looking at the last independent variable in the model, GDP per capita has a negative coefficient estimate, but the estimates are only statistically significant for overall patenting outcomes at the main patent offices. This finding is somewhat surprising, but could be driven by the unit of analysis. If firm groups with fewer subsidiaries are disproportionately located in richer countries, and patent less than firm groups located in many countries,



**Table 2:** THE INDUCED INNOVATION HYPOTHESIS (POISSON, ENERGY PRICE LEVELS).

	Clean		Dirty		All	
	All	Main	All	Main	All	Main
Energy prices	-0.175*	-0.140	-0.298***	-0.097	0.061	0.051
	(0.097)	(0.110)	(0.103)	(0.065)	(0.099)	(0.094)
Patent stocks clean	1.022***	1.002***	-0.014	-0.024	-0.026	-0.052**
	(0.055)	(0.054)	(0.028)	(0.032)	(0.021)	(0.026)
Patent stocks dirty	0.102***	0.131***	1.174***	1.162***	0.061***	0.060***
	(0.027)	(0.031)	(0.036)	(0.036)	(0.018)	(0.017)
Patent stocks all	0.152***	0.145***	0.196***	0.153***	1.117***	1.157***
	(0.037)	(0.040)	(0.041)	(0.032)	(0.015)	(0.018)
Spillovers country clean	0.102	0.128	0.373*	0.554***	-0.024	0.049
	(0.186)	(0.239)	(0.209)	(0.210)	(0.131)	(0.149)
Spillovers country dirty	-0.170*	-0.291***	0.166	0.101	0.143	0.030
	(0.094)	(0.110)	(0.145)	(0.135)	(0.089)	(0.079)
Spillovers country all	0.116	0.215	-0.465*	-0.592**	-0.070	-0.046
	(0.177)	(0.217)	(0.274)	(0.245)	(0.155)	(0.168)
GDP per capita	-0.063	-0.085	-0.058	-0.095*	-0.028	-0.089***
	(0.043)	(0.052)	(0.044)	(0.052)	(0.035)	(0.033)
Constant	-1.675	-1.983*	0.120	-0.082	-2.560***	-1.885**
	(1.069)	(1.157)	(0.970)	(0.759)	(0.738)	(0.768)
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Pre-sample controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	43,412	29,115	26,690	18,497	190,012	132,144
Unique firm groups	2,417	1,633	1,504	1,053	10,341	7,213
Pseudo $R^2$	0.609	0.614	0.837	0.829	0.855	0.855

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* Pseudo-Poisson maximum-likelihood estimation results. The dependent variable is the count of successful patent applications on the family level. Headers indicate the type of technology and the patent office. Independent patent variables are also applied for at the patent office indicated by the header. Each regression includes pre-sample controls as specified in Equation 3. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm group level.

this negative relationship between a firm groups weighted average GDP per capita exposure and patenting outcomes could occur. As all presented independent variables in the regression are log-transformed the coefficients have an elasticity interpretation, similar to the energy price coefficient. The magnitudes are therefore comparable across variables. This clearly shows that the firm’s own patent stocks are the main driver of further patenting.

The Poisson results from the specification with energy price changes are presented in [Table 3](#). These results are noticeably different, and more in line with the induced innovation hypothesis. The energy price coefficient estimates are positive for clean patenting and negative for dirty patenting. Furthermore, the coefficient is negative for overall patenting as well. This indicates that firms do respond to energy price changes with more clean and less dirty and patenting. The negative effects on overall patenting suggest that innovation activities might be crowded out by higher energy prices. It could be that higher energy prices squeeze the innovation budget. While the signs of the coefficient estimates are interesting, only two energy price coefficient estimates are statistically significant, namely for the regression on clean patenting at the main patent offices and dirty patenting at all offices.

The coefficient estimates for the firm group’s own patent stocks are highly similar to the Poisson regression specification with energy price levels. The spillover results, however, are different. The clean patent stocks in the countries the firm group is active in reduce the number of patent applications in all technology classes. This effect is strongest for the regression on all patent classes. Dirty spillovers are more in line with the results from the energy levels specification, as they aid dirty patenting. They also aid overall patenting, but they do not affect clean patenting outcomes. Overall country spillovers affect all types of patent outcomes positively, although this effect is statistically insignificant for dirty patent applications at the main patent offices. This differs from the previous findings where overall spillovers negatively impacted dirty patent outcomes, but did not affect other patenting outcomes.

Both specifications can also be fitted with a negative binomial regression model, allowing for overdispersion of the patenting outcomes. The results for patent applications at the main patent offices are presented in [Table 4](#). The results are much in line with the findings in [Table 2](#) and [3](#). The specification with energy price levels differs slightly for overall patenting, where the negative coefficient estimate for energy prices is now marginally statistically significant. The sign for the specification with dirty patenting becomes positive in the negative binomial specification, but the estimate remains statistically indistinguishable from zero. The specification with energy price changes yields highly similar estimates. All signs are the same to those from the Poisson estimates. In this specification energy prices induce clean innovation. Dirty innovation is reduced, but the estimate remains statistically insignificant. Overall patenting is again reduced, and now is statistically significant at the 5% level. The results for all patent offices are presented in [Table 17](#) in [Appendix D](#). Also there the

**Table 3:** THE INDUCED INNOVATION HYPOTHESIS (POISSON, ENERGY PRICE CHANGES).

	Clean		Dirty		All	
	All	Main	All	Main	All	Main
$\Delta$ Energy prices	0.267 (0.219)	0.567** (0.265)	-0.523*** (0.142)	-0.164 (0.207)	-0.226 (0.325)	-0.309 (0.198)
Patent stocks clean	1.014*** (0.059)	0.993*** (0.059)	0.005 (0.036)	0.025 (0.035)	-0.043* (0.026)	-0.042 (0.026)
Patent stocks dirty	0.103*** (0.028)	0.139*** (0.032)	1.080*** (0.048)	1.070*** (0.048)	0.054*** (0.016)	0.048*** (0.018)
Patent stocks all	0.149*** (0.038)	0.137*** (0.043)	0.170*** (0.044)	0.099*** (0.031)	1.103*** (0.011)	1.115*** (0.018)
Spillovers country clean	-0.562*** (0.172)	-0.770*** (0.225)	-0.826*** (0.314)	-0.574*** (0.137)	-1.111*** (0.154)	-1.005*** (0.138)
Spillovers country dirty	-0.054 (0.073)	-0.077 (0.083)	0.219* (0.122)	0.221** (0.110)	0.241*** (0.077)	0.162** (0.070)
Spillovers country all	0.647*** (0.160)	0.875*** (0.204)	0.570** (0.254)	0.323 (0.203)	0.859*** (0.132)	0.815*** (0.132)
GDP per capita	-0.081* (0.042)	-0.123** (0.051)	-0.148*** (0.058)	-0.198*** (0.063)	-0.105*** (0.032)	-0.175*** (0.036)
Constant	-4.219*** (0.731)	-4.490*** (0.872)	-2.911** (1.162)	-1.453 (1.183)	-3.840*** (0.612)	-2.840*** (0.634)
Year FEs	No	No	No	No	No	No
Pre-sample controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	40,798	27,327	25,038	17,323	179,208	124,523
Unique firm groups	2,417	1,633	1,504	1,053	10,341	7,213
Pseudo $R^2$	0.600	0.603	0.831	0.823	0.849	0.847

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* Pseudo-Poisson maximum-likelihood estimation results. The dependent variable is the count of successful patent applications on the family level. Headers indicate the type of technology and the patent office. Independent patent variables are also applied for at the patent office indicated by the header. Each regression includes pre-sample controls as specified in Equation 3. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm group level.

negative binomial results are similar to the Poisson regression results.

Again there is strong path dependence from the firm’s own patent stocks in the negative binomial results. All coefficient estimates are positive, but there is clear path dependence in the technology direction. For all specifications the coefficient estimate of the firm’s own patent stock in the same technology is around 1.2 and is statistically significant. The spillover estimates are again somewhat different, indication that they are not robust against changes in regression specifications and estimation techniques.

As the energy price could be endogenous, the shift-share instrument supplied with the energy prices database from Sato et al. (2019), and discussed in Section 4.3, could be used to instrument the energy prices. Whereas the energy price variable changes with both the energy price of each fuel type and the composition of fuels used in each country-sector-year combination, the instrument keeps the composition fixed at the 1995 level. In line with the Bartik instrument literature, I focus on the regression specification with energy price changes.

Table 5 presents the results. It provides estimates from a GMM specification that assumes a Poisson distribution, and the two-step control function specification detailed in Equation 4 and 5 which is based on Dugoua and Gerarden (2023). The table presents the second stage only, which is estimated using pseudo-Poisson maximum likelihood and includes the first-stage’s fitted residuals. The results are highly similar between the GMM and control function specifications for the clean and dirty regression specifications. The relationship between energy prices and clean patenting is positive and statistically significant. The relationship with dirty and overall patenting is statistically insignificant. The size of the coefficient is also largest for the clean specification. The only difference between the GMM and control function estimates is that the coefficient estimate of energy prices in the overall patenting specification is positive for the control function method and negative for the GMM method. But both estimates are statistically insignificant. Note that for all control function specifications the first-stage  $F$ -statistic is rather high, suggesting that the instrument explains sufficient variation in the endogenous energy prices. Compared to the Poisson results from Table 3 the energy price coefficients are now positive instead of negative in the dirty patenting specifications. But none of these coefficient estimates are statistically different from zero. The instrumental variable results for the specification with energy price levels can be found in Table 18. These results yield no statistically significant coefficient estimates for the energy prices.

Instead of instrumenting the energy prices, one could use the shift-share instrument directly. As the instrument arguably does not suffer from endogeneity, it provides exogenous variation that can be studied directly in a reduced form specification. The results, presented in Table 19 in Appendix D, are similar to the IV results of Table 5. While some estimates differ across estimation methods, these differences are often not statistically significant. Overall the conclusion is therefore that there are few differences in the results across methodologies. Whether a Poisson distribution is

**Table 4:** THE INDUCED INNOVATION HYPOTHESIS (NEGATIVE BINOMIAL).

	Levels			Changes		
	Clean	Dirty	All	Clean	Dirty	All
Energy prices	-0.092 (0.059)	0.009 (0.082)	0.046* (0.025)			
$\Delta$ Energy prices				0.620*** (0.213)	-0.312 (0.229)	-0.158** (0.078)
Patent stocks clean	1.148*** (0.033)	0.031 (0.030)	0.013 (0.016)	1.152*** (0.034)	0.030 (0.032)	-0.001 (0.016)
Patent stocks dirty	0.087*** (0.020)	1.230*** (0.038)	0.029** (0.014)	0.092*** (0.021)	1.243*** (0.038)	0.030** (0.014)
Patent stocks all	0.280*** (0.022)	0.198*** (0.027)	1.246*** (0.009)	0.279*** (0.022)	0.190*** (0.027)	1.238*** (0.009)
Spillovers country clean	0.099 (0.115)	0.087 (0.163)	0.105** (0.051)	-0.251** (0.103)	-0.432*** (0.141)	-0.588*** (0.041)
Spillovers country dirty	-0.124** (0.062)	0.148* (0.087)	-0.049* (0.029)	-0.061 (0.061)	0.185** (0.086)	0.054* (0.027)
Spillovers country all	0.058 (0.115)	-0.213 (0.174)	0.037 (0.051)	0.345*** (0.105)	0.252* (0.150)	0.591*** (0.043)
GDP per capita	-0.034 (0.035)	-0.070* (0.037)	-0.074*** (0.013)	-0.037 (0.035)	-0.094** (0.039)	-0.131*** (0.014)
Constant	-2.737*** (0.627)	-1.951** (0.830)	-2.607*** (0.257)	-4.318*** (0.580)	-3.062*** (0.711)	-3.520*** (0.235)
Year FEs	Yes	Yes	Yes	No	No	No
Pre-sample controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	29,115	18,497	132,144	27,327	17,323	124,523
Unique firm groups	1,633	1,053	7,213	1,633	1,053	7,213
Pseudo $R^2$	0.236	0.272	0.267	0.230	0.266	0.262

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* Negative binomial maximum-likelihood estimation results. The dependent variable is the count of successful patent applications on the family level. Headers indicate the regression specification and the type of technology. Only patents applied for at the main patent offices are considered. Each regression includes pre-sample controls. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm group level.

**Table 5:** THE INDUCED INNOVATION HYPOTHESIS (INSTRUMENTED ENERGY PRICES).

	GMM			Control function		
	Clean	Dirty	All	Clean	Dirty	All
$\Delta$ Energy prices	0.889** (0.423)	0.165 (0.274)	-0.202 (0.229)	0.789** (0.351)	0.103 (0.384)	0.134 (0.101)
First-stage residual				0.598 (0.425)	-0.899 (0.606)	-0.128 (0.251)
Patent stocks clean	0.991*** (0.060)	0.018 (0.036)	-0.043 (0.027)			
Patent stocks dirty	0.139*** (0.032)	1.074*** (0.046)	0.048*** (0.018)			
Patent stocks all	0.137*** (0.043)	0.102*** (0.031)	1.116*** (0.018)			
Spillovers country clean	-0.757*** (0.218)	-0.542*** (0.143)	-1.001*** (0.137)			
Spillovers country dirty	-0.078 (0.083)	0.212** (0.105)	0.162** (0.070)			
Spillovers country all	0.863*** (0.199)	0.297 (0.207)	0.810*** (0.130)			
GDP per capita	-0.126** (0.052)	-0.205*** (0.063)	-0.176*** (0.036)	2.952*** (0.608)	2.482*** (0.594)	1.297*** (0.340)
Constant	-4.422*** (0.869)	-1.297 (1.209)	-2.814*** (0.625)	-28.208*** (6.076)	-22.213*** (6.002)	-8.874** (3.632)
Year FEs	No	No	No	No	No	No
Firm FEs	No	No	No	Yes	Yes	Yes
Pre-sample controls	Yes	Yes	Yes	No	No	No
Observations	27,327	17,323	124,523	25,443	15,876	119,057
Unique firm groups	1,633	1,053	7,213	1,503	956	6,871
Pseudo $R^2$				0.643	0.841	0.867
First-stage F-stat				19.68	62.60	73.13

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* GMM and control function results. The GMM specification assumes a Poisson distribution. The control function specification is in line with Dugoua and Gerarden (2023) and includes the first-stage error term in a second-stage pseudo-Poisson maximum likelihood model. The dependent variable is the count of successful patent applications on the family level. Headers indicate the estimation method and the type of technology. Only patents applied for at the main patent offices are considered. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm group level. For the control function regressions the standard errors are bootstrapped.

used or a negative binomial distribution, or whether energy prices are instrumented or not does not seem to matter much for testing the induced innovation hypothesis.

Mind that each regression above is run on the sample of firms that has at least one patent of the category of the dependent variable. For example, when studying the firm group's clean patent applications at the main offices, only firm groups are selected that during the sample period have at least one successful clean patent application at one of the main offices. This leads to different samples for each regression specification. The differences in coefficient estimates are therefore not purely driven by differences in dependent variable outcomes, but also in differences in the composition of firm groups in the sample. One might wonder how the results look if this composition effect is filtered out. By restricting the sample to firm groups that have both at least one clean and one dirty patent application during the sample period, the sample becomes the same across regression specifications.

Poisson regression results of this exercise are presented in [Table 6](#) for patent applications at the main patent offices. These results are highly similar to the results with compositional changes from [Table 2](#) and [3](#). Perhaps the main difference is that the energy price coefficient in the specification with energy price changes and clean patent applications becomes statistically significant at the 10% level instead of the 5% level. As the samples are the same across specifications, these results allow drawing conclusions on the direction of innovation that firms take. When considering the specification with energy price levels, the conclusion is that firms reduce both clean and dirty patenting when energy prices are higher, but that dirty patenting is reduced more. Overall patenting is unaffected. Note that the energy price coefficients are likely not statistically significantly different from one another. When considering the specification with energy price changes, the conclusion is that firms respond to energy price increases with more clean patenting and less dirty patenting. Overall patenting is also reduced. Here only the coefficient in the clean patenting specification is marginally statistically significant. While I do not explicitly test for it, the energy price coefficient in the clean specification is most likely statistically significantly different from the coefficient in the dirty and overall specification.

Lastly, one could study the within-firm choices in the direction of innovation, as detailed in [Equation 6](#). The dependent variable in this specification no longer reflects the magnitude of the innovation activities but rather the technological direction taken in the innovation decisions. [Table 7](#) presents the OLS results for this specification, where energy prices are instrumented with the shift-share instrument. Different specifications are provided, where the first column is in line with one of the specifications in [Aghion et al. \(2016\)](#). The results for the energy price coefficients are characterized by statistical insignificance when studying the clean versus dirty direction. In general the energy prices steer in the direction of relatively more clean patenting, compared to dirty patenting. This does not hold for the most elaborate levels specification, where the coefficient estimate is negative. Studying the ratio between clean and all



**Table 6:** THE INDUCED INNOVATION HYPOTHESIS FOR THE SAME SAMPLES (POISSON).

	Levels			Changes		
	Clean	Dirty	All	Clean	Dirty	All
Energy prices	-0.103 (0.138)	-0.126* (0.070)	0.003 (0.153)			
$\Delta$ Energy prices				0.601* (0.321)	-0.218 (0.226)	-0.394 (0.251)
Patent stocks clean	1.018*** (0.059)	-0.024 (0.033)	-0.044 (0.031)	1.004*** (0.067)	0.014 (0.035)	-0.030 (0.033)
Patent stocks dirty	0.127*** (0.033)	1.175*** (0.035)	0.046** (0.022)	0.143*** (0.036)	1.096*** (0.040)	0.031 (0.023)
Patent stocks all	0.054 (0.050)	0.081** (0.037)	1.098*** (0.029)	0.046 (0.056)	0.001 (0.036)	1.045*** (0.029)
Spillovers country clean	0.259 (0.359)	0.381* (0.227)	-0.046 (0.244)	-0.995*** (0.314)	-0.626*** (0.172)	-1.197*** (0.239)
Spillovers country dirty	-0.402*** (0.148)	0.112 (0.143)	0.052 (0.121)	-0.091 (0.104)	0.224** (0.106)	0.233** (0.117)
Spillovers country all	0.218 (0.316)	-0.450* (0.253)	0.003 (0.247)	1.125*** (0.279)	0.343 (0.227)	0.902*** (0.240)
GDP per capita	-0.056 (0.061)	-0.075 (0.060)	-0.112** (0.055)	-0.100 (0.061)	-0.182** (0.072)	-0.209*** (0.064)
Constant	-2.343* (1.417)	-0.083 (0.821)	-0.934 (1.033)	-5.245*** (1.129)	-1.045 (1.340)	-2.076* (1.228)
Year FEs	Yes	Yes	Yes	No	No	No
Pre-sample controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	9,674	9,674	9,674	9,025	9,025	9,025
Unique firm groups	565	565	565	565	565	565
Pseudo $R^2$	0.667	0.850	0.881	0.651	0.844	0.870

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* Pseudo-Poisson maximum-likelihood estimation results. The dependent variable is the count of successful patent applications on the family level. Headers indicate the regression specification and the type of technology. Only patents applied for at the main patent offices are considered. Each regression includes pre-sample controls as specified in [Equation 3](#). Only firms that patented in both the clean and dirty technology category during the time period of the sample are included, forcing the same samples across regressions. Standard errors are clustered at the firm group level.

patenting, the results are positive and larger. The energy price coefficient estimate in the levels specification is also statistically significant at the 5% level.

## 5.2 Competition and innovation

To test the relationship between competition and innovation and to explicitly test for an inverted-U relationship, the competition coefficients from [Equation 3](#) need to be fitted. An inverted-U relationship appears if the coefficient to the competition term is positive and the coefficient to the quadratic term is negative.

[Table 8](#) presents the results from a Poisson regression. Across all specifications an inverted-U relationship between competition and innovation appears. But the competition coefficients are statistically insignificant for clean and dirty patenting. For overall patenting the coefficients are marginally statistically significant when considering patenting at the main offices only. Testing the joint significance of the two competition terms with a Wald test results in the  $p$ -values reported in the bottom of the table. For the two regressions on clean innovation the null hypothesis of the test is rejected, meaning that the competition coefficients are not both zero. For the other regressions the null hypothesis cannot be rejected.

The bottom of the table also reports the markup at which the elasticity between patenting and competition is at its maximum. Estimates range from 1.23 to 1.51 for all patent offices and 1.57 to 2.59 for the main patent offices. Remember the markup is measured as a multiplier, so values are generally greater than one. To put the numbers in perspective, the 10, 50th and 90th percentile for the entire sample are 1.08, 1.19 and 1.52, respectively. This means that the positive effects from competition peak towards the very end of the markup distribution and hence at the start of the competition distribution when considering patenting at the main patent offices only. Innovation is according to these results at its maximum under extremely low levels of market competition. Note that there is substantial statistical uncertainty around these optimal competition numbers due to the statistical uncertainty around the competition coefficient estimates. When considering all patent offices the range of maximums corresponds to the 11-36th percentile range of the competition distribution. It therefore still indicates that competition maximally induces innovation around the first quartile of competition.

Results for a specification without year fixed effects are presented in [Table 21](#) in [Appendix D](#). One might leave out year fixed effects to exploit the time variation in the regressors more. As year fixed effects could correlate with the independent variables, they could obscure the relationship between the independent and the dependent variables. This could lead to the incidental parameter problem discussed in [Section 4](#). On the other hand, one might want to exploit cross-sectional variation more than time variation when it comes to the competition measure. The main message is that the competition coefficient estimates are qualitatively the same. The specification without year fixed effects results in more statistical precision for most coefficient estimates,

**Table 7:** THE INDUCED DIRECTION OF INNOVATION (OLS).

	Clean vs dirty				Clean vs all	
	Levels	Levels	Changes	Changes	Levels	Changes
Energy prices	0.065 (0.059)	-0.031 (0.078)			0.181** (0.087)	
$\Delta$ Energy prices			0.031 (0.076)	0.047 (0.077)		0.129 (0.089)
Patent stocks clean	0.160*** (0.031)	0.151*** (0.032)	0.134*** (0.033)	0.127*** (0.033)	0.164*** (0.034)	0.150*** (0.036)
Patent stocks dirty	-0.100*** (0.033)	-0.108*** (0.035)	-0.064* (0.034)	-0.070* (0.037)	0.063* (0.032)	0.061* (0.034)
Patent stocks all		0.018 (0.019)		0.012 (0.021)	-0.461*** (0.028)	-0.440*** (0.030)
Spillovers country clean	0.094 (0.077)	0.248* (0.140)	0.138** (0.070)	0.156** (0.077)	0.528*** (0.120)	0.646*** (0.104)
Spillovers country dirty	-0.136* (0.074)	-0.086 (0.097)	-0.136* (0.082)	-0.102 (0.110)	-0.307** (0.128)	-0.282** (0.135)
Spillovers country all		-0.062 (0.175)		-0.040 (0.123)	-0.168 (0.153)	-0.210 (0.151)
GDP per capita		-0.102 (0.114)		-0.096 (0.111)	0.084 (0.149)	0.103 (0.154)
Constant	-0.063 (0.235)	0.571 (2.026)	-0.036 (0.239)	0.925 (0.967)	-1.911 (1.255)	-1.756 (1.343)
Year FEs	No	Yes	No	No	No	No
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes
Pre-sample controls	No	No	No	No	No	No
Observations	9,674	9,674	9,025	9,025	9,674	9,025
Unique firm groups	565	565	565	565	565	565
R <sup>2</sup> (within)	0.023	0.027	0.018	0.018	0.164	0.132
First-stage F-stat	6.89	6.88	6.74	6.73	8.47	8.46

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* OLS regressions with the energy prices instrumented with a shift-share instrument. The dependent variable is the log difference between the count of two technology types, e.g. clean and dirty, as indicated by the table header. Patent counts are incremented with one before taking logs. Only successful patent applications on the family level are considered. Headers indicate the regression specification and the comparison between technology types. Only patents applied for at the main patent offices are considered. Only firms that patented in both the clean and dirty technology category during the time period of the sample are included, forcing the same samples across regressions. Standard errors are clustered at the firm group level.

**Table 8:** TESTING THE INVERTED-U HYPOTHESIS (POISSON).

	All offices			Main offices		
	Clean	Dirty	All	Clean	Dirty	All
Competition (Markup)	0.194 (0.416)	0.410 (0.435)	0.174 (0.273)	0.572 (0.580)	0.458 (0.516)	0.643* (0.357)
Competition <sup>2</sup>	-0.465 (0.329)	-0.498 (0.384)	-0.300 (0.271)	-0.633 (0.457)	-0.241 (0.484)	-0.604* (0.327)
Patent stocks clean	1.000*** (0.035)	0.026 (0.039)	-0.011 (0.032)	0.982*** (0.036)	0.023 (0.050)	-0.042 (0.038)
Patent stocks dirty	0.098*** (0.034)	1.216*** (0.042)	0.072*** (0.018)	0.130*** (0.040)	1.221*** (0.049)	0.072*** (0.019)
Patent stocks all	0.145*** (0.045)	0.078* (0.044)	1.076*** (0.017)	0.138** (0.057)	0.053 (0.050)	1.119*** (0.023)
Spillovers country clean	0.132 (0.193)	0.731** (0.355)	0.352** (0.165)	0.053 (0.222)	0.906** (0.365)	0.354* (0.196)
Spillovers country dirty	-0.051 (0.103)	0.066 (0.143)	0.048 (0.097)	-0.096 (0.159)	0.003 (0.173)	-0.045 (0.144)
Spillovers country all	-0.084 (0.239)	-0.750** (0.350)	-0.375 (0.241)	0.044 (0.321)	-0.845** (0.406)	-0.266 (0.312)
GDP per capita	-0.159*** (0.059)	-0.045 (0.067)	-0.052 (0.052)	-0.159** (0.079)	-0.048 (0.073)	-0.056 (0.067)
Constant	-0.974 (1.157)	-0.790 (0.792)	-0.902 (0.942)	-1.334 (1.640)	-0.448 (1.055)	-1.455 (1.335)
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Pre-sample controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	22,279	13,085	112,152	13,628	8,330	72,117
Unique firm groups	2,844	1,680	13,865	1,796	1,111	9,018
Pseudo $R^2$	0.685	0.836	0.848	0.714	0.845	0.855
Markup at min/max	1.23	1.51	1.34	1.57	2.59	1.70
Joint test competition ( $p$ )	0.000	0.168	0.362	0.045	0.347	0.177

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* Pseudo-Poisson maximum-likelihood estimation results. The dependent variable is the count of successful patent applications on the family level. Headers indicate the type of technology and the considered patent offices. Independent patent variables are also applied for at the patent office indicated by the header. Each regression includes pre-sample controls. Only firms that patented in the category of the dependent variable during the time period of the sample are included. The joint test presents the  $p$ -value from a Wald test with the null hypothesis stating that both competition terms are zero. Standard errors are clustered at the firm group level.

and especially for the squared competition term.

### 5.3 Induced innovation and competitive pressures

The interaction between the induced innovation hypothesis and the inverted-U between competition and innovation can be tested by estimating the full model from [Equation 3](#). As the results contain multiple interactions, marginal effects are produced to test the effects of energy prices on patenting under different levels of market competition.

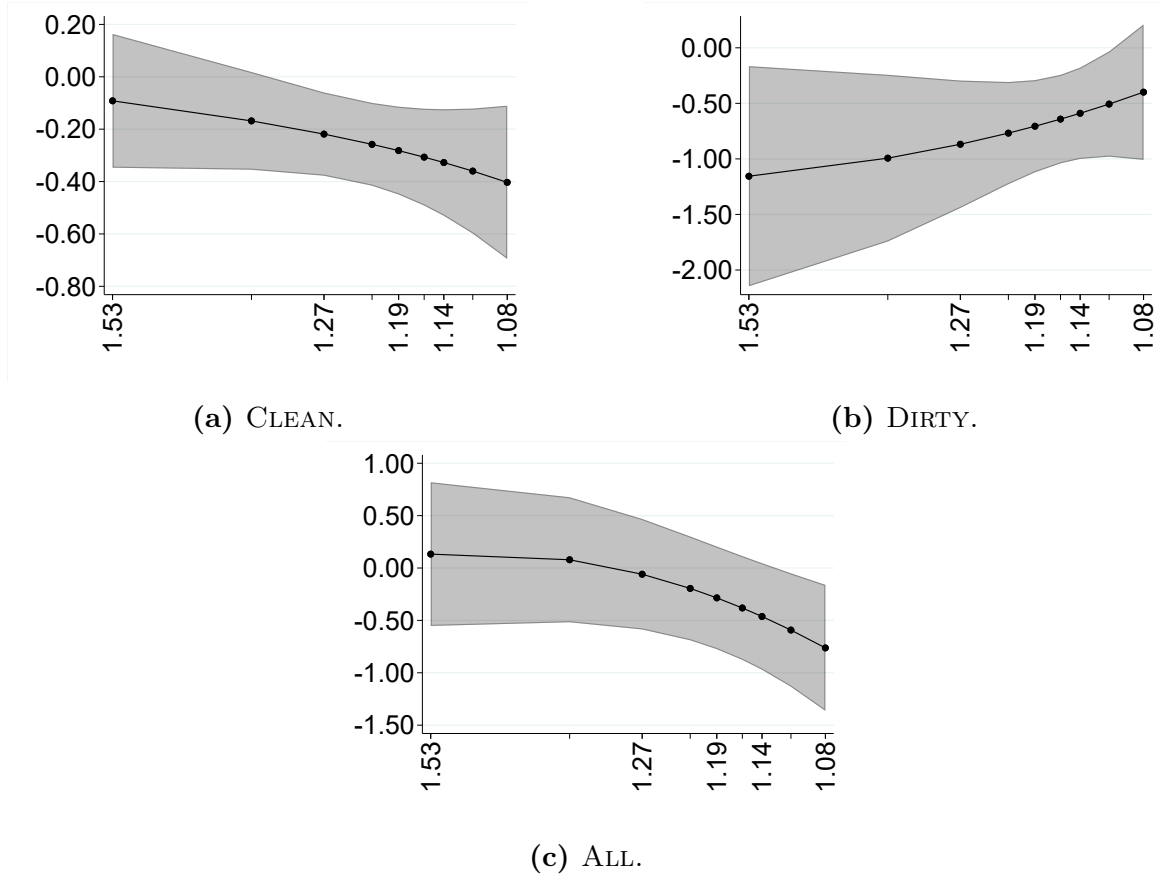
Average marginal effects are calculated using the fitted coefficients and the observed values of the independent variables. To obtain the marginal effects of energy prices on patenting under different levels of competition, the competition variable is fixed at different deciles of the competition distribution. Mind that the marginal effects cannot directly be read from the regression output, as the Poisson model is not a linear model (see Ai & Norton, 2003, for an illustration). Standard errors are calculated using the delta method.

[Figure 7](#) presents these margin plots for the full Poisson specification. The separate figures present findings for the different technology types. The vertical axes present the marginal effects of energy prices on patenting, which can be interpreted as elasticities due to the log-log relationship. The horizontal axes describe different levels of market competition. A marginal effect is presented for the 10th to the 90th competition decile. To give the horizontal axis a readable interpretation, the nominal values of the markups are printed. For example, the marginal effect of energy prices on clean patenting is around  $-0.4$  at the 90th percentile of competition. So, when markups are 1.08 (or 8%), a 1% energy price increase leads to a  $\exp(-0.4) - 1 \approx -0.33$  percentage change in patenting, or a 0.33% decrease.

The findings show that energy prices mostly negatively affect the number of patent applications. For the range of competition considered for these figures, the elasticities are negative for both clean and dirty innovation. The figures however show that the elasticity is smaller in absolute terms for clean innovation as it is for dirty innovation, for all considered values of competition. Further, the way competition affects the patent elasticities is different. Whereas the elasticity grows in absolute size for clean innovation when competition increases, it reduces in absolute size for dirty innovation. For many of the competition deciles the elasticities are statistically significant. Note that the magnitudes are rather large, especially for dirty patenting.

The elasticity of overall innovation with respect to energy prices also differs with market competition. Here the elasticity starts off slightly positive and becomes negative for higher levels of competition. Only for the last two competition deciles the elasticities are statistically significant.

Analogously, the same plots can be produced for the specification with energy price changes. The elasticities are presented in [Figure 8](#). The effect of energy prices on clean innovation is positive for low levels of competition, but turns negative for



**Figure 7:** MARGINAL EFFECTS OF ENERGY PRICES ON INNOVATION (VERT. AXIS) FOR DIFFERENT LEVELS OF COMPETITION (HOR. AXIS).

*Note:* The vertical axis represents the marginal effect of energy prices on patenting, which can be interpreted as an elasticity. The horizontal axis describes the level of market competition. Competition is measured on a log scale ( $-\log(\text{Markup})$  to be exact), and for readability the untransformed values of markup are indicated. The values correspond with the 10th to the 90th percentile of the markup variable in the total sample. Lower markups mean more competition. Patenting is measured at the main offices. The separate plots split estimates by technology type. The regressions are estimated using pseudo-Poisson maximum likelihood for the full model of Equation 3. Shaded areas represent 95% confidence intervals.

above-median levels of competition. Interestingly, the opposite picture occurs for dirty patenting, although the switch from negative to positive occurs only after the 60th percentile of competition. Overall patenting sees a positive effect from energy prices for all levels of market competition, except for the last decile, although these estimates are not statistically significant. Also the dirty patenting elasticities are not statistically significant. The estimates for clean patenting are statistically significant only for low levels of competition, where the elasticity is positive.

These results are to some extent in line with the findings in [Table 3](#), where the energy price coefficient was positive for clean patenting and negative for dirty and overall patenting. These findings are repeated for low levels of competition for clean and dirty innovation, but not for higher levels of market competition. For overall innovation, the addition of the competition variables has turned the energy price coefficient positive. However, the main take-away from this combined hypothesis, is that estimates are statistically uncertain. It is therefore difficult to draw strong conclusions from these findings.

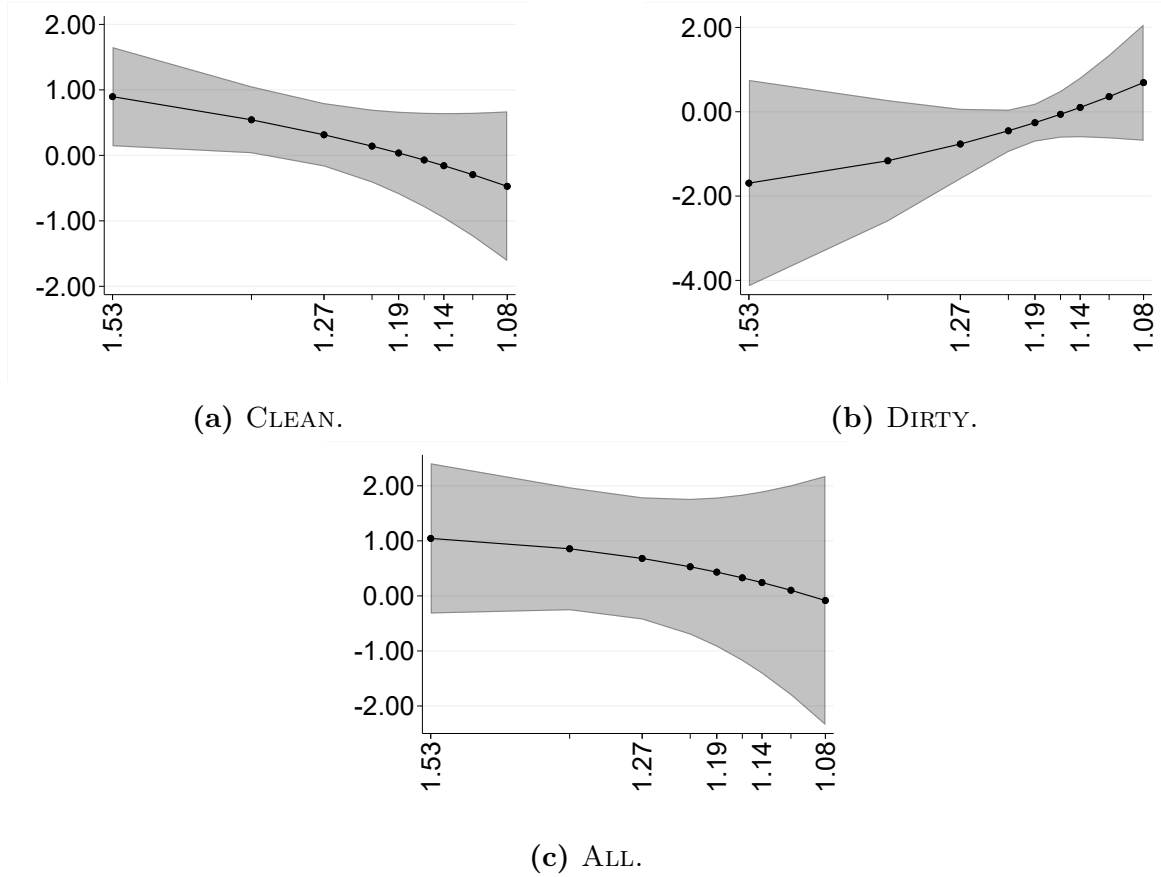
## 5.4 Robustness

Several additional analyses can be performed to complement the above main findings. First, the analysis can be performed on the level of the individual firm instead of the firm group. Second, patent quality can be restricted further by studying Triadic patent families, patents that are applied for at all three main patent offices in the US, Europe and Japan. So far the focus has been on patent families of which at least one patent is applied for at any of the main patent offices. Third, I fit a regression specification for the induced innovation hypothesis in which both energy price levels and energy price changes are included. Fourth, I study the patenting responses to energy prices of firm groups in industries with different energy intensities. Fifth, I relax the quadratic form of the relationship between competition and innovation.

The results of the firm-level analyses are presented in [Table 9](#) for the separate hypotheses. For this specification only patents at the main patent offices are considered. The results for the specification with all patent offices are reported in [Table 22](#) in [Appendix D](#). The regression specifications are the same as in [Table 2](#) and [3](#), but for conciseness not all coefficient estimates are reported. These results are qualitatively in line with the firm group results for the induced innovation hypothesis. The analysis using energy price changes shows a clear story confirming the induced innovation hypothesis. The effect of energy price changes on clean patenting is quantitatively larger than in the regressions for firm groups. The results for the energy price levels are more conclusive than in the firm group analyses. The coefficient estimate corresponding to dirty patenting is negative and statistically significant when considering individual firms.

The results for the inverted-U hypothesis are reported in the last two columns of [Table 9](#). These firm-level results are noticeably different from the firm group-





**Figure 8:** MARGINAL EFFECTS OF ENERGY PRICE CHANGES ON INNOVATION (VERT. AXIS) FOR DIFFERENT LEVELS OF COMPETITION (HOR. AXIS).

*Note:* The vertical axis represents the marginal effect of the first-differenced log energy prices on patenting, which can be interpreted as an elasticity. The horizontal axis describes the level of market competition. Competition is measured on a log scale ( $-\log(\text{Markup})$  to be exact), and for readability the untransformed values of markup are indicated. The values correspond with the 10th to the 90th percentile of the markup variable in the total sample. Lower markups mean more competition. Patenting is measured at the main offices. The separate plots split estimates by technology type. The regressions are estimated using Poisson regression for the full model from [Equation 3](#). Shaded areas represent 95% confidence intervals.

level results, as both competition coefficients estimates are now negative. This does not yield the inverted-U relationship that I hypothesize. While it technically yields an inverted-U relationship, the maximum is outside the domain that is reasonably considered.

**Table 9:** INDUCED INNOVATION AND THE INVERTED-U (FIRM-LEVEL).

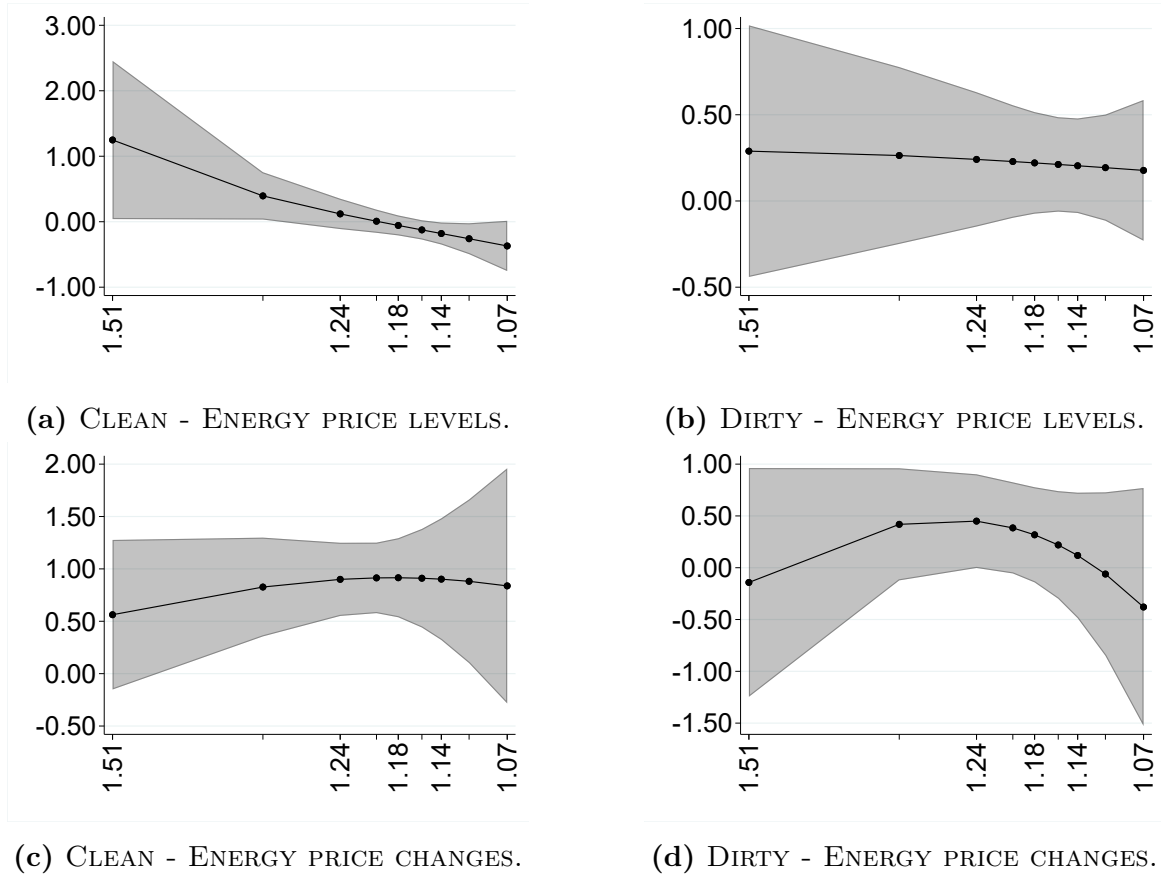
	Induced innov				Inverted-U	
	Clean	Dirty	Clean	Dirty	Clean	All
Energy prices	-0.058 (0.104)	-0.215** (0.100)				
$\Delta$ Energy prices			1.194*** (0.261)	-0.140 (0.257)		
Competition (Markup)					-0.145 (0.307)	-0.192 (0.289)
Competition <sup>2</sup>					-0.046 (0.249)	-0.043 (0.189)
Year FEs	Yes	Yes	No	No	Yes	Yes
Pre-sample controls	Yes	Yes	Yes	Yes	Yes	Yes
<i>PS, SPILL, X</i>	Yes	Yes	Yes	Yes	Yes	Yes
Observations	35,322	23,635	33,456	22,386	24,846	136,006
Unique firms	1,866	1,249	1,866	1,249	2,356	13,453
Pseudo $R^2$	0.608	0.763	0.597	0.753	0.666	0.835

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* Pseudo-Poisson maximum-likelihood estimation results. The dependent variable is the count of successful patent applications on the family level. Headers indicate the type of technology. Only patents at the main patent offices are considered. Each regression includes patent stocks, spillovers, pre-sample controls and controls as indicated in the footer of the table. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm level.

The results for the combined hypotheses are presented in [Figure 9](#). Here the analysis with energy price changes is most pronounced, with positive and statistically significant effects of energy prices on clean innovation for most levels of market competition. For moderate levels of competition energy prices also seem to induce dirty innovation, although these effects are smaller and not statistically significant. The analysis with energy price levels yields somewhat different results. Dirty patenting seems to be induced by energy prices, although these findings are not statistically significant. Energy prices induce clean innovation when market competition is low, but not when market competition is high. These firm-level results are also rather different from the firm group-level results in [Figure 7c](#) and [8c](#).

Measuring innovation, energy prices and competition on the firm level ignores two important channels. First, firm groups likely coordinate innovation decisions. If



**Figure 9:** MARGINAL EFFECTS OF ENERGY PRICES ON INNOVATION (VERT. AXIS) FOR DIFFERENT LEVELS OF COMPETITION (HOR. AXIS) ON THE FIRM LEVEL.

*Note:* The vertical axis represents the marginal effect of energy prices on patenting, which can be interpreted as an elasticity. The horizontal axis describes the level of market competition. Competition is measured on a log scale ( $-\log(\text{Markup})$  to be exact), and for readability the untransformed values of markup are indicated. The values correspond with the 10th to the 90th percentile of the markup variable in the total sample. Lower markups mean more competition. Patenting is measured at the main offices. The separate plots split estimates by technology type and the form of the energy prices (levels or first-differenced). The regressions are estimated using Poisson regression for the full model from Equation 3. Shaded areas represent 95% confidence intervals.

one firm in the group faces higher energy prices, another firm in the group might perform the innovation activities in response. Second, merging the PATSTAT and Amadeus databases on the firm level is more challenging than merging on the firm group level. Names within firm groups are often similar, making it difficult to link a patent applicant in PATSTAT to the correct individual firm in the Amadeus data, as discussed in detail in Bremer (2023). Linking to the wrong individual firm leads to a disconnect between the relationships studied in Table 9. Studying innovation behavior on the firm group levels is therefore preferred.

When only considering Triadic patents, meaning patent families consisting of patent applications that are applied for at all three major patent offices in the world, the USPTO, JPO and EPO, the number of observed patent applications drops sharply. Only about a hundred firm groups in the sample patent clean or dirty technologies in this way. The results of these analyses are presented in Table 10. The separate hypothesis tests are all statistically insignificant, likely due to the low number of observations.

**Table 10:** INDUCED INNOVATION AND THE INVERTED-U (TRIADIC PATENTS).

	Induced innov				Inverted-U	
	Clean	Dirty	Clean	Dirty	Clean	All
Energy prices	0.152 (0.305)	-0.223 (0.439)				
$\Delta$ Energy prices			1.408 (1.009)	1.365 (1.034)		
Competition (Markup)					0.757 (2.359)	0.472 (1.077)
Competition <sup>2</sup>					0.469 (2.019)	-0.555 (0.857)
Year FEs	Yes	Yes	No	No	Yes	Yes
Pre-sample controls	Yes	Yes	Yes	Yes	Yes	Yes
<i>PS, SPILL, X</i>	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,990	1,404	1,855	1,296	835	6,269
Unique firm groups	118	88	118	88	123	868
Pseudo $R^2$	0.224	0.164	0.184	0.142	0.292	0.322

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* Pseudo-Poisson maximum-likelihood estimation results. The dependent variable is the count of successful patent applications on the family level. Headers indicate the type of technology. Only Triadic patents are considered, meaning patent families applied for at all three main patent offices. Each regression includes patent stocks, spillovers, pre-sample controls and controls as specified in Equation 3 and indicated in the footer of the table. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm group level.

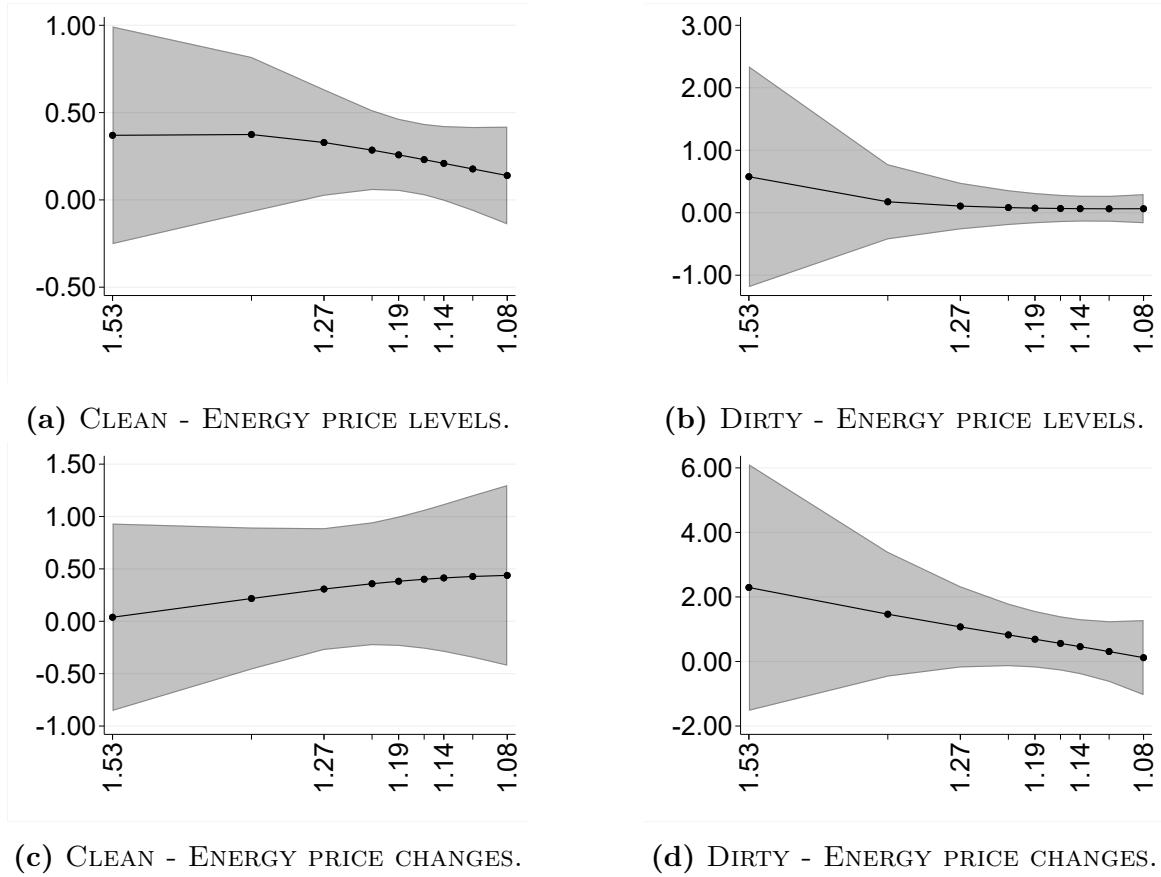
The combined hypotheses results are presented in [Figure 10](#). These results again reflect large statistical uncertainty, with nearly all estimates being statistically indistinguishable from zero. Only the marginal effect of energy prices on clean patenting is significant for moderate levels of market competition when considering the specification with energy price levels.

There is a clear tradeoff between the number of observations and patent quality. Also, there are likely sample selection effects, as not all firms pursue Triadic patents. While the related literature often prefers Triadic patents, this paper prefers studying patent families at the main patent offices. It does not constrain the quality as much, but it does result in more observations and hence more statistical power.

In the main findings section I presented two regression specifications for the induced innovation hypothesis, one with energy price levels and another with energy price changes. As the two variable specifications exploit different types of variation in the energy price data, those results provided different insights. However, there is no reason to keep the two specifications separated from each other. I here therefore present results for a regression specification in which both the energy price levels as well as the energy price changes are considered simultaneously. While not common in the induced innovation literature, the combination of energy price levels and changes in regressions is used in explaining labor market outcomes (Marin & Vona, 2019).

[Table 11](#) presents the results of a Poisson estimation. The results for clean patenting, in columns 1 and 2, are highly similar to the results from the specifications that used only one measure for energy prices. There is a negative effect of the energy price level on clean patent applications of about -0.19. This estimate is marginally statistically significant for the specification considering patent applications at the main patent offices. This estimate is close to the statistically insignificant estimate of -0.14 from the regression with only energy price levels ([Table 2](#)). The coefficient estimate for the energy price changes is again positive and statistically significant, and with 0.72 somewhat larger than the 0.57 from the regression with only energy price changes ([Table 3](#)). For the effects on dirty patenting, in column 4 for the main patent offices, the energy price changes are statistically insignificant, in line with the findings from the regression with only energy price changes. In this combined specification the coefficient estimate for the energy price level is now statistically significant and greater in absolute terms. The same pattern is found for overall patenting at the main offices. The energy price level coefficient estimate in column 6 is negative and statistically significant, compared to a statistical zero in the regression with only energy price levels. Together the findings in columns 2, 4 and 6 confirm the induced innovation hypothesis when studying energy price changes, and they suggest that energy price levels negatively impact dirty and overall patenting.

Next, I inspect the heterogeneity stemming from energy intensity. One can imagine that energy price increases are a greater challenge for firms in energy intensive industries. I therefore create an additional variable for energy intensity. Using data from



**Figure 10:** MARGINAL EFFECTS OF ENERGY PRICES ON INNOVATION (VERT. AXIS) FOR DIFFERENT LEVELS OF COMPETITION (HOR. AXIS) (TRIADIC PATENTS).

*Note:* The vertical axis represents the marginal effect of energy prices on patenting, which can be interpreted as an elasticity. The horizontal axis describes the level of market competition. Competition is measured on a log scale ( $-\log(\text{Markup})$  to be exact), and for readability the untransformed values of markup are indicated. The values correspond with the 10th to the 90th percentile of the markup variable in the total sample. Lower markups mean more competition. Only Triadic patents are considered. The separate plots split estimates by technology type and the form of the energy prices (levels or first-differenced). The regressions are estimated using Poisson regression for the full model from Equation 3. Shaded areas represent 95% confidence intervals.

**Table 11:** THE INDUCED INNOVATION HYPOTHESIS (POISSON, ENERGY PRICE LEVELS AND CHANGES SIMULTANEOUSLY).

	Clean		Dirty		All	
	All	Main	All	Main	All	Main
Energy prices	-0.190** (0.087)	-0.191* (0.106)	-0.543*** (0.111)	-0.391*** (0.063)	-0.267*** (0.074)	-0.278*** (0.075)
$\Delta$ Energy prices	0.427** (0.199)	0.723*** (0.244)	-0.144 (0.143)	0.091 (0.230)	0.002 (0.321)	-0.079 (0.221)
Patent stocks clean	1.020*** (0.058)	0.998*** (0.058)	-0.019 (0.038)	0.020 (0.036)	-0.041* (0.025)	-0.041* (0.024)
Patent stocks dirty	0.109*** (0.029)	0.146*** (0.034)	1.141*** (0.037)	1.109*** (0.043)	0.060*** (0.020)	0.053*** (0.020)
Patent stocks all	0.145*** (0.038)	0.134*** (0.044)	0.184*** (0.042)	0.112*** (0.031)	1.101*** (0.013)	1.120*** (0.017)
Spillovers country clean	-0.392** (0.154)	-0.603*** (0.200)	-0.695*** (0.152)	-0.587*** (0.121)	-0.911*** (0.163)	-0.792*** (0.132)
Spillovers country dirty	-0.122 (0.081)	-0.155* (0.093)	0.033 (0.115)	0.091 (0.147)	0.173* (0.095)	0.096 (0.084)
Spillovers country all	0.556*** (0.151)	0.796*** (0.193)	0.666*** (0.201)	0.480*** (0.166)	0.734*** (0.139)	0.672*** (0.124)
GDP per capita	-0.074* (0.041)	-0.115** (0.051)	-0.104** (0.045)	-0.145** (0.058)	-0.084** (0.035)	-0.156*** (0.037)
Constant	-2.844*** (0.880)	-3.154*** (1.025)	-0.595 (0.918)	-0.103 (1.027)	-1.922*** (0.614)	-0.777 (0.629)
Year FEs	No	No	No	No	No	No
Pre-sample controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	40,798	27,327	25,038	17,323	179,208	124,523
Unique firm groups	2,417	1,633	1,504	1,053	10,341	7,213
Pseudo $R^2$	0.600	0.604	0.835	0.825	0.850	0.849

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* Pseudo-Poisson maximum-likelihood estimation results. The dependent variable is the count of successful patent applications on the family level. Headers indicate the type of technology and the patent office. Independent patent variables are also applied for at the patent office indicated by the header. Each regression includes pre-sample controls as specified in Equation 3. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm group level.



EUROSTAT on domestic energy use and value added, I construct energy intensities on the industry level.<sup>4</sup> In order to limit endogeneity concerns, I ignore country-level variation as well as time variation for this measure. Instead, I take the industry-level energy use and value added for the entire EU in 2014. To respect the firm group structure, I again weight this measure by the number of employees of the individual firms in the firm group.<sup>5</sup> Unfortunately, energy intensities are not available for all industries for which I have energy price data. This leads to the exclusion of several industries, like construction, textile and leather manufacturing, food and tobacco, and mining and quarrying.

To study the heterogeneity, the energy intensity measure is split in terciles to create three groups of firms according to their energy intensity, denoted as *EI*. The Poisson estimation results are presented in Table 12 for three technology types and the two specifications of the energy price measures. Only results for patents at the main patent offices are presented. Columns 1-3 show that the results of the specification with energy price levels are much in line with the Poisson results without interactions in Table 2, both in coefficients size as well as statistical significance. For all three types of technologies the energy price level coefficient estimate is negative but statistically insignificant.

Turning to the results for the specification with energy price changes does yield new insights. Column 4 shows that clean induced innovation is driven by firm groups active in industries with relatively low energy intensities. The effect of energy price changes on clean patenting are statistically indistinguishable from zero for firm groups in industries with medium or high energy intensity. Interestingly, these results are flipped for dirty and overall patenting in column 5 and 6. Under low and medium energy intensity, dirty and overall patenting do not respond to energy price changes. But firm groups in energy-intensive industries strongly and negatively alter dirty and overall patenting in response to energy price changes.

These findings are in line with cost-savings responses of firms. When energy prices increase, firms in energy-intensive industries see their costs increase more than other firms. This cost increase seems to invoke a negative response in patenting outcomes, which could mean that innovation budgets are squeezed. However, columns 4-6 of Table 12 show that this squeeze takes place for dirty and overall innovation, but not for clean innovation. While the coefficient estimate of the interaction between energy prices and high energy intensity is negative in column 4, it is not statistically significant. Firms in industries with low energy intensities likely experience less pressure on operating costs when energy prices change. These firms do not reduce dirty or overall patenting output. They are, however, the firms that respond with clean patenting.

<sup>4</sup>Domestic energy use is measured in terajoules and gross value added is measured in millions of 2005 Euros.

<sup>5</sup>To be precise, energy use,  $EN$ , and value added,  $VA$ , are weighted by the number of employees,  $empl$ , of each subsidiary  $s$  in firm group  $j$ , such that energy intensity is defined as  $EI_j = (\sum_{s \in \mathcal{S}_j} empl_s EN_s) (\sum_{s \in \mathcal{S}_j} empl_s VA_s)^{-1}$ , where  $\mathcal{S}_j$  is firm group  $j$ 's set of subsidiaries.

**Table 12:** THE INDUCED INNOVATION HYPOTHESIS AND ENERGY INTENSITY (POISSON).

	Levels			Changes		
	Clean	Dirty	All	Clean	Dirty	All
Energy prices $\times$ Low EI	-0.180 (0.130)	-0.094 (0.091)	-0.049 (0.090)			
Energy prices $\times$ Medium EI	-0.216 (0.133)	-0.089 (0.096)	-0.060 (0.091)			
Energy prices $\times$ High EI	-0.207 (0.134)	-0.109 (0.095)	-0.072 (0.092)			
$\Delta$ Energy prices $\times$ Low EI				1.861*** (0.452)	-0.026 (0.383)	0.125 (0.322)
$\Delta$ Energy prices $\times$ Medium EI				-0.380 (0.401)	0.133 (0.240)	0.003 (0.321)
$\Delta$ Energy prices $\times$ High EI				-0.280 (0.331)	-1.135*** (0.392)	-1.018** (0.395)
Year FEs	Yes	Yes	Yes	No	No	No
Pre-sample controls	Yes	Yes	Yes	Yes	Yes	Yes
$PS, SPILL, X$	Yes	Yes	Yes	Yes	Yes	Yes
Observations	25,432	15,784	107,156	23,852	14,760	100,911
Unique firm groups	1,433	906	5,868	1,433	906	5,868
Pseudo $R^2$	0.620	0.789	0.844	0.609	0.783	0.835

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* Pseudo-Poisson maximum-likelihood estimation results. EI refers to energy intensity on the industry level. EI is split in three terciles, indicated by low, medium and high. The dependent variable is the count of successful patent applications on the family level at the main patent offices. Headers indicate the type of technology. Independent patent variables are also applied for at the main patent offices. Each regression includes pre-sample controls as specified in Equation 3, as well as patent stocks, spillovers, GDP per capita and a control, as indicated in the bottom of the table. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm group level.

It is this low energy intensity group that experiences induced innovation from energy price changes.


Lastly, I relax the quadratic relationship between competition and innovation. While I find some evidence of an inverted-U relationship between competition and innovation in the main findings, the coefficient estimates are often not statistically significant. This begs the question whether a quadratic relationship is describing the relationship between competition and innovation best. I therefore estimate the relationship more freely using competition deciles. I construct competition deciles based the markups that I use in my main analysis. For this analysis I also remove the time variation from this markup data in order to avoid composition effects over time and to measure the longer-term competitive pressures that a firm experiences. The average markup of each firm is taken over the period 2009-2014, and for these averages the percentiles are determined. These percentiles then enter the regression instead of the competition terms in [Equation 3](#). Besides the competition terms these regressions contain the same variables as in [Table 8](#). Margin plots will be used to present the findings. The findings for clean, dirty and all patent technologies are presented in [Figure 11](#) for the main patent offices. The analogous findings for all patent offices are presented in [Figure 14 Appendix D](#).

The results show that firm groups' patenting outcomes only slightly vary with market competition. The patenting rates are rather similar across competition deciles for all three technology types. These figures thereby do not support the inverted-U hypothesis. This somewhat contrasts the findings in [Table 8](#), although those were statistically uncertain. It does however show that when a quadratic relationship is modelled, it does follow an inverted-U. When the form of the relationship is relaxed, I do not find evidence for an inverted-U shape.

## 5.5 Discussion


The results have presented mixed evidence on the induced innovation hypothesis. The main driver of these different findings is the specification of the energy price variable. When using energy price levels, I find no evidence for induced clean innovation. When using energy price changes, measured as a growth rate, the results provide strong evidence for induced clean innovation. These two opposing results are highly robust against different estimation methods, like Poisson, negative binomial and instrumental variable regressions.

This raises the question which of the two variable specifications provides the best answer to the induced innovation hypothesis. To answer this question, I refer back to [Figure 2](#) and [3](#) which visualize the distribution and time variation of the two variable specifications. Where the cross-sectional differences are clearly dominating the time variation for the energy price levels, the opposite holds for the energy price changes. The question then becomes which type of variation one would like to exploit to answer the induced innovation hypothesis. Cross-sectional variation or time variation?



(a) CLEAN.

(b) DIRTY.



(c) ALL.

**Figure 11:** EFFECTS OF COMPETITION DECILES ON INNOVATION AT THE MAIN PATENT OFFICES.

*Note:* The vertical axis represents the effect of the competition decile on patenting. The horizontal axis describes the ranges of competition percentiles considered. For example, 90-100 refers to the highest competition decile. Competition is time invariant and derived from the average firm groups' markups over the period 2009-2014. The regressions are estimated using Poisson regression for the full model from Equation 3. Shaded areas represent 95% confidence intervals. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm group level.

When using energy price levels, the energy price coefficient from [Equation 3](#) picks up correlations between clean patent outcomes and energy prices. As energy prices are measured on the country-sector-year level some of their variation comes from country- or sector-specific characteristics. An example of sector characteristics would be the fuel mix, while an example of country characteristics would be the transport costs, the presence of natural resources, or energy policies. If for any reason these country or sector characteristics in the fuel price correlate with patenting outcomes, the energy price coefficient in the regressions might not only pick up the induced innovation effect.

On the other hand, the above country- or sector-driven correlations could also provide an answer to the induced innovation hypothesis. But instead of answering to the question what happens if energy prices change, it answers the question what happens if energy prices are structurally high (or low). The latter question is interesting in itself and likely captures innovation decisions in response to structural fuel mix and fuel price differences. The energy price coefficients estimate mainly reflects comparisons across industries and countries.

Perhaps the purer interpretation of the induced innovation hypothesis is to answer what happens to clean innovation in response to input price changes. This is more closely captured by the regression specifications with energy price changes. The three drawbacks of this specification are that (1) it ignores the price level, (2) it relies on patent outcomes responding to short-term energy price pressures, and (3) it might suffer from omitted variable bias in the time dimension. The first point could be challenged if energy price changes matter more or less depending on the energy price level it changed from. One could imagine that a 10% decrease in already low energy prices might trigger no innovation response, but that a 10% increase in already high energy prices might induce innovation. The second point is problematic if firms only change innovation strategies after multiple years of energy price changes induce them to do so. The third point is similar to the drawback of the energy price level specification. But instead of unobserved correlations over the cross-section, the energy price changes might correlate naturally with patent outcomes due to unobserved time events. For example, macro-economic events could both affect energy prices and innovation decisions.

When comparing to the literature, two key differences need to be highlighted. First, many papers that I discussed in [Section 2.1](#) consider specific industries when testing the induced innovation hypothesis (for example [Aghion et al., 2016](#); [Linn, 2008](#); [Newell et al., 1999](#); [Rozendaal & Vollebergh, forthcoming](#)). My sample tries to cover a broader range of industries, although data limitations mean that manufacturing, construction and wholesale are overrepresented (see [Table 15](#) in [Appendix C.5](#)). While a broader sample allows to test for a universal relationship between energy input prices and innovation, it might hide heterogeneity. For example, [Table 12](#) showed how firms in energy-intensive industries respond differently to energy price changes compared to

firms in industries with low energy intensity. Also, the algorithm proposed in Bremer (2023) increases the innovation coverage of my sample. Second, not all works consider the same geographical area or make use of firm-level patenting. For example, the early work of Popp (2002) considers US patents and studies them on the technology-year level, not on the firm-level.

Furthermore, the role of patent stocks have been discussed. Overall I found a clear relationship between the firm group's own patent stocks and further patenting, especially within technology classes. I have not presented consistent proof for knowledge spillovers, as measured by the relationship between a country's patent stocks and the firm group's patenting. The coefficients estimates differed between regression specifications and estimation techniques.

One reason for the lack of consistent findings might be found in the definition of the patent stocks that are available to the firm group. I have considered all country-level patent stocks in a country where the firm group is located to be available to the firm group. This assumes that firm groups located in many countries tap into a large pool of knowledge. This ignores the firm group's ability to utilize this knowledge. One can imagine that a multinational firm group has locations where local knowledge does not find its way into the organization. The location of the R&D unit might be more important than the location of a sales office. It is possible that this measure overestimates the available knowledge disproportionately for larger firms that also patent more. In that case some of the negative spillover effects could be explained. Furthermore, this simple measure ignores much of the discussion in Section 2.3 on how knowledge diffuses.

Turning to the estimation methods, I have employed three estimation methods that use the shift-share instrument for the energy prices (Table 5 and 7). For the instrument to be valid, the instrument needs to (1) correlate with the endogenous energy price variable and (2) be uncorrelated with the error term of the first stage. The first requirement is likely to hold. This requirement can also be tested with the first-stage's  $F$ -statistic. For the results in Table 5 and 18 there is no need to worry as the  $F$ -stat is well above 10, a commonly used threshold. The results in Table 7 suggest the OLS regressions might suffer from a weak instrument.

The second requirement cannot be tested statistically. Instead, one needs to argue that the instrument affects the outcome variable only through its effects on the endogenous regressor. While this might be the case, the above discussed drawbacks of the energy price level and energy price changes specifications might hint at a violation of the second requirement. While the composition effect is held constant in the shift-share instrument, the price variation is exploited. For both the specification with energy price levels as well as energy price changes, this might be a problem. For energy price changes this is problematic if fuel price changes correlate with patenting outcomes, something that could happen if both are driven by macro-economic events. For energy price levels the instrument might correlate with the error term if fuel price

levels correlate with patent outcomes. This could happen if specific sector or country characteristics explain both the energy prices as well as the patent outcomes, like, for example, country-specific policy changes.

The shift-share instrument might therefore not fully solve the potential energy price endogeneity issue entirely. However, the fact that results are highly consistent across estimation methods is reassuring to some extent.

For the inverted-U hypothesis the same question arises regarding the type of variation one should exploit. Since competition is based on markups, it contains both cross-sectional as well as time variation (see [Figure 4](#)). While competition should be a dynamic variable, as competition can change over time, one could argue that short-term time variation is somewhat undesirable. One might not want to confuse changes in short-term performance of firms, picked up by markups, for changes in competition. The question of which variation to exploit therefore also holds for the competition measure. This issue is likely exacerbated by the use of markups. While I use markups to measure competition, markups surely reflect profitability. The competition measure might therefore pick up macro trends as well. I would therefore argue that competition should mostly reflect cross-sectional differences. The depiction of the competition variable in [Figure 4](#) is therefore reassuring.

Studying the combined hypothesis has unfortunately lead to few insights. The estimations of the elasticity between patenting and energy prices for different levels of market competition have yielded mostly statistically insignificant results. From the estimation of the separate hypotheses, one can likely infer that this is driven by the uncertainty around the relationship between competition and innovation. As that relationship is statistically highly uncertain, this likely transfers to the estimates of the combined hypothesis.

## 6 Conclusion

This paper studies whether and how energy input prices and market competition induce clean innovation. These questions stem from two strands of literature, namely the induced innovation hypothesis, initiated by Hicks (1932), and the inverted-U relationship between competition and innovation as documented by Aghion et al. (2005). Specifically, I empirically test how energy prices affect clean innovation under different market competition intensities.

I have constructed a panel data set of firm groups covering mostly European countries over the period 1995-2014. Leveraging the work in Bremer (2023), I am able to merge a multitude of firms between the patent data source and the firm financials data sources, compared to simpler merging techniques. This allows me to study more firms, gaining a more complete overview of innovation behavior. To quantify innovation outcomes, I count successful patent applications on the firm level. Patents are categorized by their technology classes into clean and dirty innovations.



Energy prices from Sato et al. (2019) are used, which take into account multiple relevant energy types in each country, industry and year, and their respective prices. It further comes with a shift-share instrument, which can be used to tackle endogeneity issues. Market competition is measured by average markups in the respective country, industry and year.

In order to estimate the effects of energy prices and competition on patent counts, I resort to several empirical count models used in the literature. The base specification makes use of pseudo-Poisson maximum likelihood estimation, but I also estimate negative binomial regressions, and I instrument the energy prices with a shift-share instrument using GMM and control function estimates. Besides the estimation technique, I separately consider energy price levels and energy price changes in the induced innovation regressions. As there is no consensus in the induced innovation literature, a discussion of the results from these different regression specifications and estimates might be insightful.

I find consistent evidence for induced innovation from energy prices across estimation techniques, but only when considering energy price changes as the variable of interest. When considering energy price levels, I find no strong support of induced innovation. I have thoroughly discussed how the two energy price measures exploit different types of variation. The energy price changes exploit time variation, and energy price levels mostly exploit cross-sectional variation. While both are valid, the regressions with different variable specifications likely answer different questions. They also come with different challenges. The energy price levels could be correlated with industry or country characteristics, picking up more than just energy input price pressures. On the other hand, the energy price changes ignore level effects and only focus on short-term price changes. The question is whether firms' innovation outcomes respond to short-term incentives.

When considering the energy intensity of the industries the firms operate in, I find evidence in line with a story of budget constraints, in combination with directed technical change. Firms with high energy intensity reduce their overall innovation output in response to energy input price increases. These firms especially reduce dirty innovation, but not clean innovation. And firms operating in industries with low energy intensity do not significantly reduce overall or dirty patenting, but they do increase clean patenting.

While I do find an inverted-U relationship between competition and innovation, the estimates are statistically uncertain. Furthermore, the maximum of the inverted-U is reached at low levels of competition, effectively making the relationship between competition and innovation linear and negative. While this is in line with the Schumpeterian effect, the statistical uncertainty indicates a weak relationship between competition and innovation. When the quadratic relationship is relaxed, the results do not hint at an inverted-U relationship, nor at any other particular relationship. Patenting rates seem to be rather similar across firms experiencing different levels of market

competition.

When interacting the two hypotheses, the estimates are inconclusive. Overall, the induced innovation conclusions remain, although their estimates are statistically insignificant for most levels of market competition. Also, the marginal effects of energy prices on dirty and overall innovation are statistically indistinguishable from zero. The results make it difficult to conclude that the elasticity between patenting and energy prices systematically differs with the level of market competition.

These findings are to some extent in line with the literature. When considering energy price changes, I find evidence for induced clean innovation from energy price changes. And although not statistically significant, I do find some evidence for an inverted-U relationship between competition and innovation, for clean, dirty and overall innovation. However, the induced innovation findings do not replicate when considering energy price levels, while several other studies do find this evidence using such a specification. Furthermore, the inverted-U between competition and innovation peaks at relatively low levels of competition. Future studies could aim to shed light on these two matters. Which regression specification should one use to address the induced innovation hypothesis? And where does the inverted-U between competition and innovation peak?

# Appendices

## A Classification codes and clean patents

In order to identify environment-related patents, IPC and CPC classification codes are used. The OECD (2016) provides a list of codes that are considered environment-related technologies. I use their list from 2016. As the list is not machine-readable, as it for instance consists of ranges of classification codes, I combine it with the exhaustive lists of IPC and CPC codes to construct a full set of IPC and CPC codes that are considered clean technologies.<sup>6</sup>

Since the OECD's list contains ranges of patent codes, there are inevitably new or old patent codes within these ranges that were not considered when the list was created. Although the list seems to be based on the IPC and CPC codes from the 2015-12-31 versions, I do not distinguish between versions.<sup>7</sup> Hereby I assume that newly introduced IPC and CPC codes within the ranges mentioned by the OECD are of a similar category and are therefore likely to also be in the environment-related category. I argue the same for codes that have expired.

## B Data cleaning

The Amadeus Financials database is cleaned through the following steps.

1. Any flow variables, like operating revenue and cash flow is converted to its 12-month equivalent. For example, if a firm reports figures for the last 18 months, a cash flow of 1.5 million USD is converted to 1 million USD.
2. All monetary values are converted to USD.
3. Duplicates within a firm and consolidation level might occur due to the availability of multiple reporting bases and accounting practices for a few firms. These duplicates are dropped by taking the reporting basis and accounting practice with most observations for the variables number of employees, total assets and EBITDA margin.
4. Because not all firms report at year-end, the financials are not comparable. For the observations that are not reported at year-end, the values at year-end are interpolated or extrapolated. To do so, exact reporting dates are discarded and replaced with the quarter of the reporting date. This avoids interpolating values close to year-end. For example, if a firm reports in November, its financial values

---

<sup>6</sup>The exhaustive lists of IPC and CPC codes can be found on the websites of the IPC and CPC respectively.

<sup>7</sup>The IPC and CPC versions are not mentioned. It seems to be based on the 2015-12-31 version, but some expired patent codes are still mentioned.

are considered year-end. Any firm that does not have a balanced set of year-end observations, has its year-end observations interpolated and extrapolated. When interpolating, at most 8 quarters are filled. When extrapolating, at most 2 quarters are filled. While interpolating and extrapolating have their downsides, it avoids discarding information from firms that report in quarters other than quarter four of a calendar year. For the interpolation and extrapolation a second-order polynomial is fitted through the existing data points.

5. Lastly, I deflate all monetary values to 2010 USD equivalents.

The Amadeus Subsidiaries data is cleaned through the following steps.

1. Duplicates in the Amadeus Subsidiaries data are removed after sorting the data. Ownership links between the same firms that are reported closer to January 1, 2016 and that report a larger ownership share are kept.
2. Subsidiaries that are owned for more than 100% are removed from the data.
3. Any circles are removed from the ownership structure by removing any downstream link where the subsidiary is also a parent in the ownership structure.
4. All ownership links are followed down until the subsidiaries are found that are not a parent themselves. Also all ownership shares are tracked and multiplied in order to represent the share that the ultimate parent has in the respective subsidiary. For example, if firm A owns 30% of firm B and firm B owns 80% of firm C, then the data reports that firm A owns 30% of firm B and 24% (30% of 80%) of firm C.


## C Descriptives

### C.1 Publication lags

The publication lag distribution is presented in [Figure 12](#). These lags are calculated based on patent families with an earliest application year between 2000 and 2010. The sample is truncated at the year 2010 to allow for long publication lags. The choice to only consider applications as of 2000 is motivated by potential dynamics in the publication lag over time. Publication lags might have been different in the 20th century. As truncation of new patents is a potential problem, it is irrelevant to consider unrepresentative lag times.

### C.2 Patent counts


[Figure 13](#) presents the yearly successful patent counts for different technology types. Note that this figure plots these counts on a logarithmic scale in order to reveal



**Figure 12:** PUBLICATION LAG DISTRIBUTION.

*Note:* The difference in days is calculated between the patent family’s earliest filing date and earliest publication date. The histogram describes the lags in years (considered to be 365 days). The vertical axis describes the number of observations in millions. Only granted patents are considered. And only patent families with the earliest filing year between 2000 and 2010 are considered. The histogram is truncated between 0 and 5 years.

information on their growth rates. It shows that the growth rate of clean innovations have sped up since 1990, while that of dirty and other innovations have remained constant, with the exception of small burst around 1996. Note that this figure plots the same data as [Figure 1](#) in the main text.



**Figure 13:** CLEAN, GRAY, DIRTY AND OTHER PATENT COUNTS OVER TIME (LOG SCALE).

*Note:* Patent counts are on the family level. Counts are aggregated within technology class and within year. Only granted patents are considered. And only patents are considered that are applied for by firms covered in this paper. Patents at all patent offices are considered. The other patent category is plotted on the right axis.

### C.3 Double counting of patents

[Table 13](#) reports the distribution of the number of applicants per patent application. It shows that the vast majority of patents is applied for by one single applicant. This therefore reduces the concern over double counting of innovations.

**Table 13:** DOUBLE COUNTING OF PATENT APPLICATIONS.

Number of applicants	Frequency
1	94.19%
2	4.54%
3	0.64%
4+	0.63%

## C.4 Firm-level sample

The summary statistics of the firm-level sample are presented in [Table 14](#).

**Table 14:** SUMMARY STATISTICS (FIRM-LEVEL SAMPLE).

	Mean	SD	Min	Median	Max	Obs
Patents clean (all)	0.069	1.50	0.00	0.00	294.00	623,200
Patents dirty (all)	0.093	4.13	0.00	0.00	1,176.00	623,200
Patents all (all)	1.380	23.49	0.00	0.00	4,039.00	623,200
Patents clean (main)	0.044	1.12	0.00	0.00	239.00	623,200
Patents dirty (main)	0.056	2.37	0.00	0.00	583.00	623,200
Patents all (main)	0.886	15.60	0.00	0.00	1,876.00	623,200
Patents clean (triadic)	0.001	0.04	0.00	0.00	11.00	623,200
Patents dirty (triadic)	0.001	0.06	0.00	0.00	15.00	623,200
Patents all (triadic)	0.016	0.35	0.00	0.00	58.00	623,200
Patent stocks clean (main)	0.260	5.86	0.00	0.00	1,126.09	623,200
Patent stocks dirty (main)	0.372	16.10	0.00	0.00	2,597.37	623,200
Patent stocks all (main)	5.844	106.98	0.00	0.00	11,213.77	623,200
Spillovers country clean (main)	1,456.301	1,814.59	0.00	709.25	7,734.02	623,200
Spillovers country dirty (main)	1,653.107	1,772.58	0.00	721.66	6,708.64	623,200
Spillovers country all (main)	29,584.896	29,471.17	0.00	17,599.50	109,338.06	623,200
Energy prices (2010USD/toe)	843.377	420.30	212.05	760.49	2,805.88	301,317
Markup	1.261	0.32	0.54	1.18	5.33	232,270
GDP per capita (PPP 2017\$)	40,956.450	10,082.29	2,096.56	42,145.68	120,647.82	622,584

*Note:* All patent statistics are on the DOCDB family level. The term in brackets for the patent variables indicates the office at which patents are applied for, where the main offices are any of USPTO, EPO and JPO, and triadic refers to patent families applied for at all three main offices.

## C.5 Industry composition of the Amadeus databases

[Table 15](#) describes the industry compositions of the raw Amadeus Financials and Amadeus Subsidiaries databases, and the sample.

**Table 15:** INDUSTRY COMPOSITION IN THE AMADEUS DATABASES.

NACE description	Unique firms		Observations	
	Ama Fin	Ama Sub	Sample	%
A Agriculture, forestry and fishing	9,185	15,605	232	2.5
B Mining and quarrying	3,479	9,873	610	17.5
C Manufacturing	89,315	129,396	72,583	81.3
D Electricity, gas, steam and air conditioning supply	9,520	28,717	306	3.2
E Water supply; sewerage; waste management	5,779	8,851	475	8.2
F Construction	33,673	65,222	6,831	20.3
G Wholesale and retail trade; repair of motor vehicles	105,427	145,721	7,931	7.5
H Transporting and storage	21,461	43,976	840	3.9
I Accommodation and food service activities	8,817	21,422	437	5.0
J Information and communication	17,129	48,151	1,237	7.2
K Financial and insurance activities	60,589	90,917	1,377	2.3
L Real estate activities	39,109	106,564	1,026	2.6
M Professional, scientific and technical activities	44,768	101,336	4,263	9.5
N Administrative and support service activities	34,112	61,789	1,332	3.9
O Public administration and defence; social security	3,437	1,044	74	2.2
P Education	5,702	11,490	81	1.4
Q Human health and social work activities	15,136	19,821	260	1.7
R Arts, entertainment and recreation	4,408	10,220	205	4.7
S Other services activities	4,975	17,625	246	4.9
T Activities of households as employers	57	2,184	18	31.6
U Activities of extraterritorial organisations and bodies	48	51	0	0.0

*Note:* The industry codes refer to the NACE rev.2 classification. For compactness I shortened a few descriptions. Ama Fin and Ama Sub refer to the Amadeus Financials and Subsidiaries databases, respectively. Those numbers refer to unique firms. These numbers do not consider missing data or data availability over time. The sample column, however, does consider data availability and provides the number of observations that have non-missing values for the energy price, competition and gdp per capita. It thereby roughly aligns with the sample used for testing the combined hypotheses. Note that the sample contains firms groups, but that the NACE code of the largest firm in the group is used for counting here. The percentage column represents the ratio between the sample column and the Ama Fin column. As it compares total observations with the number of unique firms, it is not bounded by 100%, but it does show for which industries energy price and competition data is missing.

## C.6 Examples of patent technologies

Table 16 presents examples of common patent technologies. The table is split by technology types (clean, dirty, gray).

## D Further results

Table 17 presents regression results for the induced innovation hypothesis when considering patent applications at all patent offices, and when fitting the count model with a negative binomial distribution. Table 18 presents induced innovation results with energy price levels as the variable of interest and instrumented energy prices. Reduced form specifications results, where the shift-share instrument is used directly, are provided in Table 19. Table 20 presents Poisson regression results for the induced innovation hypothesis with gray patent counts as the dependent variable. Table 21 presents results for the inverted-U hypothesis from a regression without year fixed effects. Table 22 presents firm-level results for the induced innovation and the inverted-U hypotheses. Figure 14 presents free-form estimates of the relationship between competition and innovation at all patent offices.



**Table 16: EXAMPLES OF COMMON PATENT TECHNOLOGIES BY TECHNOLOGY TYPE.**

Type	Group	Freq	Specific tech	Freq	Description
Clean	Y02T0010	953,328	Y02T0010700500	77,937	Energy storage systems for electromobility, e.g. batteries
Clean	C02F0001	648,519	C02F0001280000	59,455	Treatment of water, waste water, or sewage by sorption (using ion-exchange; sorbent compositions)
Clean	Y02E0010	643,902	Y02E0010500000	95,418	Photovoltaic (PV) energy
Clean	Y02P0020	296,224	Y02P0020550000	65,845	Technologies relating to chemical industry - Design of synthesis routes, e.g. reducing the use of auxiliary or protecting groups
Clean	F01N0003	293,852	F01N0003280000	44,284	Exhaust or silencing apparatus having means for purifying, rendering innocuous, or otherwise treating exhaust - Construction of catalytic reactors
Dirty	F02D0041	994,888	F02D0041140000	88,149	Electrical control of supply of combustible mixture or its constituents - Circuit arrangements for generating control signals - Introducing closed-loop corrections
Dirty	C10L0001	396,803	C10L0001180000	19,078	Liquid carbonaceous fuels - containing oxygen
Dirty	E21B0043	268,042	E21B0043000000	24,399	Methods or apparatus for obtaining oil, gas, water, soluble or meltable materials or a slurry of minerals from wells
Dirty	F23G0005	233,356	F23G0005000000	40,221	Incineration of waste; Incinerator constructions; Details, accessories or control therefor
Dirty	C09K0008	194,252	C09K0008800000	8,258	Compositions for drilling of boreholes or wells; Compositions for treating boreholes or wells, e.g. for completion or for remedial operations - containing organic compounds
Gray	Y02E0020	160,708	Y02E0020348000	34,694	Combustion technologies with mitigation potential - Indirect CO2 mitigation, i.e. by acting on non CO2 directly related matters of the process, e.g. pre-heating or heat recovery
Gray	Y02T0010	77,192	Y02T0010146000	71,260	Climate change mitigation technologies related to transport - Internal combustion engine (ICE) based vehicles
Gray	Y02P0020	59,564	Y02P0020129000	43,472	Technologies relating to chemical industry - Energy recovery, e.g. by cogeneration, H2 recovery or pressure recovery turbines
Gray	Y02P0040	51,570	Y02P0040570000	47,858	Technologies relating to the processing of minerals - Improving the yield, e.g. reduction of reject rates
Gray	Y02P0010	49,408	Y02P0010253000	32,342	Technologies related to metal processing - Process efficiency

*Note:* Technologies are identified through their IPC or CPC classes. For each technology type (clean, dirty, gray) the five most common 8-digit technologies are counted (group) and within that the most common specific technology is counted are described (Specific tech, Description). Note that the technology groups (Group) can span multiple technology types. Description data source: Espacenet's CPC classification search.

**Table 17:** THE INDUCED INNOVATION HYPOTHESIS (NEGATIVE BINOMIAL, ALL PATENT OFFICES).

	Levels			Changes		
	Clean	Dirty	All	Clean	Dirty	All
Energy prices	-0.082 (0.052)	0.009 (0.071)	-0.047** (0.022)			
$\Delta$ Energy prices				0.402** (0.183)	-0.303 (0.196)	-0.496*** (0.063)
Patent stocks clean	1.164*** (0.030)	0.032 (0.026)	0.042*** (0.013)	1.167*** (0.030)	0.034 (0.027)	0.024* (0.013)
Patent stocks dirty	0.075*** (0.017)	1.253*** (0.032)	0.032** (0.013)	0.075*** (0.018)	1.264*** (0.032)	0.030** (0.013)
Patent stocks all	0.263*** (0.019)	0.207*** (0.023)	1.228*** (0.007)	0.267*** (0.019)	0.200*** (0.023)	1.224*** (0.008)
Spillovers country clean	0.091 (0.101)	0.113 (0.138)	0.236*** (0.045)	-0.108 (0.088)	-0.301*** (0.116)	-0.403*** (0.036)
Spillovers country dirty	-0.040 (0.057)	0.133* (0.073)	-0.078*** (0.027)	-0.017 (0.056)	0.144** (0.072)	0.007 (0.025)
Spillovers country all	-0.022 (0.097)	-0.205 (0.143)	-0.048 (0.043)	0.154* (0.086)	0.179 (0.115)	0.475*** (0.035)
GDP per capita	-0.047 (0.029)	-0.034 (0.033)	-0.079*** (0.012)	-0.049* (0.030)	-0.052 (0.034)	-0.132*** (0.012)
Constant	-2.322*** (0.523)	-2.497*** (0.668)	-1.798*** (0.210)	-3.505*** (0.476)	-3.469*** (0.569)	-3.290*** (0.201)
Year FEs	Yes	Yes	Yes	No	No	No
Pre-sample controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	43,412	26,690	190,012	40,798	25,038	179,208
Unique firm groups	2,417	1,504	10,341	2,417	1,504	10,341
Pseudo $R^2$	0.226	0.266	0.256	0.220	0.262	0.251

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* Negative binomial maximum-likelihood estimation results. The dependent variable is the count of successful patent applications on the family level. Headers indicate the regression specification and the type of technology. Patents applied for at any patent office are considered. Each regression includes pre-sample controls. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm group level.

**Table 18:** THE INDUCED INNOVATION HYPOTHESIS (INSTRUMENTED ENERGY PRICES).

	GMM			Control function		
	Clean	Dirty	All	Clean	Dirty	All
Energy prices	0.074 (0.114)	-0.092 (0.112)	0.155 (0.104)	-0.219 (0.393)	0.017 (0.499)	0.174 (0.342)
First-stage residual				-0.775 (0.812)	-0.618 (0.729)	-0.346 (0.532)
Patent stocks clean	0.999*** (0.054)	-0.024 (0.033)	-0.054* (0.027)			
Patent stocks dirty	0.125*** (0.030)	1.162*** (0.037)	0.059*** (0.017)			
Patent stocks all	0.150*** (0.039)	0.154*** (0.033)	1.159*** (0.017)			
Spillovers country clean	0.142 (0.243)	0.559** (0.225)	0.072 (0.156)			
Spillovers country dirty	-0.254** (0.106)	0.103 (0.145)	0.041 (0.077)			
Spillovers country all	0.164 (0.232)	-0.598** (0.287)	-0.077 (0.175)			
GDP per capita	-0.083 (0.054)	-0.096* (0.053)	-0.090*** (0.034)	-0.067 (0.801)	-0.811 (1.022)	0.269 (0.622)
Constant	-3.115*** (0.917)	0.451 (0.729)	-1.942*** (0.733)	3.542 (8.458)	11.329 (12.020)	0.254 (7.042)
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Firm FEs	No	No	No	Yes	Yes	Yes
Pre-sample controls	Yes	Yes	Yes	No	No	No
Observations	29,115	18,497	132,144	27,654	17,432	127,977
Unique firm groups	1,633	1,053	7,213	1,537	985	6,963
Pseudo $R^2$				0.651	0.846	0.867
First-stage F-stat				20.31	62.55	76.00

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* GMM and control function results. The GMM specification assumes a Poisson distribution. The control function specification is in line with Dugoua and Gerarden (2023) and includes the first-stage error term in a second-stage pseudo-Poisson maximum likelihood model. The dependent variable is the count of successful patent applications on the family level. Headers indicate the estimation method and the type of technology. Only patent applied for at the main patent offices are considered. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm group level.

**Table 19:** THE INDUCED INNOVATION HYPOTHESIS (POISSON, REDUCED FORM).

	Levels			Changes		
	Clean	Dirty	All	Clean	Dirty	All
Energy prices	0.057 (0.088)	-0.078 (0.096)	0.115 (0.081)			
$\Delta$ Energy prices				0.624** (0.288)	0.128 (0.219)	-0.147 (0.169)
Patent stocks clean	0.998*** (0.053)	-0.027 (0.032)	-0.051* (0.027)	0.995*** (0.059)	0.020 (0.036)	-0.044* (0.026)
Patent stocks dirty	0.125*** (0.030)	1.160*** (0.036)	0.059*** (0.017)	0.139*** (0.032)	1.072*** (0.046)	0.049*** (0.018)
Patent stocks all	0.150*** (0.039)	0.158*** (0.033)	1.158*** (0.017)	0.139*** (0.042)	0.102*** (0.031)	1.116*** (0.018)
Spillovers country clean	0.141 (0.242)	0.580*** (0.215)	0.068 (0.153)	-0.758*** (0.223)	-0.547*** (0.143)	-1.001*** (0.137)
Spillovers country dirty	-0.256** (0.104)	0.107 (0.137)	0.028 (0.072)	-0.085 (0.081)	0.215** (0.106)	0.163** (0.070)
Spillovers country all	0.167 (0.228)	-0.625** (0.267)	-0.059 (0.168)	0.873*** (0.204)	0.300 (0.208)	0.809*** (0.130)
GDP per capita	-0.084 (0.053)	-0.100* (0.051)	-0.097*** (0.035)	-0.118** (0.050)	-0.204*** (0.064)	-0.177*** (0.037)
Constant	-3.139*** (0.870)	-0.077 (0.676)	-2.167*** (0.650)	-4.552*** (0.885)	-1.321 (1.224)	-2.792*** (0.633)
Year FEs	Yes	Yes	Yes	No	No	No
Pre-sample controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	29,115	18,497	132,144	27,327	17,323	124,523
Unique firm groups	1,633	1,053	7,213	1,633	1,053	7,213
Pseudo $R^2$	0.614	0.828	0.855	0.604	0.823	0.847

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* Pseudo-Poisson maximum-likelihood estimation results. The energy price variable is the shift-share instrument, making this a reduced form specification. The dependent variable is the count of successful patent applications on the family level. Headers indicate the type of technology and the patent office. Independent patent variables are also applied for at the patent office indicated by the header. Each regression includes pre-sample controls. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm group level.

**Table 20:** THE INDUCED INNOVATION HYPOTHESIS FOR GRAY PATENTS (POISSON).

	Levels		Changes	
	All	Main	All	Main
Energy prices	-0.434 (0.273)	-0.513 (0.319)		
$\Delta$ Energy prices			-0.848** (0.421)	-0.250 (0.459)
Year FEs	Yes	Yes	No	No
Pre-sample controls	Yes	Yes	Yes	Yes
<i>PS, SPILL, X</i>	Yes	Yes	Yes	Yes
Observations	4,962	3,847	4,623	3,583
Unique firm groups	292	227	292	227
Pseudo $R^2$	0.376	0.349	0.365	0.326

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* Pseudo-Poisson maximum-likelihood estimation results. The dependent variable is the count of successful patent applications on the family level. Headers indicate how the energy price variable enters the equation and the patent office. Independent patent variables are also applied for at the patent office indicated by the header. Each regression includes pre-sample controls as specified in [Equation 3](#). Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm group level.

**Table 21:** TESTING THE INVERTED-U HYPOTHESIS (POISSON, NO YEAR FIXED EFFECTS).

	All offices			Main offices		
	Clean	Dirty	All	Clean	Dirty	All
Competition (Markup)	0.271 (0.412)	0.660 (0.461)	0.407 (0.296)	0.713 (0.571)	0.764 (0.563)	0.933** (0.364)
Competition <sup>2</sup>	-0.556* (0.323)	-0.827** (0.409)	-0.496* (0.282)	-0.764* (0.455)	-0.648 (0.501)	-0.827*** (0.319)
Patent stocks clean	0.971*** (0.037)	0.036 (0.042)	-0.022 (0.032)	0.953*** (0.037)	0.065 (0.048)	-0.036 (0.039)
Patent stocks dirty	0.106*** (0.034)	1.179*** (0.042)	0.071*** (0.018)	0.137*** (0.039)	1.157*** (0.050)	0.067*** (0.020)
Patent stocks all	0.156*** (0.042)	0.079* (0.045)	1.077*** (0.016)	0.149*** (0.054)	0.033 (0.052)	1.101*** (0.024)
Spillovers country clean	-0.651*** (0.169)	-0.120 (0.318)	-0.505*** (0.151)	-0.795*** (0.214)	-0.126 (0.267)	-0.541*** (0.188)
Spillovers country dirty	-0.081 (0.071)	0.030 (0.141)	0.043 (0.085)	0.013 (0.114)	0.142 (0.176)	-0.004 (0.129)
Spillovers country all	0.699*** (0.207)	0.084 (0.303)	0.442** (0.192)	0.741*** (0.266)	-0.019 (0.312)	0.548** (0.275)
GDP per capita	-0.144*** (0.055)	-0.054 (0.071)	-0.050 (0.051)	-0.140* (0.074)	-0.047 (0.081)	-0.031 (0.066)
Constant	-3.214*** (1.139)	-2.645*** (1.005)	-2.980*** (0.955)	-3.251** (1.565)	-2.237** (1.041)	-3.767*** (1.367)
Year FEs	No	No	No	No	No	No
Pre-sample controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	22,279	13,102	112,152	13,645	8,339	72,117
Unique firm groups	2,844	1,680	13,865	1,796	1,111	9,018
Pseudo $R^2$	0.675	0.831	0.843	0.701	0.836	0.849
Markup at min/max	1.28	1.49	1.51	1.59	1.80	1.76
Joint test competition ( $p$ )	0.000	0.005	0.184	0.043	0.399	0.032

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

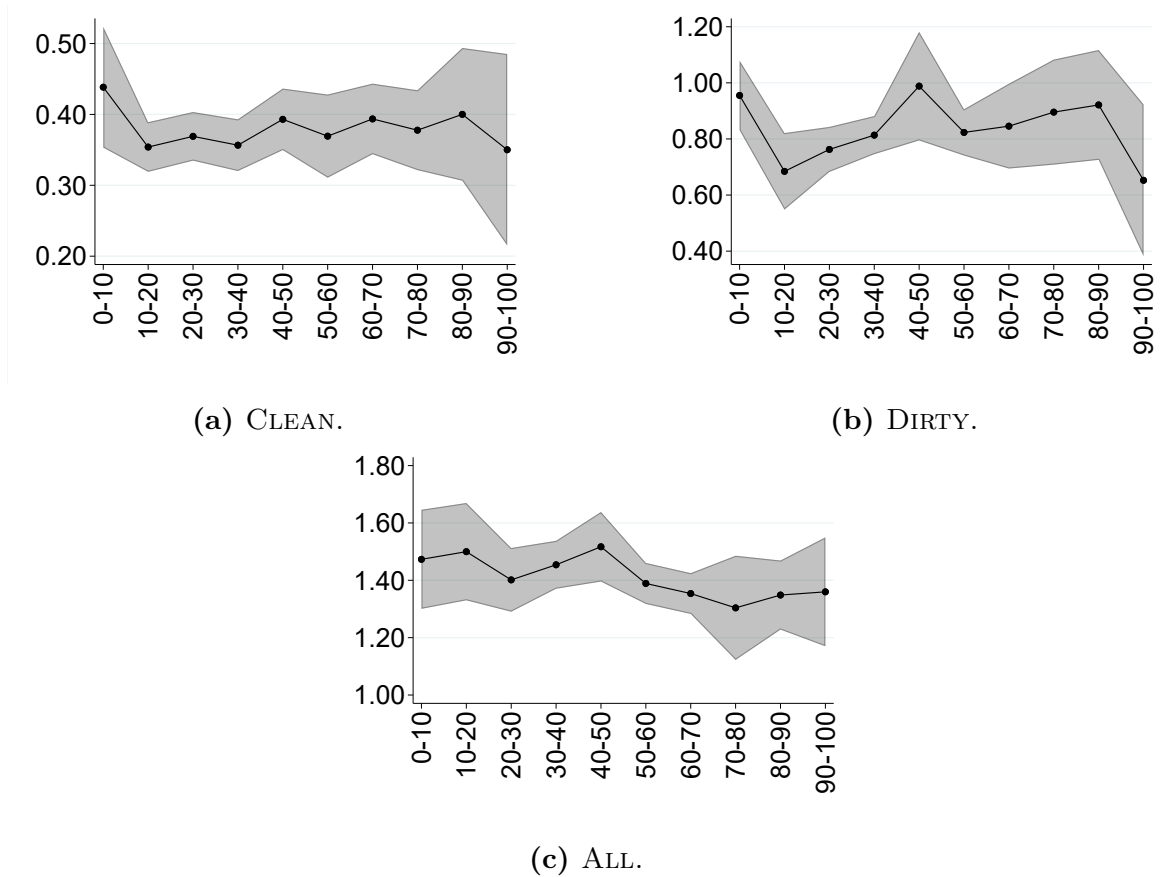
*Note:* Pseudo-Poisson maximum-likelihood estimation results. The dependent variable is the count of successful patent applications on the family level. Headers indicate the type of technology and the considered patent offices. Independent patent variables are also applied for at the patent office indicated by the header. Each regression includes pre-sample controls. Only firms that patented in the category of the dependent variable during the time period of the sample are included. The joint test presents the  $p$ -value from a Wald test with the null hypothesis stating that both competition terms are zero. Standard errors are clustered at the firm group level.

**Table 22:** INDUCED INNOVATION AND THE INVERTED-U (FIRM-LEVEL, ALL OFFICES).

	Induced innov				Inverted-U	
	Clean	Dirty	Clean	Dirty	Clean	All
Energy prices	-0.049 (0.102)	-0.206 (0.135)				
$\Delta$ Energy prices			0.770*** (0.199)	-0.483** (0.213)		
Competition (Markup)					-0.244 (0.263)	-0.089 (0.212)
Competition <sup>2</sup>					0.009 (0.211)	-0.105 (0.150)
Year FEs	Yes	Yes	No	No	Yes	Yes
Pre-sample controls	Yes	Yes	Yes	Yes	Yes	Yes
<i>PS, SPILL, X</i>	Yes	Yes	Yes	Yes	Yes	Yes
Observations	51,961	33,479	49,206	31,698	37,834	211,033
Unique firms	2,755	1,781	2,755	1,781	3,657	21,328
Pseudo $R^2$	0.612	0.776	0.605	0.769	0.671	0.834

Clustered standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:* Pseudo-Poisson maximum-likelihood estimation results. The dependent variable is the count of successful patent applications on the family level. Headers indicate the type of technology. Patents at any patent office are considered. Each regression includes patent stocks, spillovers, pre-sample controls and controls as indicated in the footer of the table. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm level.



**Figure 14:** EFFECTS OF COMPETITION DECILES ON INNOVATION AT ALL PATENT OFFICES.

*Note:* The vertical axis represents the effect of the competition decile on patenting. The horizontal axis describes the ranges of competition percentiles considered. For example, 90-100 refers to the highest competition decile. Competition is time invariant and derived from the average firm groups' markups over the period 2009-2014. The regressions are estimated using Poisson regression for the full model from [Equation 3](#). Shaded areas represent 95% confidence intervals. Only firms that patented in the category of the dependent variable during the time period of the sample are included. Standard errors are clustered at the firm group level.



## References

- Abramovitz, M. (1986). Catching up, forging ahead, and falling behind. *Journal of Economic History*, 46(2), 385–406.
- Abrams, D. S., Akcigit, U., & Grennan, J. (2013, November). *Patent value and citations: Creative destruction or strategic disruption?* (Working Paper No. 19647). National Bureau of Economic Research.
- Aghion, P., Bénabou, R., Martin, R., & Roulet, A. (2023). Environmental preferences and technological choices: Is market competition clean or dirty? *American Economic Review: Insights*, 5(1), 1–20.
- Aghion, P., Bloom, N., Blundell, R., Griffith, R., & Howitt, P. (2005). Competition and innovation: An inverted-U relationship. *Quarterly Journal of Economics*, 120(2), 701–728.
- Aghion, P., Blundell, R., Griffith, R., Howitt, P., & Prantl, S. (2009). The effects of entry on incumbent innovation and productivity. *Review of Economics and Statistics*, 91(1), 20–32.
- Aghion, P., Dechezleprêtre, A., Hemous, D., Martin, R., & Van Reenen, J. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1), 1–51.
- Ai, C., & Norton, E. C. (2003). Interaction terms in logit and probit models. *Economics Letters*, 80(1), 123–129.
- Ambec, S., Cohen, M. A., Elgie, S., & Lanoie, P. (2013). The Porter Hypothesis at 20: Can environmental regulation enhance innovation and competitiveness? *Review of Environmental Economics and Policy*, 7(1), 2–22.
- Bao, C. G., & Chen, M. X. (2018). Foreign rivals are coming to town: Responding to the threat of foreign multinational entry. *American Economic Journal: Applied Economics*, 10(4), 120–57.
- Bartik, T. J. (1991). *Who benefits from state and local economic development policies?* W.E. Upjohn Institute for Employment Research.
- Blundell, R., Griffith, R., & Van Reenen, J. (1995). Dynamic count data models of technological innovation. *Economic Journal*, 105(429), 333–344.
- Blundell, R., Griffith, R., & Van Reenen, J. (1999). Market share, market value and innovation in a panel of British manufacturing firms. *Review of Economic Studies*, 66(3), 529–554.
- Bremer, L. (2023). *Fuzzy firm name matching: Merging Amadeus firm data to PATSTAT* (Discussion Paper No. 23-055/VIII). Tinbergen Institute.
- Calel, R., & Dechezleprêtre, A. (2016). Environmental policy and directed technological change: Evidence from the European carbon market. *Review of Economics and Statistics*, 98(1), 173–191.
- Cockburn, I., & Griliches, Z. (1988). Industry effects and appropriability measures in the stock market’s valuation of R&D and patents. *American Economic Review*, 78(2), 419–423.

- Coe, D. T., & Helpman, E. (1995). International R&D spillovers. *European Economic Review*, 39(5), 859–887.
- CompNet. (2022). User guide for the 9th vintage of the CompNet dataset.
- Correia, S., Guimarães, P., & Zylkin, T. (2020). Fast poisson estimation with high-dimensional fixed effects. *Stata Journal*, 20(1), 95–115.
- De Loecker, J., Eeckhout, J., & Unger, G. (2020). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics*, 135(2), 561–644.
- De Rassenfosse, G., & Jaffe, A. B. (2018). Econometric evidence on the depreciation of innovations. *European Economic Review*, 101, 625–642.
- Dechezleprêtre, A., Martin, R., & Mohnen, M. (2017). *Knowledge spillovers from clean and dirty technologies* (Working Paper No. 135). Grantham Research Institute on Climate Change and the Environment.
- Dechezleprêtre, A., Ménière, Y., & Mohnen, M. (2017). International patent families: From application strategies to statistical indicators. *Scientometrics*, 111(2), 793–828.
- Dechezleprêtre, A., Muckley, C. B., & Parvati, N. (2019). *Is firm-level clean or dirty innovation valued more?* (Centre for Climate Change Economics and Policy Working Paper No. 358). London School of Economics and Political Science. London.
- Dernis, H., & Khan, M. (2004). *Triadic patent families methodology* (Working Paper). Paris, OECD Publishing.
- Dugoua, E., & Gerarden, T. (2023, October). *Induced innovation, inventors, and the energy transition* (Working Paper No. 31714). National Bureau of Economic Research.
- Fischer, C., Parry, I. W., & Pizer, W. A. (2003). Instrument choice for environmental protection when technological innovation is endogenous. *Journal of Environmental Economics and Management*, 45(3), 523–545.
- Ganapati, S., Shapiro, J. S., & Walker, R. (2020). Energy cost pass-through in US manufacturing: Estimates and implications for carbon taxes. *American Economic Journal: Applied Economics*, 12(2), 303–42.
- Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8), 2586–2624.
- Griliches, Z. (1990). Patent statistics as economic indicators: A survey. *Journal of Economic Literature*, 28(4), 1661–1707.
- Grossman, G. M., & Helpman, E. (1991). Trade, knowledge spillovers, and growth. *European Economic Review*, 35(2-3), 517–526.
- Grubb, M., Drummond, P., Poncia, A., McDowall, W., Popp, D., Samadi, S., Penasco, C., Gillingham, K. T., Smulders, S., Glachant, M., Hassall, G., Mizuno, E., Rubin, E. S., Dechezleprêtre, A., & Pavan, G. (2021). Induced innovation in energy technologies and systems: A review of evidence and potential im-

- plications for CO2 mitigation. *Environmental Research Letters*, 16(043007), 1–48.
- Hall, B. H., Jaffe, A., & Trajtenberg, M. (2005). Market value and patent citations. *RAND Journal of Economics*, 36(1), 16–38.
- Hašič, I., & Migotto, M. (2015). *Measuring environmental innovation using patent data* (Working Paper No. 89). OECD. Paris, OECD Publishing.
- Hausman, J., Hall, B. H., & Griliches, Z. (1984). Econometric models for count data with an application to the patents-R&D relationship. *Econometrica*, 52(4), 909–938.
- Hicks, J. R. (1932). *The theory of wages* (1st ed.). London: Macmillan.
- IEA. (2021, April). *Methodology for identifying fossil fuel supply related technologies in patent data* (tech. rep.).
- Jacobs, B., Nahuis, R., & Tang, P. J. G. (2002). Sectoral productivity growth and R&D spillovers in the Netherlands. *De Economist*, 150, 181–210.
- Jaffe, A. B., Newell, R. G., & Stavins, R. N. (2003). Chapter 11 - Technological change and the environment. In K.-G. Mäler & J. R. Vincent (Eds.), *Environmental degradation and institutional responses* (pp. 461–516, Vol. 1). Elsevier.
- Jaffe, A. B., & Palmer, K. (1997). Environmental regulation and innovation: A panel data study. *Review of Economics and Statistics*, 79(4), 610–619.
- Jaffe, A. B., Trajtenberg, M., & Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics*, 108(3), 577–598.
- Jee, S. J., & Srivastav, S. (2024). Knowledge spillovers between clean and dirty technologies: Evidence from the patent citation network. *Ecological Economics*, 224, 108310.
- Johnstone, N., Hašič, I., & Popp, D. (2010). Renewable energy policies and technological innovation: Evidence based on patent counts. *Environmental and Resource Economics*, 45(1), 133–155.
- Keller, W. (1998). Are international R&D spillovers trade-related? Analyzing spillovers among randomly matched trade partners. *European Economic Review*, 42(8), 1469–1481.
- Kelly, B., Papanikolaou, D., Seru, A., & Taddy, M. (2021). Measuring technological innovation over the long run. *American Economic Review: Insights*, 3(3), 303–20.
- Lanoie, P., Laurent-Lucchetti, J., Johnstone, N., & Ambec, S. (2011). Environmental policy, innovation and performance: New insights on the Porter hypothesis. *Journal of Economics & Management Strategy*, 20(3), 803–842.
- Linn, J. (2008). Energy prices and the adoption of energy-saving technology. *Economic Journal*, 118(533), 1986–2012.

- Marin, G., & Vona, F. (2019). Climate policies and skill-biased employment dynamics: Evidence from EU countries. *Journal of Environmental Economics and Management*, *98*, 102253.
- Marin, G., & Vona, F. (2021). The impact of energy prices on socioeconomic and environmental performance: Evidence from french manufacturing establishments, 1997-2015. *European Economic Review*, *135*, 103739.
- Nesta, L., Vona, F., & Nicolli, F. (2014). Environmental policies, competition and innovation in renewable energy. *Journal of Environmental Economics and Management*, *67*(3), 396–411.
- Newell, R. G., Jaffe, A. B., & Stavins, R. N. (1999). The induced innovation hypothesis and energy-saving technological change. *Quarterly Journal of Economics*, *114*(3), 941–975.
- Nickell, S., Nicolitsas, D., & Dryden, N. (1997). What makes firms perform well? [Paper and Proceedings of the Eleventh Annual Congress of the European Economic Association]. *European Economic Review*, *41*(3), 783–796.
- Nickell, S. J. (1996). Competition and corporate performance. *Journal of Political Economy*, *104*(4), 724–746.
- Nicolli, F., & Vona, F. (2016). Heterogeneous policies, heterogeneous technologies: The case of renewable energy. *Energy Economics*, *56*, 190–204.
- Noailly, J., & Shestalova, V. (2017). Knowledge spillovers from renewable energy technologies: Lessons from patent citations. *Environmental Innovation and Societal Transitions*, *22*, 1–14.
- Noailly, J., & Smeets, R. (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.
- Noailly, J., & Smeets, R. (2016). *Financing energy innovation: The role of financing constraints for directed technical change from fossil-fuel to renewable innovation* (EIB Working Papers No. 2016/06). Luxembourg, European Investment Bank (EIB).
- OECD. (2009). *OECD patent statistics manual*.
- OECD. (2016). *Patent search strategies for the identification of selected environment-related technologies (ENV-TECH)* (tech. rep.).
- OECD. (2021). *Methodologies to measure market competition* (Policy Paper). OECD. Paris.
- Popp, D. (2002). Induced innovation and energy prices. *American Economic Review*, *92*(1), 160–180.
- Popp, D., & Newell, R. (2012). Where does energy R&D come from? examining crowding out from energy R&D. *Energy Economics*, *34*(4), 980–991.
- Popp, D., Newell, R. G., & Jaffe, A. B. (2010). Energy, the environment, and technological change. In *Handbook of the economics of innovation* (pp. 873–937, Vol. 2). Elsevier.

- Porter, M. E., & Van der Linde, C. (1995a). Green and competitive: Ending the stalemate. In *On competition*. Harvard Business School Publishing.
- Porter, M. E., & Van der Linde, C. (1995b). Toward a new conception of the environment-competitiveness relationship. *Journal of Economic Perspectives*, 9(4), 97–118.
- Rozendaal, R. (2024). *Market power, innovation, and the green transition* (tech. rep.).
- Rozendaal, R., & Vollebergh, H. (forthcoming). Policy-induced innovation in clean technologies: Evidence from the car market. *Journal of the Association of Environmental and Resource Economists*.
- Sarr, M., & Noailly, J. (2017). Innovation, diffusion, growth and the environment: Taking stock and charting new directions. *Environmental and Resource Economics*, 66(3), 393–407.
- Sato, M., Singer, G., Dussaux, D., & Lovo, S. (2019). International and sectoral variation in industrial energy prices 1995-2015. *Energy Economics*, 78, 235–258.
- Scherer, F. M. (1967). Market structure and the employment of scientists and engineers. *American Economic Review*, 57(3), 524–531.
- Schumpeter, J. A. (1950). *Capitalism, socialism and democracy* (Third). Harper & Row.
- Shapiro, J. S., & Walker, R. (2018). Why is pollution from US manufacturing declining? The roles of environmental regulation, productivity, and trade. *American Economic Review*, 108(12), 3814–54.
- Verdolini, E., & Galeotti, M. (2011). At home and abroad: An empirical analysis of innovation and diffusion in energy technologies. *Journal of Environmental Economics and Management*, 61(2), 119–134.
- Vives, X. (2008). Innovation and competitive pressure. *Journal of Industrial Economics*, 56(3), 419–469.
- Wooldridge, J. M. (2015). Control function methods in applied econometrics. *Journal of Human Resources*, 50(2), 420–445.