Shock Therapy for Clean Innovation: Within-firm Reallocation of R&D Investments*

Esther Ann Bøler[†] Katinka Holtsmark[‡] Karen Helene Ullveit-Moe[§]

December 2024

Abstract

We analyze how a major negative shock to the producers of fossil fuels may lead to a shift from dirty to clean R&D along the supply chain. First, we develop a theoretical framework of directed technical change, showing that adjustment costs in R&D activity can lead fossil energy sector suppliers to shift their R&D activity towards clean innovation more than other firms, as a consequence of a negative oil price shock. Second, we investigate the impact of a major drop in the oil price in 2014 on clean R&D. Relying on rich firm level trade data, we propose a novel method of identifying firms' exposure to the price shock. We find that more exposed firms increased their clean R&D investments more than less exposed firms. Our findings contribute to the understanding of the drivers of clean technological change, which is vital to assess the effectiveness of different climate policy measures, including carbon pricing.

JEL codes: D25, F18, O31, Q55, Q58.

Keywords: Clean innovation, supply chains, carbon pricing.

*We thank Lassi Ahlvik, Clare Balboni, Eugenie Dugoua, Rikard Forslid, Bård Harstad, Linnea Lorentzen, Christoph Meinerding, Sjak Smulders, and Andre Trindade, as well as seminar audiences at University of Cambridge, Geneva Graduate Institute, University of Surrey, OsloMET and participants at CESifo workshop on Trade, Market Failures and Development, CITP Academic Conference, Tilburg Growth, International Macro, Finance and Trade Conference, CESifo Global Economy Conference, CITP/Sussex Workshop on IP Rights, EAERE Leuven, Imperial/LSE Workshop in Environmental Economics, NAERE Bergen, 14th Toulouse Conference on the Economics of Energy and Climate, Uppsala Transport Economics & the Environment Workshop, Technology and International Trade, CEP Annual Conference, CESifo Venice Summer Institute, 15th Nordic Conference on Register Data and Economic Modelling, Tax Research Conference 2024, and CCA Workshop on International Economics for valuable comments and discussion. This project has received funding from the Norwegian Ministry of Trade, Industry and Fisheries (INNOPROD). Data from Statistics Norway has been crucial for this project.

[†]Imperial College Business School, CEP, CEPR & CESifo; e.boler@imperial.ac.uk

[‡]University of Oslo & CESifo; k.k.holtsmark@econ.uio.no

[§]University of Oslo, CEPR & CESifo; k.h.ulltveit-moe@econ.uio.no

1 Introduction

The transition from fossil fuels to clean energy requires a large-scale shift in technological development from fossil to clean technologies. A major challenge to policy makers is to incentivize investments in clean technologies. In this paper we ask whether leveraging the technological capabilities present in the fossil energy supply chain can be part of the solution. Specifically, we investigate the extent to which firms that supply inputs to oil producers react to a negative oil price shock by switching their resources from dirty to clean research and development (R&D) activities.[1]

While there is an extensive literature investigating how clean R&D responds to changes in energy prices, the focus has generally been on the effects of *increased* fossil energy prices, which strengthens incentives for clean innovation through higher demand for renewable energy and clean technologies. What has received less attention in the literature to date, is whether *reduced* fossil energy prices may induce increased investments in clean R&D through reallocation of resources in the fossil energy supply chain. Moreover, there is little empirical evidence on the role of within-firm reallocation in shifting from fossil to clean technologies.

Figure 1: An oil price shock and the share of firms with clean R&D

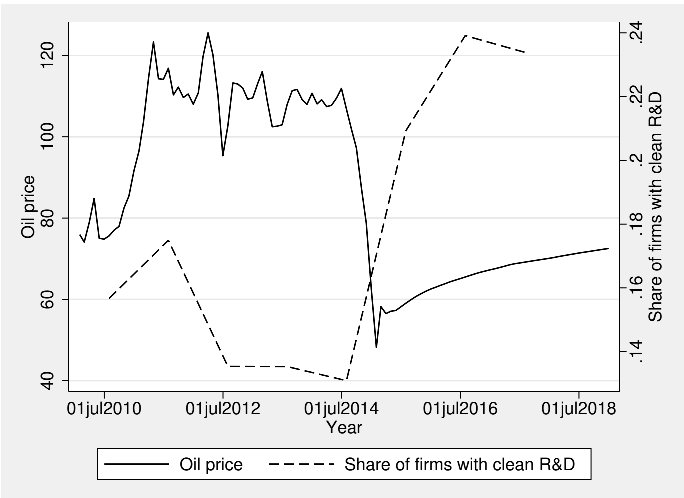

Note: The table shows the share of firms in the baseline sample that reports investments in clean R&D, the evolution of the oil price until Q1 of 2015, and the oil future price for the remaining time period. The sample covers the years 2007 to 2017.

Our paper is motivated by a striking empirical finding, shown in Figure 1. The figure depicts a sharp increase in the share of firms with clean R&D following a fall in the oil

[1] Clean R&D refers to R&D expenditures directed towards renewable energy or other environment-related energy. More details can be found in Section 2.2.

price. The main contribution of the paper is to show that lower profitability in the fossil energy sector can lead to reallocation of R&D resources within firms in the supply chain and trigger a shift towards clean innovation.[2]

We start by developing a theoretical framework building on the literature on directed technical change, and in particular on Acemoglu et al. (2012). In Acemoglu et al. (2012), a fall in the fossil energy price will give an incentive for firms to increase their clean R&D activity. We introduce within-firm dynamics by adding two key features to the model, which give rise to additional responses to price changes. These extensions are motivated by three stylized facts based on the rich data set at hand: First, in the aftermath of a large negative shock to the oil price in 2014, the share of firms in the oil producers' supply chain that conducts clean R&D activity increases sharply. Second, there is strong persistence in R&D activities at the firm level over time.[3] Third, the majority of the increase in clean R&D after the shock comes from firms that also have dirty R&D investments. Summarizing, these facts suggest that there is a non-trivial relationship between clean and dirty R&D, and that within-firm dynamics may play a role in explaining firms' reactions to changes in relative energy prices.

In the first extension to the model, we formalize the idea of adjustment costs for R&D activity proposed by Bloom et al. (2013b). The presence of adjustment costs implies that in the aftermath of a negative shock to expected returns to dirty innovation, firms that are more exposed to the shock will face a stronger incentive to divert their R&D activities towards clean innovation. In the second extension, we allow for spillovers from the more mature, i.e., dirty, R&D activity towards clean R&D. This extension works in the opposite direction: After a negative shock to the returns to dirty innovation, the presence of spillovers will incentivize firms exposed to the shock to scale down both dirty and clean R&D. The theoretical model allows us to develop a set of predictions which we take to the data to determine which of these forces dominate. We do so by studying the impact of a negative shock to fossil energy prices on clean R&D activity in the supply chain of fossil energy producers.

Our empirical analysis exploits a large, negative shock to the current and future price of oil, as oil prices dropped by about 60 per cent in the latter half of 2014. This quasi-natural experiment allows us to study the dynamics of resource reallocation within firms that supply the oil industry. We use detailed micro data from Norway. The oil industry

[2]Examples of green technology areas where capabilities and knowledge originating in the oil sector are particularly relevant include wind and geothermal energy. As an example, the U.S. Department of Energy launched an initiative in September 2024 called "Geothermal Energy from Oil and Gas Demonstrated Engineering", specifically directed at "leveraging the extensive knowledge, technology, skill, and experience of the oil and gas sector" to develop the geothermal industry (see https://www.geode.energy for more information).

[3]This finding is in line with existing empirical evidence, see e.g. Bloom (2007).

holds a central role in the Norwegian economy, and accounted for around 60 per cent of Norwegian exports and 14 per cent of private sector employment in 2014.[4]

To estimate the causal effect of the oil price shock on upstream firms' clean R&D investments, we propose a novel method for identification where we rely on firm-level trade data to compute a measure of firm exposure to the oil price shock. This firm-level exposure measure allows us to compare firms in the same industry that have similar characteristics at baseline, but that are differentially exposed to the shock due to their product mix. We find that firms that are more exposed to the 2014 fall in oil prices increase their clean R&D efforts more than firms that are less exposed to the shock. This is true if we look at the likelihood of investing in clean R&D, the share of clean R&D, or the value of clean R&D expenditures. We also document that more exposed firms have higher probabilities of starting new investments in clean R&D, and of switching from non-clean to clean R&D.

Aiming to understand the mechanisms at work, we investigate whether the firms we identify as being more exposed to the negative oil price shock experience subsequent decreases in sales and profits. This is indeed the case, as sales per employee decreases, as does profits. We also explore whether the exposed firms simply intensify their R&D investments across the board, but this is not the case. We find no effect on total R&D investments, the number of R&D employees or on the share of R&D employees in total employment. Our empirical findings support the hypothesis that within-firm reallocation of R&D investments lead upstream firms already engaged in dirty innovation to redirect their innovation efforts towards clean technologies in response to a negative shock to future profits in the downstream industry.

Our findings imply that a negative shock to profitability in the fossil energy sector may propagate through the supply chain and lead to relatively stronger reallocation of resources from fossil to clean R&D among the exposed firms. Based on our theoretical framework, we argue that the results may be driven by within-firm adjustment costs in R&D activity. Our results have implications for the understandings of the effects of climate policy in general, and carbon pricing, specifically. As carbon pricing will lower fossil energy sector profitability, our findings suggest an additional channel through which carbon pricing will shift resources from fossil to clean R&D activity.

The paper relates to several strands of literature. The theoretical framework builds directly on the well-established literature on directed technical change and climate, in particular Acemoglu et al. (2012). Dechezleprêtre and Hémous (2022) provide a recent review.[5] We add to this literature by including within-firm reallocation of R&D invest-

[4] See https://www.norskpetroleum.no/en/production-and-exports/exports-of-oil-and-gas/ and Prestmo et al. (2015).

[5] Acemoglu et al. (2012) underline the key role of R&D subsidies in climate policy, in addition to

ments, in a model allowing for innovation to take place both in clean and dirty technology at the same time. Our paper shares the feature of parallel innovation in both sectors with Fried (2018), who presents a dynamic general equilibrium model with endogenous innovation, which is calibrated using historical oil shocks. Our paper complements Fried's in several ways. First, by modeling a specific channel through which a negative oil price shock induces within-firm reallocation of R&D investments. Second, by providing a well-identified empirical investigation of the effects of a negative shock to the expected future profits from oil extraction.

On the empirical side, the paper contributes to the literature investigating the relationship between relative energy prices and clean innovation. The findings of Newell et al. (1999) indicate that higher energy prices will induce more rapid development of energy-saving technologies. Similarly, Popp (2002) finds that higher energy prices are associated with a significant increase in energy-saving innovations at the industry-level. Aghion et al. (2016) show that firms in the auto industry tend to innovate more in clean and less in dirty technologies when they face higher tax-inclusive fuel prices. A similar pattern, with a positive relationship between energy prices and energy-saving innovation, is also found by Crabb and Johnson (2010), Calel and Dechezleprêtre (2016), Ley et al. (2016) and Hu et al. (2023), while Acemoglu et al. (2023) shows how lower natural gas prices are generally negatively related to clean innovation activity. Dugoua and Gerarden (2024), Dugoua and Dumas (2021) and Dechezleprêtre and Kruse (2022) provide recent contributions to the literature on energy prices, environmental policies and clean innovation.[6] The main focus of the existing empirical literature has been on the demand-side effect of higher fossil energy prices on clean innovation. We focus on the supply side effect of lower fossil energy prices, a channel which has received less attention in the literature to date.

The paper also contributes to the literature on the role of carbon taxation in the green transition.[7] We show that carbon pricing may have a positive impact on the development

a carbon price, when there is path dependency in technical change. Acemoglu et al. (2016) estimate a similar model using micro data from US energy markets and conclude that, although the carbon tax plays an important role, subsidies for clean R&D are key to achieving the shift from dirty to clean technologies. Similar conclusions are drawn in Greaker et al. (2018). The key role of R&D subsidies in optimal climate policy is also underlined by Casey (2024). Hart (2019) use a model structure closer to that normally used in integrated assessment models of climate change (see e.g. Golosov et al. (2014)), resulting in a more important role for the carbon price. By allowing for complementarities between technology and energy inputs in production, Lemoine (2024) reaches a similar conclusion with regards to the importance of sufficient carbon pricing.

[6] See Popp et al. (2010) and Popp (2019) for reviews of the literature on environmental policy and innovation more broadly.

[7] See Golosov et al. (2014) for a key contribution to the literature on optimal carbon pricing, and Timilsina (2022) and Köppl and Schratzenstaller (2023) for reviews of this literature. Important contributions to the broader literature on optimal climate policy are, among others, Gerlagh et al. (2009), Stern and Valero (2021) and Blanchard et al. (2023).

of clean technologies not only through increased demand for these technologies, but also through increased supply, as profitability falls in the fossil energy sector.

Lastly, the paper speaks to the literature on shock propagation in networks. Barrot and Sauvagnat (2016) identify firm-specific shocks among supplier firms, and find that their domestic buyers experience substantial losses as a consequence. Dhyne et al. (2022) focus on the opposite direction and examine how suppliers are affected when their buyers are hit by shocks. Most of this literature is focused on short-run effects of short-run shocks; we contribute by focusing on a large and persistent shock and its effect on investments.

The paper is organized as follows. In Section 2, we describe the context and data used for the empirical analysis, and present a set of stylized facts. In Section 3, we present the theoretical model of directed technical change with within-firm adjustment costs and technology spillovers. The empirical strategy is described in Section 4, while the empirical results are presented in Section 5. In Section 6, we provide some concluding remarks.

2 Data and Stylized Facts

In this section, we first describe the large, negative oil price shock that occurred in the second half of 2014. We will exploit this large decline in the profitability of fossil energy production to identify the consequences for investments in clean R&D.[8] We then present the rich data at hand, and document three stylized facts, which guide our theoretical model and motivate our empirical design.

2.1 The oil shock

Our analysis relies on a quasi-natural experiment created by the large and sudden drop in the oil price. After peaking at \$107.95 a barrel on June 20 2014, the oil price plunged to \$44.08 a barrel by January 28, 2015, a drop of 59.2 percent in a little over 7 months. The main reason was a combination of an oversupply in the market and low global demand for oil. The increased production of shale oil in the US increased substantially, while the OPEC countries contributed to the oversupply by maintaining high production levels despite the falling price. At the same time, the economic slowdown in China and Europe reduced the global demand for oil, further exacerbating the supply-demand imbalance.

The result was a large, negative shock to the current and future price of oil. Figure 2 shows the development of the price of brent oil over the time period we study. As can be seen from the graph, oil prices plummeted starting in June 2014. Firms did most likely

[8]Lorentzen (2024) studies the impact of the same shock on sectoral reallocation of workers and subsequent effects on earnings in destination sectors.

not expect this dramatic decline to happen. In Figure 3 the dotted line shows the future price of oil in Q2 2014, just before the price drop started. It shows that the price was expected to remain high for the coming years. The dashed line on the other hand, shows the future price of oil as estimated in Q4 of 2014, and clearly illustrates that after most of the shock had materialized, markets anticipated that prices would remain low for the foreseeable future.

The drop in future prices is potentially important for our analysis, as the decisions firms are making about investments in R&D are forward-looking in nature, and firms are unlikely to respond much to temporary fluctuations in the price. The oil price is set in the world market, and is completely exogenous from the perspective of firms in a small open economy. Hence, the 2014 oil price drop serves an ideal quasi-natural experiment that allows us to study resource reallocation within firms that are exposed to the shock through their supply linkages to the oil extraction industry.

The fall in the oil price in 2014 had a substantial negative impact on expected output growth in the firms in Norway supplying the oil extraction industry. Figure 4 reports the results from the business sector survey conducted regularly by the Norwegian Central bank for the period January to August 2015. According to the survey, expectations varied significantly across industries, with a positive outlook in most sectors, while the industries serving and supplying the oil industry stand out, with a very negative outlook for both domestic sales and exports.[9]

[9]It should be noted that the technology, intermediates and capital goods used for offshore extraction of oil differs substantially from the technology used for onshore oil extraction. It is therefore unlikely that the firms supplying the Norwegian offshore industry experienced increased demand for their product as a consequence of the positive shale oil shock in the US. We can see this reflected in Figure 4.

Figure 2: Oil price: Brent Blend

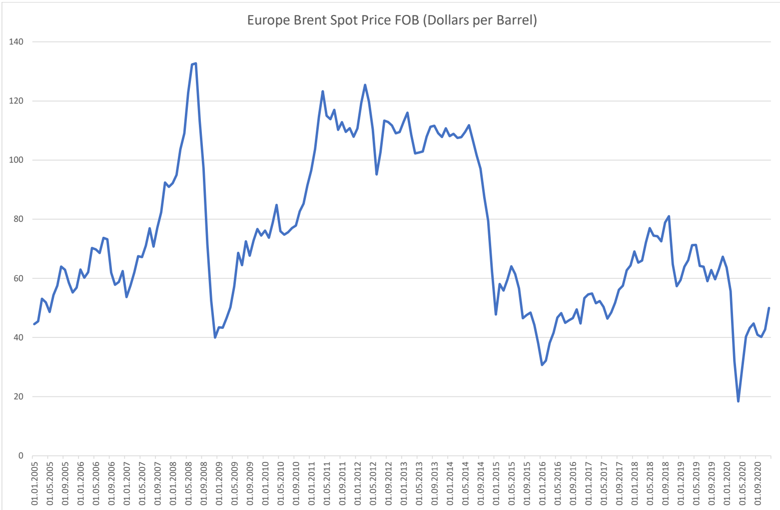

Note: The figure plots the European Brent Spot Price for oil per quarter for the period for the period 2005 to 2020. Source: US Energy Information, https://www.eia.gov/dnav/pet/hist/LeafHandler.ashx?n=PET&s=RBRTE&f=D.

Figure 3: Crude Oil Price and Oil Future Prices

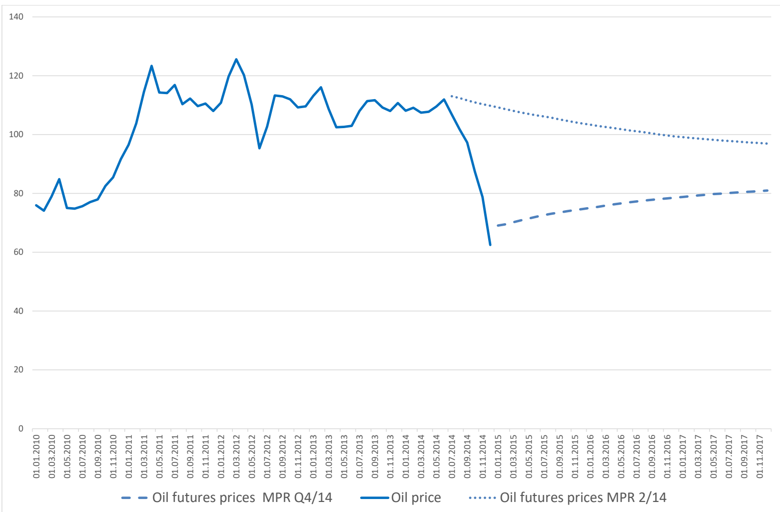

Note: The figure plots the Crude Oil Price until end 2014 and Oil Future Prices from 2014 to 2017. Source: The Norwegian Central Bank Monetary Policy Reports 2/2014 and 4/2014.

Figure 4: Expected Output Growth

Note: The figure shows expected annual output growth the coming six months by sector in the Norwegian economy based on the Norwegian Central bank regular survey of the business sector. Source: The Norwegian Central Bank Monetary Policy Report 3/2015.

2.2 Data

Our empirical analysis is based on five data sets. The first data set is administrative firm register data from Statistics Norway, which covers the universe of firms across all sectors. The firm register provides information on the date of the entry and exit of each individual firm, allowing us to calculate the firm's age. The register also holds data on firms' number of employees.

The second data set is income statement and balance sheet data from Statistics Norway for all private non-financial joint-stock companies. Since 90 percent of Norwegian firms with one or more employees are joint-stock firms, this means that it covers almost the entire universe of firms in Norway. The income statement and balance sheet data are based on data from annual accounting reports that, according to Norwegian law, must be filed with the public Register of Company Accounts.

The third data set is the R&D survey. The survey provides information about the value of R&D expenditures, the number of and the wage bill of R&D personnel. Importantly, the survey also gives information about the share of R&D related to clean energy. Specifically, firms are asked to report what percentage of R&D expenditures are directed towards renewable energy or other environment-related energy. We combine these two categories to create a measure of clean innovation.[10] More detailed descriptions of the R&D data can be found in the Appendix, Section A.

The fourth data set we use is product-level trade data for the universe of firms in Norway. The data is at the HS8 product-level and covers all goods that are imported or exported by Norwegian firms.

The fifth data set provides complete information on all direct support from the government for R&D and innovation to firms.[11] We link all data sets with a unique firm identifier.

2.3 Sample selection

For our analysis, we focus on all joint stock firms in manufacturing (defined by the NACE industries #10 to 35) that are covered by the R&D survey. We do not include the firms that were directly hit by the fall in oil price, namely the oil producers. This gives us an unbalanced panel of approximately 1,300 firms per year. We focus on the manufacturing industry, as these firms are responsible for the vast majority of trade in goods. We use information about which goods are imported by oil producers (NACE 06) to construct our exposure measure as described below in Section 4.1. Our sample is constructed to cover the years 2007 to 2017. The years 2007-2013 define the pre-shock period, and 2014-2017 define the post-shock period.

2.4 Stylized facts on R&D and clean R&D

We present three stylized facts on the development of R&D and clean R&D in the manufacturing industry in the aftermath of the oil price shock that motivates our analysis and which will guide our theoretical model.

Fact 1: After the oil price shock, the share of firms investing in clean R&D increased both in absolute and relative terms. In 2013, the share of firms with positive R&D in our sample was 40 percent and the share of firms investing in clean R&D was only 5 percent. Fast-forward to 2017, and the numbers have increased to 46 percent for overall R&D, whereas the share of firms investing in clean R&D has more than doubled, to 11 percent. Figure 5 shows the increase in the share of firms with clean R&D expenditure relative to all firms engaged in R&D.

[10] Firms are also asked about the percentage of R&D expenditures that are directed towards energy efficiency. We choose to leave this category out, as it is likely to capture process innovations related to reducing the firm's own emissions as well as innovation related to technologies that will reduce emissions of the downstream firms.

[11] For more details see https://www.ssb.no/en/teknologi-og-innovasjon/forskning-og-innovasjon-i-naeringslivet/statistikk/naeringspolitiske-virkemidler.

Figure 5: Share of R&D firms with clean R&D

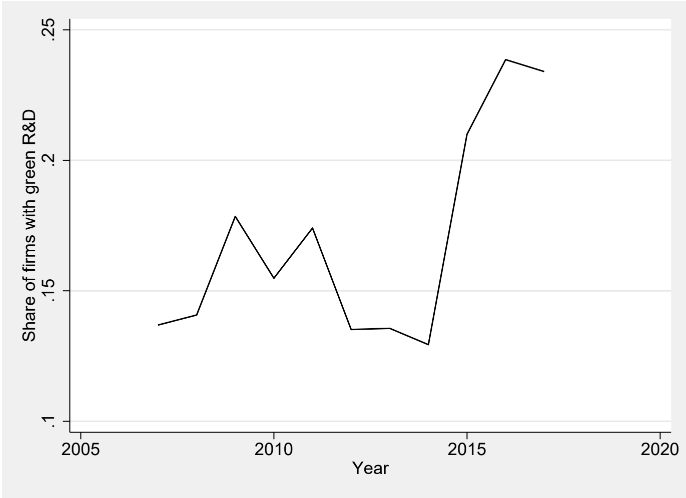

Note: The table shows the share of firms that report investments in clean R&D relative to all firms engaged in R&D.. The sample covers the years 2007 to 2017.

Fact 2: The majority of the increase in clean R&D comes from firms that have ongoing investments in non-clean R&D. Figure 6 shows how the aggregate change in clean R&D can be decomposed into 3 categories: firms starting to do R&D (including new entrants), firms who stop doing R&D (including exiting firms), and continuing R&D performers. For each of these categories, we split firms into two groups, depending on whether their R&D investments are 100 percent clean or a combination of clean and non-clean. As the graph shows, the majority of the increase in clean R&D after the oil price shock comes from firms with both clean and non-clean R&D.

Figure 6: Decomposition of Clean R&D

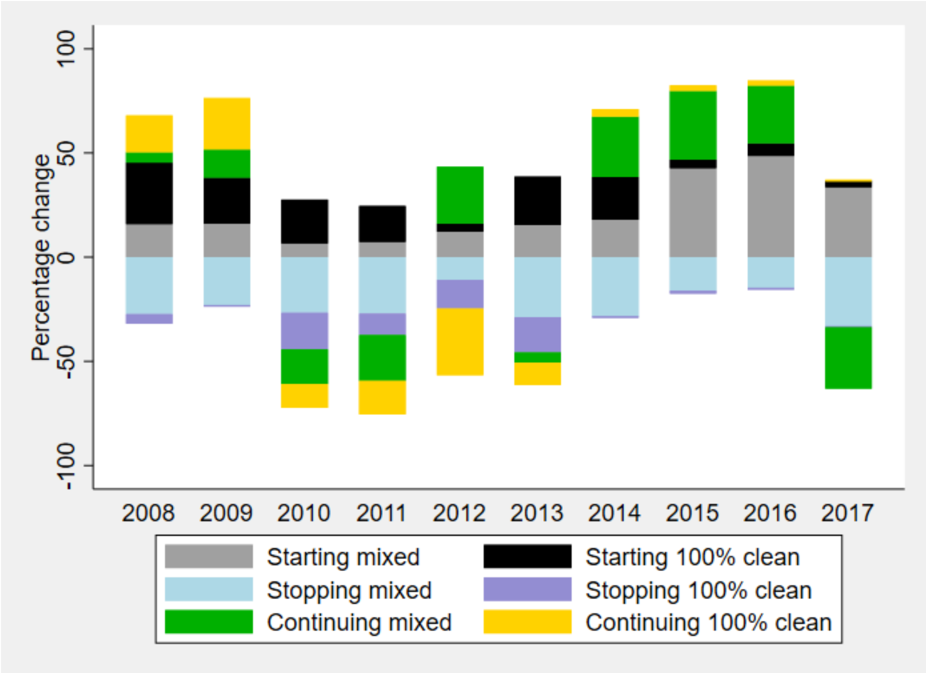

Note: The figure shows the annual percentage change in clean R&D – decomposed according to the status of the firm, i.e. whether the firm starts, stops or continues investment in both clean and dirty R&D ("mixed"), or starts, stops or continues investment in just clean R&D. The sample covers the years 2007 to 2017.

Fact 3: There is persistence in R&D activities at the firm-level. A firm that has positive R&D expenditures in one period, has a 90 percent probability of continuing in the subsequent period. This inertia in R&D activities is in line with the empirical findings of Bloom (2007), who finds that R&D investments are highly persistent across business cycles, and supports the hypothesis that the rescaling of R&D activity is subject to adjustment costs.

Table 1: Annual transition rates

Status year t	Status year t+1	
	No R&D	R&D
No R&D	0.862	0.138
R&D	0.098	0.902

Note: The table shows transition rates for firms in the baseline sample and covers the years 2007 to 2017.

These stylized facts indicate that within-firm dynamics in R&D activity may be relevant for understanding reallocation of resources between clean and dirty R&D in response to changes in relative prices. Based on these observations, in the next section we introduce within-firm spillovers and adjustment costs into a framework of directed technical change.

3 A Theoretical model of Directed Technical Change

In this section, we develop a stylized model of directed technical change, to guide our empirical analysis of the impact of a permanent drop in the price of fossil fuels on investments in clean R&D. We build on Acemoglu et al. (2012), and develop a model that allows us to investigate how a negative price shock to a dirty final good may propagate through the supply chain and induce a diversion of R&D investments from dirty to clean R&D among input producers. The model allows us to illustrate how firms supplying inputs to the fossil energy industry may react differently to a drop in the oil price as compared to firms that do not deliver their inputs to the fossil fuel producers.

Informed by the three stylized facts presented above, we extend the model by Acemoglu et al. (2012) by allowing firms to differ in their innovation productivity with respect to dirty and clean R&D, implying that firms will make heterogenous choices regarding R&D investments. Moreover, we introduce two types of within-firm dynamics in the model. First, we allow for within-firm spillovers from the more mature R&D activity (dirty) to the less mature (clean) activity. [12]Second, in line with Bloom et al. (2013b), we include adjustment costs in the R&D activity within the firms. We model adjustment costs by including a convex cost of rescaling the size of the R&D activity within the firm.

3.1 The Model

Consider a discrete-time, infinite-horizon economy where two final goods, $j = c, d$, are produced in each time period. The clean good, denoted by subscript c, in quantity Y_{ct}, and the dirty good, denoted by subscript d, in quantity Y_{dt}. We consider a small, open economy and therefore the final-good prices are exogenously given.

Each final good is produced by competitive producers using a continuum of inputs that are specific to the type of final good produced, with the production technology:

$$Y_{jt} = \int_0^1 A_{jit}^{1-\alpha} x_{jit}^{\alpha} di, j = c, d, \tag{1}$$

with $\alpha \in (0, 1)$ where x_{jit} is the quantity of input ji . A_{jit} denotes the productivity of

[12]Between-firm spillovers are well-documented in the literature, see specifically Bloom et al. (2013a). Many of the mechanisms proposed in this literature are likely present also within firms. Nix (2020) proposes mechanisms that may lead to a similar pattern within firms.

the input. We abstract away from the use of other inputs such as labor in the final good sector, in order to keep the framework as simple as possible.

There is a mass of 1 firms that can produce inputs. The unit cost in production of inputs is constant and independent of both type and productivity of the input, given by ϕ. Without loss of generality, we follow Acemoglu et al. (2012) and let $\phi = \alpha^2$.

Innovation takes place within the input-producing firms. Define A_{jt} as the aggregate state of the technology for the final good type j at time t across all firms producing inputs, with:

$$A_{jt} \equiv \int_0^1 A_{jit} di. \tag{2}$$

A successful innovation for input x_{jit} will increase the productivity of the input, A_{jit}, by a common factor $(1 + \gamma) > 1$. Let subscript k denote a firm k. When an innovation for input ji happens in firm k, the firm obtains a one-period patent and is allocated the monopoly rights for input x_{ji} for the current time period. We follow Acemoglu et al. (2012) and assume that for all inputs with no successful innovation, the monopoly rights are allocated to a random firm.

To innovate, the firms must hire scientists to conduct R&D activity. There is a mass of 1 scientists available in the market, and they have skills that can be used for both clean and dirty R&D. A firm hiring a scientist chooses whether to put a scientist to work on dirty or clean R&D. A scientist that is hired for type-j R&D activity is randomly assigned to an input of that type. The input producers act competitively in the market for scientists. Let s_{kjt} be the number of scientists hired by firm k for type-j R&D activity in period t, and let $s_{kt} = s_{kdt} + s_{kct}$, be the total number of scientists hired by firm k in period t.

Firms are heterogeneous and the probability of innovation of type j in firm k depends on a firm and final good specific innovation productivity, $\eta_{kj} \in \{0, \eta_j\}$ with $\eta_j \in (0, 1)$, and on a concave innovation production function $g^j(\cdot)$. We allow for within-firm technological spillovers from the more mature (dirty) R&D activity towards the less mature (clean) R&D activity. For the relevant range of s_{kct} and s_{kdt}, the clean-innovation production function is therefore given by $g^c(s_{kct}, s_{kdt})$,with $g^c(0, s_{kdt}) = 0$, $g^c(1, 1) \leq 1$, $\partial g^c/\partial s_{kct} > \partial g^c/\partial s_{kdt} \geq 0$, $\partial^2 g^c/\partial s_{kct}^2 < 0$, $\partial^2 g^c/\partial s_{kdt}^2 = 0$ and $\partial^2 g^c/\partial s_{kct}\partial s_{kdt} \geq 0$, while all third derivatives are zero.

For dirty innovation, only the number of scientists working on dirty R&D is relevant, since there are no technological spillovers from the less mature towards the more mature R&D activity. The dirty-innovation production function is given by $g^d(s_{kdt})$, with $g^d(0) =$

0, $g^{d\prime}(\cdot) > 0$, $g^{d\prime\prime} < 0$, $g^{d\prime\prime\prime} = 0$ and $g^d(1) \leq 1$.[13]

Firms will not engage in R&D of type j if their innovation productivity for that type is zero. We can divide the input producers into four groups, depending on η_{kc} and η_{kd}:

- Group I: $\eta_{kc} = \eta_c$, $\eta_{kd} = 0$. Engage only in clean R&D.
- Group II: $\eta_{kc} = 0$, $\eta_{kd} = \eta_d$. Engage only in dirty R&D.
- Group III: $\eta_{kc} = \eta_c$, $\eta_{kd} = \eta_d$. Engage in both types of R&D.
- Group IV: $\eta_{kc} = 0$, $\eta_{kd} = 0$. Are not active (can be disregarded).

Let $\lambda^l \in [0, 1]$ denote the share of firms that belong to group l, with $\lambda^I + \lambda^{II} + \lambda^{III} + \lambda^{IV} = 1$.

We further assume that the rescaling of R&D activity is subject to adjustment costs. We think of these adjustment costs as capturing not only direct costs of firing scientists, but also other economic and behavioral responses that might arise within a team of highly specialized workers when downscaling is considered. For instance, managers might be reluctant to let go of scientists with specialized skills that will be hard to replace. They may also prefer to avoid disrupting the dynamics within a team and rather keep a well-composed team intact. This "team cohesion effect" occurs when the collective performance of a team is greater than the sum of its individual members' performances, and removing one or more of the scientists will lead to productivity losses for the remaining team members. Managers might also be worried about a "layoff contagion effect", where laying off one scientist will make colleagues worried about their own job security.[14] Managers might also care about their own careers and be reluctant to let their own teams shrink, even if that contradicts the best interest of the firm. To reflect such adjustment costs, we include a convex cost of rescaling the size of the R&D activity within the firm.[15] To keep the framework as simple as possible, the adjustment cost is separable from the innovation productivity and given by:

$$R_k = c \cdot \left(\frac{s_{kt} - s_{kt-1}}{s_{kt-1}} \right)^2, \tag{3}$$

[13] To ensure an interior solution with respect to s_{kjt} to the problem of a firm with $\eta_{kj} = \eta_j$, we assume $\lim_{s_{kdt} \to 0} g^d(s_{kdt}) = \infty$ and $\lim_{s_{kct} \to 0} g^c(s_{kct}, s_{kdt}) = \infty$.

[14] A related argument is that a team of workers may be of greater value than the sum of its individuals in the labor market, as pointed to by a recent literature on co-mobility of workers and "aqui-hiring" (see for example Marx and Timmermans (2017).) If managers take this into account, it might make them less willing to lay off part of a team as it might make the rest of the team less productive, or induce them to leave.

[15] The assumption of convex adjustment costs is standard in the labor market literature, see e.g. Hamermesh (1995)

for firm k, with $c > 0$. Note that a key to the results presented in the following is that the cost of rescaling depends on the overall level of R&D activity in the firm, not on whether the activity is related to clean or dirty inputs.

The aggregate productivity of type-j inputs will develop over time according to:

$$A_{ct} = \left(1 + \gamma \int_0^1 \eta_{kc} g^c(s_{kct}, s_{kdt}) dk\right) A_{ct-1} \tag{4}$$

$$A_{dt} = \left(1 + \gamma \int_0^1 \eta_{kd} g^d(s_{kdt}) dk\right) A_{dt-1} \tag{5}$$

with $\int_0^1 s_{kt} = 1$ in equilibrium. Allocation of scientists will therefore drive the technological state of the economy in the long run. In the following, we make one further simplifying assumption, namely that all input producers are myopic in the sense that they do not take future adjustment costs into account when they make their R&D investment decision.[16] The value of potential innovation as a result of investment in R&D arises in the current time period, and is of course included in the firms' objectives.

3.2 Market Equilibrium

Let p_{jt} denote the exogenous price of final good j in period t, while p_{jit} denotes the price of input ji. The producer of final good type j will choose inputs to maximize profits:

$$\max_{x_{jit}} p_{jt} \int_0^1 A_{jit}^{1-\alpha} x_{jit}^{\alpha} di - \int_0^1 p_{jit} x_{jit} di. \tag{6}$$

Solving the maximization problem gives the downward-sloping demand for each input:

$$p_{jit}(x_{jit}) = \alpha p_{jt} A_{jit}^{1-\alpha} x_{jit}^{\alpha-1}. \tag{7}$$

The monopolistic producer of good x_{jit} will maximize profits subject to the final good producers' demand for inputs. In equilibrium, production of, and profits related to, input ji are given by

$$x_{jit} = A_{jit} p_{jt}^{\frac{1}{1-\alpha}} \tag{8}$$

$$\pi_{jit} = \alpha(1-\alpha) A_{jit} p_{jt}^{\frac{1}{1-\alpha}} \tag{9}$$

[16] Forward-looking firms will take into account not only the current adjustment costs, but also those expected to be incurred in the future. For a firm adjusting its R&D up or down, the expected future adjustment costs provide an incentive to rescale somewhat quicker, relative to the optimal rescaling for a myopic firm. However, as long as the future adjustment costs are discounted, the key effect of including the adjustment costs on the firms' decisions is captured in the myopic problem presented in the following. Hence, we make the assumption about myopic behavior in order to keep the framework as simple as possible.

The expected within-period revenue from hiring a scientist for type-j R&D activity for firm k in period t, denoted I_{kjt}, is given by:

$$I_{kct} = \eta_{kc} g^c(s_{kct}, s_{kdt}) \int_0^1 \pi_{cit} di = \eta_{kc} g^c(s_{kct}, s_{kdt}) \alpha(1-\alpha) p_{ct}^{\frac{1}{1-\alpha}} (1+\gamma) A_{ct-1} \tag{10}$$

$$I_{kdt} = \eta_{kd} g^d(s_{kdt}) \int_0^1 \pi_{dit} di = \eta_{kd} g^d(s_{kdt}) \alpha(1-\alpha) p_{dt}^{\frac{1}{1-\alpha}} (1+\gamma) A_{dt-1} \tag{11}$$

Because scientists are randomly assigned to a variety within the type they are hired for, it is the aggregate input quality for that type, A_{jt-1}, that determines the expected revenue from hiring.

Scientists are paid wage w_t, and firm k thus hires scientists to solve the following problem in period t:

$$\max_{s_{kdt}, s_{kct}} \left\{ \alpha(1-\alpha)(1+\gamma) \left(\eta_{kc} p_{ct}^{\frac{1}{1-\alpha}} A_{ct-1} g^c(s_{kct}, s_{kdt}) + \eta_{kd} p_{dt}^{\frac{1}{1-\alpha}} A_{dt-1} g^d(s_{kdt}) \right) - w_t s_{kt} - c \cdot \left(\frac{s_{kt} - s_{kt-1}}{s_{kt-1}} \right)^2 \right\}, \tag{12}$$

subject to $s_{kt} = s_{kct} + s_{kdt}$.

We let the resulting demand for scientists from a firm k in group I, which only conducts clean R&D activity, be given by $s^I_{kct}(p_{ct}, w_t)$. Demand for scientists from firm k in group II, which only conducts dirty R&D activity, is given by $s^{II}_{kdt}(p_{dt}, w_t)$, and demand for scientists from firm k in group III, which conducts both types of R&D activity, is given by $s^{III}_{kct}(p_{ct}, p_{dt}, w_t)$ and $s^{III}_{kdt}(p_{ct}, p_{dt}, w_t)$. The fourth group of firms does not conduct any research and does therefore not demand any scientists. Equilibrium in the market for scientists requires that the demand for scientists in each period t from each group of input producers is equal to the supply of scientists, which we have set to one:

$$\lambda^I s^I_{kct}(p_{ct}, w_t) + \lambda^{II} s^{II}_{kdt}(p_{dt}, w_t) + \lambda^{III} s^{III}_{kct}(p_{ct}, p_{dt}, w_t) + \lambda^{III} s^{III}_{kdt}(p_{ct}, p_{dt}, w_t) = 1 \tag{13}$$

We note that group I is not directly affected by a change in the price of dirty goods, while group II is not directly affected by a change in the price of clean goods.

3.3 Comparative Statics

Let us now consider the effect of a fall in the price of the dirty final good on clean innovationin upstream firms. Specifically, we want to compare the effect for a firm in group III, that is engaged in both clean and dirty R&D, to that of a firm in group I, that is only engaged in clean R&D. Because we want to compare otherwise similar firms, assume now that all firms have the same total number of scientists hired in the previous period, i.e. $s_{kt-1} = s_{t-1}$ for all k. To get at these effects, we must consider all general-equilibrium effects in the model.

We show in the Appendix Section B that the direct effect of a fall in the price of the dirty final good on demand for scientists is negative for the upstream firms in group II and III, which supply the producers of the dirty final good with inputs: Their revenue falls, and returns to dirty R&D innovation fall. As a consequence, these firms will reduce their dirty R&D activity and thus their demand for scientists, which in turn leads to a decline in wages (see equations 29, 32 and 37).The decline in wages will induce all firms that are active in R&D, including the firms in group I, to expand their R&D activity (see equations 33-36).

Firms in group III are active in both dirty and clean R&D. The impact of a price shock for dirty goods on their clean R&D activities depends both on the wage effect, which is also experienced by firms in group *I*, as well on within-firm allocation of R&D investments related to technological spillovers and adjustment costs: The presence of technological spillovers implies that reduced dirty R&D activity leads to lower marginal productivity of scientists working on clean innovation. This in turn leads to reduced clean R&D activity. Adjustment costs, on the other hand, imply that these firms increase the resources they allocate towards clean innovation, relative to firms in group I, because the net marginal cost of their clean R&D activity decreases as they downscale their dirty R&D. In other words, the adjustment costs make firms shift some of their scientist from dirty to clean R&D rather than firing them. The effect of a change in the price of the dirty final good on the demand for scientists for clean R&D activity in group III depends on what effects that dominate.

The following two propositions summarize our key conclusions regarding the relative impact of a price shock for firms that only conduct clean R&D (group I) and firms active in both clean and dirty R&D (group *III*), using the case without spillovers and adjustment costs as a benchmark:

Proposition 1. *When there are no within-firm technological spillovers, and there are no adjustment costs, a drop in the price of the dirty final good, p_{dt}, will induce the same change in clean R&D in all firms engaged in clean R&D, independently of whether the firms are also active in dirty R&D or not:*

$$\frac{ds^{I}_{kct}}{dp_{dt}} = \frac{ds^{III}_{kct}}{dp_{dt}}. \tag{14}$$

Proof. All calculations and expressions are provided in the Appendix, Section B, see specifically Equation 41. When $\partial g^c/\partial s_{kdt} = \partial^2 g^c/\partial s_{kct}\partial s_{kdt} = c = 0$, it follows that the direct effect on clean R&D, the first term in Equation 39, is zero for firms in both groups. Moreover, the indirect effect (the second term in Equation 39 and the full expression in Equation 40) is the same across the two groups. □

Proposition 2. *If there are within-firm technological spillovers and adjustment costs, and the effect of the adjustment costs dominates the spillover effect, firms engaged in both dirty and clean R&D will respond to a fall in the price of the dirty final good by increasing their clean R&D activity, compared to the firms only engaged in clean R&D:*

$$\frac{ds^{III}_{kct}}{dp_{dt}} > \frac{ds^{I}_{kct}}{dp_{dt}}. \tag{15}$$

If, on the other hand, the spillover effect dominates the effect of the adjustment costs, the firms with both dirty and clean R&D activity, will respond to the price fall by reducing their clean R&D activity relative to the firms only engaged in clean R&D activity:

$$\frac{ds^{III}_{kct}}{dp_{dt}} < \frac{ds^{I}_{kct}}{dp_{dt}}. \tag{16}$$

Proof. See the Appendix, Section B, specifically Equation 41.

In the next section, we take these propositions to the data. It is worth noting that in the theoretical framework, we deliberately focus on the intensive margin to avoid corner solutions. We do this to highlight the key mechanisms that we propose, which would hold also in a more complicated set-up allowing for corner solutions. In the empirics that follow, we will also consider the extensive margin, i.e. not limit the analysis to firms that were already conducting clean R&D before the shock, but also include firms that were only conducting dirty R&D before the shock. □

4 Empirical Strategy

The point of departure for our empirical analysis is the theoretical framework developed above. According to Proposition 2, the relative impact of an oil price shock on exposed firms (group III) relative to non-exposed firms (group I) is ambiguous. It thus remains an empirical question. To investigate this question we develop a novel measure of identification which lets us exploit firms' heterogeneous exposure to the shock in our analysis.

4.1 Identification: Measuring the exposure to the shock

To make progress towards estimating the causal effects of the oil price shock, we need to identify which firms are exposed to the shock. The standard approach in the literature has been to rely on sectoral input-output tables, which provide the share of production in an industry i that is bought by industry j. However, this would not be satisfactory for our purposes for two main reasons.

First, we are interested in the firm-level reaction to the shock, and we know from a large literature in international trade that firms are very heterogeneous even within narrowly defined industries. As we will describe below, our exposure measure shows that there is a lot of variation in exposure within industries. This heterogeneity NB that would be classified as either treated or control if we rely on sectoral input-output tables.

Second, if our treatment were defined at the industry-level, we would not be able to control for any other contemporaneous shocks to individual industries.

Therefore, we instead develop a novel firm-specific exposure measure leveraging the detailed trade data. We exploit the rich information we have about firm-level imports by oil producers to identify which products they use in their production, and construct our firm-specific exposure measure by following these steps:

1. Identify the HS8 products imported by the oil producing industry (NACE #6) in the pre-shock period (2007-2013).[17]

2. Identify the firms in manufacturing that export products that are imported by the oil producing industry ($j \in o$).

3. Calculate firm-level exposure x_{ok} as the share of products supplied to the oil industry in each firm's total export basket in 2013 as $x_{ok} = \sum_{j \in o} x_{kjt} / \sum_j x_{kjt}$ for $t = 2013$.

More details on the construction and underlying assumptions are given in the Appendix, Section C.

Measuring exposure to the shock in this way has some clear advantages compared to using sector-level input-output tables. First, it captures the fact that the oil price shock was a global shock. Compared to measuring the exposure only through domestic linkages, we are capturing the extent to which firms are exposed to the oil price drop through international trade. Second, and crucial for identification, it gives firm-level variation in exposure. Compared to using sectoral input-output linkages, using our exposure measure allows us to control for other, potentially time-varying, factors that are affecting all firms in an industry.

[17]We use the whole pre-shock period to account for the fact that firms' sourcing of capital goods will be subject to lumpiness over time.

Using our exposure measure does come at a cost, as it reduces the sample to only goods-exporting firms in the manufacturing sectors, and we are therefore not able to capture firms that trade in services. Moreover, our measure relies on the assumption that a firm's share of exports directed towards the oil producing industry is an appropriate measure of the share of the firm's activity devoted to these products.

Figure 7 highlights the importance of using a firm-specific measure. The box plot describes the within-industry variation in the exposure measure for all of our manufacturing industries. The median of the exposure measure is shown by the line within the box, and the box's boundaries represent the first (Q1) and third (Q3) quartiles, indicating the middle 50 percent of the data. The difference between Q3 and Q1 is the interquartile range (IQR), and the whiskers extend from the box to the smallest and largest values within 1.5 times the IQR from the quartiles, highlighting the range of the data. The points outside this range are plotted individually as outliers, indicating data points that differ significantly from the rest of the distribution.

The box plot shows that there is substantial variation in the exposure of firms within practically all of the manufacturing industries. An alternative to our approach would be to follow the commonly used classification from Statistics Norway, which defines industries with at least 50 percent of its value of production directed to the petroleum sector to be oil related.[18] This would in practice mean that we would have classified all firms in industries 25, 26, 27, 28, 30 and 33 as being affected by the shock, and firms in all other industries as not affected. We find it reassuring that these sectors also stand out as being strongly affected with our approach. However, as Figure 7 shows, there is substantial exposure also among firms in the other industries. In particular, NACE 22 (Manufacture of rubber and plastic products), NACE 29 (Manufacture of motor vehicles, trailers and semi-trailers), NACE 31 (Manufacture of furniture) and NACE 32 (Other manufacturing) stand out as having medians close to 1.

[18]See Statistics Norway for details: https://www.ssb.no/246994/naeringsundergrupper-i-standard-for-naeringsgruppering-sn2007-som-er-gruppert-som-petroleumsrettet-leverandorindustri-og-utvinningstjenester .

Figure 7: Within-industry Variation in Exposure

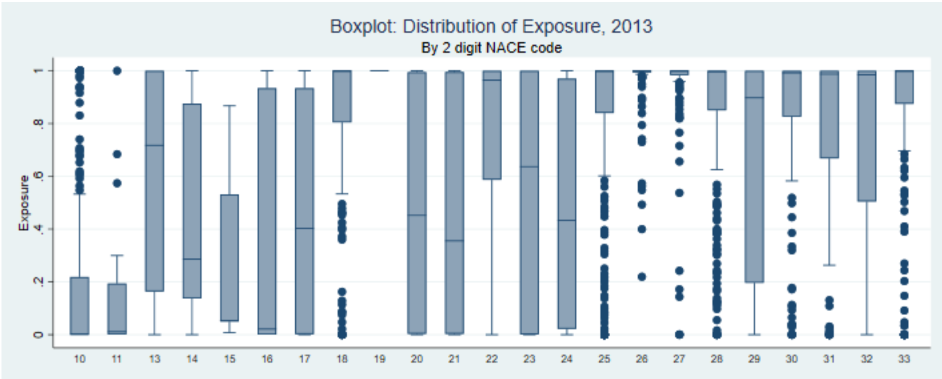

Note: The median of the exposure measure is shown by the line within the box, and the box's boundaries represent the first (Q1) and third (Q3) quartiles, indicating the middle 50 percent of the data. The difference between Q3 and Q1 is the interquartile range (IQR), and the whiskers extend from the box to the smallest and largest values within 1.5 times the IQR from the quartiles. The points outside this range are plotted individually as outliers.

4.2 Empirical Model

In order to estimate the impact of the oil price shock, we specify a difference-in-difference model based on the theoretical framework developed in Section 3. This model allows us to identify the causal effect of the oil price shock on clean R&D investments, as we compare firms that are differentially affected by the oil price shock before and after 2014. The baseline specification takes the form:

$$y_{kt} = \alpha_k + \beta x_{ok} \times Post_t + \gamma Z_{kt} + \delta_{st} + \varepsilon_{kt}, \tag{17}$$

where y_{kt} is an outcome variable related to the firm's R&D investments or performance. x_{ok} captures the firm's exposure to the shock based on the export portfolio (as described in Section 4.1).[19] α_k are firm fixed effects, capturing all time-invariant differences between the more exposed and less exposed firms. Crucially, as our exposure measure is constructed at the firm level, we are able to control for industry-specific trends, δ_{st}, at the NACE 2-digit level. Our main coefficient of interest is β, which captures how firms exposed to the shock respond in the years following the shock. $Post_t$ is a dummy variable that equals one in the the post-period ($t > 2013$). Z_{kt} is a vector of firm-level controls aimed at ensuring that the more exposed and less exposed firms are comparable. To account for differences across firms related to the size of their operations, we include log employment and log tangible assets as firm level controls. Firms may also have been

[19] To be precise, this is a difference-in-difference specification with a continuous treatment variable. We discuss the implications of having a continuous treatment variable below, in Section 5. Our approach is conceptually similar to a shift-share design, but with a one-off, common shock.

differently hit by the global oil price shock due to their exposure to trade, in particular due to a depreciation of the Norwegian exchange rate in the aftermath of the shock. To account for this, we include firms' export share (in total sales) as a further control.

As pointed out above, we limit our analysis to firms in the manufacturing sector. That is, we leave out the oil producers (NACE #6) that were directly affected by the oil price shock. The firms in our sample differ in their exposure to the shock depending on their supply linkages to the oil producers. They may, however, also differ in their exposure to the shock due to their use of energy as an input in production. Manufacturing firms in Norway typically rely on two main sources of energy: electricity based on hydro power, and fossil fuels for transport and some types of machinery. In the Appendix Section D, we show how firms' costs of energy have developed over time for each of these energy goods and in aggregate, see Figure 9. The main takeaway is that electricity prices developed smoothly over the relevant time horizon. Still, to control for potential differences in exposure to the shock on the cost side, we add firms' energy share, measured as cost of energy related to production and transportation relative to operational costs, to the specification. Finally, we also control for differences in state aid to R&D, by adding a dummy which captures whether or not the firm has received public funding for R&D and innovation.

As all of these control variables themselves are possibly affected by the oil price shock, we compute them based on the first year the firm is observed in the sample and interact them with year dummies. Standard errors are clustered at the firm-level.

5 Empirical Results on Clean Innovation

Our main outcome variables of interest are related to clean R&D activities. We start by analyzing the effect on the probability of doing clean R&D, as measured by a dummy variable that takes the value of 1 if the firm reports any clean R&D investment, and 0 otherwise. The results from estimating different versions of Equation (17) with and without controls and fixed effects can be found in Table 2.

In column (1), we only include industry-year fixed effects. We find that for a firm with average exposure, the shock leads to an increase in the likelihood of reporting any clean R&D investment of approximately 4.5 percentage points. The baseline likelihood of having any clean R&D investment is 7.5 percent, meaning that this effect is substantial and of economic importance. In column (2), we add firm fixed effects, which attenuates the coefficient slightly, but not by much. In column (3), we add our baseline control variables, except that we leave out the control for the firm's energy share in the production process. We add this control in column (4) to reach our preferred specification. We find that the energy share is the only control variable that seems to have some bite. An effect of including energy share in the regression is not surprising, as firms with substantial use of energy in production may clearly also be affected by the oil price drop through this direct channel. In column (5), we estimate our preferred specification using PPML instead of OLS.[20] Our results suggest that firms that where relatively more exposed to the oil price shock, due to their delivery of inputs to the oil producing sector, were more likely to invest in clean R&D after the shock.

Table 2: Probability of Clean R&D

Variable:	Dummy (1)	Dummy (2)	Dummy (3)	Dummy (4)	Dummy (5)
$Post_t{*}x_{ok}$	0.044**	0.039**	0.042**	0.055***	0.499**
	(0.019)	(0.020)	(0.020)	(0.020)	(0.199)
Controls ex. energy	No	No	Yes	Yes	Yes
Controls incl. energy	No	No	No	Yes	Yes
Firm FE	No	Yes	Yes	Yes	Yes
Industry*year FE	Yes	Yes	Yes	Yes	Yes
Estimator	OLS	OLS	OLS	OLS	PPML
Obs.	11,695	11,695	11,695	11,695	3,024

Note: Standard errors in parenthesis are clustered on firm. Controls include baseline levels of log employment, log tangible assets, export share, and a dummy for public funding, all interacted with year dummies. A control for energy share is included where indicated. $*p < 0.1, **p < 0.05, ***p < 0.01$.

We are interested not only in the likelihood of investing in clean R&D, but also in the share of clean R&D in each firm's total R&D, and the value of clean R&D investments. To get at the share, we replace the dummy variable with a variable capturing the value of clean R&D investments relative to the total R&D investments of the firm. The results can be found in columns (1) and (2) in Table 3 below, for a linear and PPML model respectively. We find a positive and significant effect also on this measure, although not as precisely estimated as the dummy. The baseline share of clean R&D is 6%, so again this is an economically significant effect. These findings suggest that firms more exposed to the oil price shock are increasing their clean R&D investments relative to their non-clean investments relatively more than less exposed firms. In columns (3) and (4) we replace the share of clean R&D with the value of clean R&D, in logs in column (3) and levels in column (4). Here, we find support for the hypothesis that exposed firms are increasing

[20] While it would be preferable to use a logistic model, these models struggle to converge with high-dimensional fixed effects. PPML can deal with high-dimensional fixed effects, but on the other hand assumes an exponential conditional expectation that will not automatically be valid for binary data. However, if the probabilities are all sufficiently close to zero, the exponential function is a good approximation to the logistic and therefore the use of PPML can be justified,

their clean R&D spending also in absolute terms, although the effect in column (4) is just below the threshold for significance at the 10 percent level. The coefficient from column (3) translates into an increase in the value of clean R&D spending of around 40% for firms exposed to the shock.

Table 3: Clean R&D: Share and Value

Variable:	Share (1)	Share (2)	Log Value (3)	Value (4)
$Post_t * x_{oi}$	1.575*	0.358*	0.346**	0.716
	(0.899)	(0.218)	(0.140)	(0.470)
Controls	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Ind.*year FE	Yes	Yes	Yes	Yes
Estimator	OLS	PPML	OLS	PPML
Obs.	11,695	3,024	11,695	3,024

Note: Standard errors in parenthesis are clustered on firm. Log Value is measured as log(1 + Clean R&D expenditures). Controls include baseline levels of log employment, log tangible assets, export share, energy share and a dummy for public funding, all interacted with year dummies. $*p < 0.1, **p < 0.05, ***p < 0.01$.

The findings presented in Tables 2 and 3 support the hypothesis that firms more exposed to the oil price shock through the oil producers' supply chain will react relatively stronger to the price shock in terms of investments in clean R&D.

Our exposure measure is continuous, which has been shown by Callaway et al. (2024) to complicate inference. As they demonstrate in their paper, one can summarize what they call "average level treatment effects among treated" firms by comparing the average change in outcomes for all treated firms to the average change in outcomes for untreated firms, following Sun and Shapiro (2022). The level treatment effect captures the difference between a firms's potential outcome under the exposure it receives and its untreated potential outcome – basically the equivalent of a treatment effect in a classical binary treatment set-up. To identify this effect, we rely on a standard parallel trends assumption: the average evolution of outcomes that firms with any exposure would have experienced without treatment is the same as the evolution of outcomes that firms in the non-exposed group actually experienced. We re-estimate all our results using a binary dummy variable that takes the value of one for firms with positive exposure and zero for non-exposed firms. The results are very similar to our baseline results, suggesting that the average level treatment effects are linear.[21]

[21] The results using a binary dummy variable are available upon request from the authors.

We investigate further where the increase in clean R&D investments is originating – is it mainly due to firms that have previously not invested in any R&D that start investing, or is it mainly driven by firms with pre-existing R&D investments that reallocate their R&D spending towards clean activities? To get at this, we create two dummy variables: First, "New Clean" equals one if the firm reports spending on clean R&D in period t but did not in period $t-1$. Second, "R&D to Clean" equals one if the firm reports spending on clean R&D in period t but did not in period $t-1$, conditional on the firm reporting positive R&D spending (in any type of R&D) in period $t-1$.

The results from this exercise is found in Table 4. In column (1), we see a strong and positive coefficient, indicating that firms that are more exposed to the shock are significantly more likely to start new, clean R&D projects, compared to firms that are not exposed. In columns (2) and (3) we show the results from using the "R&D to Clean" dummy. In column (2), all firms are included, regardless of whether or not they report any R&D spending, In column (3), we limit the sample to firms that report positive R&D investments before the shock. We, again, find that firms that are more exposed to the shock are more likely to start new clean R&D projects – this effect is particularly strong if we narrow the sample to include only R&D performers.

Table 4: Switching into Clean R&D

Variable:	"New Clean" (1)	"R&D to Clean" (2)	"R&D to Clean" (3)
$Post_t \times x_{oi}$	0.033***	0.018**	0.056***
	(0.011)	(0.009)	(0.022)
Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Industry*year FE	Yes	Yes	Yes
Control group	All firms	All firms	R&D firms
Observations	11,695	11,695	4,751

Note: Standard errors in parenthesis are clustered on firm. "New Clean" is a dummy that takes the value of 1 if the firm has positive investments in Clean R&D in year t, but not in t-1. "R&D to Clean" is the same as "New Clean" conditional on having invested in R&D in t-1. Controls include baseline levels of log employment, log tangible assets, export share, energy share and a dummy for public funding, all interacted with year dummies. $*p < 0.1, **p < 0.05, ***p < 0.01$.

To summarize our results: we find that firms that are more exposed to the 2014 fall in the oil price through the oil producers' supply chain increase their clean R&D more than firms that were less exposed to the shock. This is true if we look at the likelihood of investing in clean R&D, the share of clean relative to non-clean R&D, or the value of clean R&D. We also document that the firms which are most exposed to the shock have higher probabilities of starting new investments in clean R&D, and switching from non-clean to clean R&D. We will now explore what mechanisms are at play to produce these effects. Our results indicates that within-firm dynamics such as adjustment costs, are relevant for the understanding the full effect of such price drops on clean technological progress.

Pre-Trends A potential concern is that more exposed firms face different pre-trends compared to those less exposed. It would be reassuring to find similar pre-trends for firms with different exposure levels indicating that our identification strategy is solid. We investigate this by estimating the dynamic pattern of our main outcome variable, the dummy for clean R&D:

$$y_{kt} = \alpha_k + \beta x_{ok} \times \delta_t + \gamma Z_{kt} + \delta_{st} + \varepsilon_{kt}, \tag{18}$$

For ease of exposition, we show the graph for a binary treatment variable here.[22] The result from estimating this specification is illustrated in Figure 8. First, the lack of evidence of any sort of pattern in the pre-trends is reassuring, and suggests that our treatment and control firms are behaving similarly in our pre-period. We see a clear upward trajectory beginning in 2014 and turning significant from 2015.

Figure 8: Dynamic DID

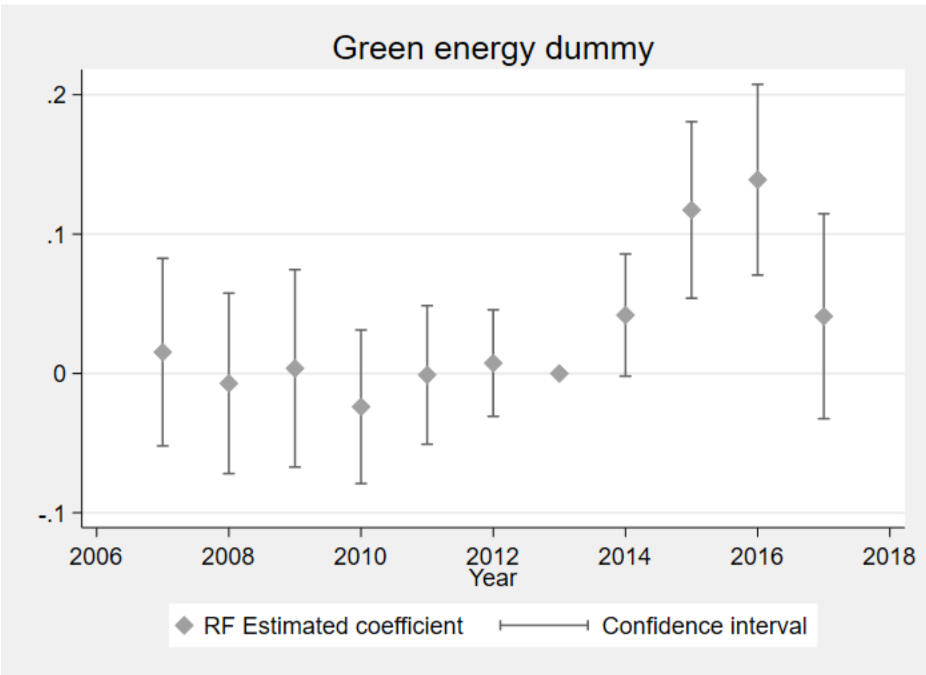

Note Each point presents the event-study coefficient estimates and 90% confidence intervals of Equation 18.

[22]The graph looks very similar with the continuous treatment variable, but is harder to interpret.

Table 5: Placebo

Variable:	Bio tech R&D (1)	ICT R&D (2)
$Post_t \times x_{oi}$	-0.015	-0.032
	(0.015)	(0.024)
Controls	Yes	Yes
Firm FE	Yes	Yes
Industry*year FE	Yes	Yes
Observations	11,695	11,695

Standard errors in parenthesis are clustered on firm. Controls include baseline levels of log employment, log tangible assets, export share, energy share and a dummy for public funding, all interacted with year dummies. $*p < 0.1, **p < 0.05, ***p < 0.01$.

Robustness Another potential concern is that firms exposed to the shock shifted their R&D investments towards other technological fields in general, but not with any specific emphasis on clean R&D. To address the concern of whether the shock actually triggered directed technical change we perform a placebo test using information on firms' investments in bio tech R&D and ICT R&D. We estimate Equation (17) with these two new variables as dependent variables. Table 5 reports the results. The estimated coefficients are not significant.

Although our dynamic difference-in-difference estimation is reassuring, one cannot exclude the possibility that our results are driven by something other than the oil price. In particular, if clean energy prices increased substantially during the time period we are looking at, this would have incentivized firms to ramp up their clean R&D investments., and exposed firms might in theory be hit differently by changes in clean energy prices, compared to less exposed firms. However, figure 10 in Appendix Section D shows how clean energy prices evolved over our sample period. If anything, they declined smoothly over the time period in question.

Mechanisms In our theoretical framework, within-firm reallocation of investments from dirty to clean R&D after a negative oil price shock is driven by the fall in expected profits from producing inputs for the oil-extracting sector. Therefore, we investigate whether the firms we identify as more exposed to the negative oil price shock experience a decrease in sales and profits.

We create two new variables, sales per employee and an indicator variable for operating profits. The indicator takes on the value of either 0, -1 or 1, depending on whether the firms makes zero, negative or positive profits. We estimate Equation (17) with these two new variables as dependent variables. The results from these regressions are presented in

Table 6. We find that sales per employee decreases significantly, as does profits, although the latter finding is not as robust as the first. Our results suggest that the mechanism suggested by the theoretical model is indeed at play.

Table 6: Sales per employee and Profits indicator

Variable:	Sales per emp. (1)	Sales per emp. (2)	Profits indicator (3)	Profits indicator (4)
$Post_t \times x_{oi}$	-0.082***	-0.043*	-0.179***	-0.081
	(0.026)	(0.026)	(0.049)	(0.060)
Controls	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	No	Yes	No
Industry FE	Yes	No	Yes	No
Industry*year FE	No	Yes	No	Yes
Observations	11,695	11,695	11,695	11,695

Standard errors in parenthesis are clustered on firm. Income is measured as operating income and profits as an indicator for operating profits. Controls include baseline levels of log employment, log tangible assets, export share, energy share and a dummy for public funding, all interacted with year dummies. The indicator for operating profits takes on 0, -1 and 1 depending on whether the firms makes zero, negative or positive profits. $*p < 0.1, **p < 0.05, ***p < 0.01$.

Finally, we explore whether the most exposed firms intensify their R&D investments across the board, relative to other firms. As shown in Table 7, we find no effects of the shock on total R&D investments, the number of R&D employees, or on the share of R&D employees in total employment. These results suggest that firms are not scaling up all R&D investments in response to the shock. Rather, in line with our theoretical predictions, they appear to reallocate resources towards clean R&D investments.

Table 7: R&D

Variable:	R&D dummy (1)	log R&D emp (2)	R&D emp share (3)
$Post_t \times x_{oi}$	-0.015	-0.008	-0.001
	(0.028)	(0.007)	(0.001)
Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Industry*year FE	Yes	Yes	Yes
Observations	11,695	11,695	11,695

Standard errors in parenthesis are clustered on firm. Log R&D employment is log(1+R&D employees). Controls include baseline levels of log employment, log tangible assets, export share, energy share and a dummy for public funding, all interacted with year dummies. $*p < 0.1, **p < 0.05, ***p < 0.01$.

6 Concluding Remarks

In this paper, we aim to shed light on a potentially important channel through which changes in the price of fossil fuels may affect investments in clean technology. We do so by developing a theoretical framework that highlights the role of within-firm reallocation of R&D investments after a negative price shock to fossil energy. We exploit the rich data at hand to develop a novel measure of shock exposure. Our results suggest that a price shock which trickles down the supply chain, combined with adjustment costs at the firm level, may induce within-firm reallocation of resources from dirty towards clean innovation.

Our findings have some key policy implications. Understanding the mechanisms through which firms react to external price shocks and the subsequent reallocation of resources is crucial for designing effective policies that promote environmentally sustainable economic growth. Our results imply that policies lowering the profitability in the oil producing sector may promote a shift towards investment in clean R&D, not only through the induced structural change which entail reallocation of resources from fossil towards clean energy producing firms, but also indirectly through within-firm reallocation of R&D investments in upstream industries. Importantly, our findings show that carbon pricing can lead to reallocation of resources in favor of clean innovation, not only through increasing demand for clean alternatives, but also through within-firm reallocation in the fossil energy sector.

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hémous, "The Environment and Directed Technical Change," *American Economic Review*, 2012, *102* (1), 131–66. 1, 5, 3, 3.1, 3.1

_ , _ , Lint Barrage, and David Hémous, "Climate Change, Directed Innovation, and Energy Transition: The Long-Run Consequences of the Shale Gas Revolution," Technical Report, National Bureau of Economic Research 2023. 1

_ , Ufuk Akcigit, Douglas Hanley, and William Kerr, "Transition to Clean Technology," *Journal of Political Economy*, 2016, *124* (1), 52–104. 5

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

Armington, Paul S., "A Theory of Demand for Products Distinguished by Place of Production," *Staff Papers (International Monetary Fund)*, 1969, *16* (1), 159–178. C

Barrot, Jean-Noel and Julien Sauvagnat, "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks," *The Quarterly Journal of Economics*, 2016, *131* (3), 1543–1592. 1

Blanchard, Olivier, Christian Gollier, and Jean Tirole, "The Portfolio of Economic Policies Needed to Fight Climate Change," *Annual Review of Economics*, 2023, *15* (1), 689–722. 7

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 2013, *81* (4), 1347–1393. 12

_ , Paul M Romer, Stephen J. Terry, and John Van Reenen, "A Trapped-Factors Model of Innovation," *American Economic Review*, 2013, *103* (3), 208–213. 1, 3

Bloom, Nick, "Uncertainty and the Dynamics of R&D," *American Economic Review*, 2007, *97*, 250–255. 3, 2.4

Calel, Raphael and Antoine Dechezleprêtre, "Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market," *Review of Economics and Statistics*, 2016, *98* (1), 173–191. 1

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna, "Difference-in-Differences with a Continuous Treatment," Technical Report 2024. 5

Casey, Gregory, "Energy Efficiency and Directed Technical Change: Implications for Climate Change Mitigation," *Review of Economic Studies*, 2024, *91* (1), 192–228. 5

Crabb, Joseph M and Daniel K.N. Johnson, "Fueling Innovation: The Impact of Oil Prices and CAFE Standards on Energy-Efficient Automotive Technology," *The Energy Journal*, 2010, *31* (1). 1

Dechezleprêtre, Antoine and David Hémous, "Directed Technical Change and Environmental Economics," Technical Report, in "The Economics of Creative Destruction", a Festschrift in honor of Philippe Aghion and Peter Howitt 2022. 1

_ **and Tobias Kruse**, "The Effect of Climate Policy on Innovation and Economic Performance Along the Supply Chain: A firm- and Sector-Level Analysis," Technical Report, OECD Environment Working Papers No. 189 2022. 1

Dhyne, Emmanuel, Ayumu Ken Kikkawa, Toshiaki Komatsu, Magne Mogstad, and Felix Tintelnot, "Foreign Demand Shocks to Production Networks: Firm Responses and Worker Impacts," Working Paper 30447, National Bureau of Economic Research 2022. 1

Dugoua, Eugenie and Marion Dumas, "Green Product Innovation in Industrial Networks: A Theoretical Model," *Journal of Environmental Economics and Management*, 2021, *107*, 102420. 1

_ **and Todd Gerarden**, "Induced Innovation, Inventors, and the Energy Transition," *American Economic Review: Insights*, 2024, *Forthcoming.* 1

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

Gerlagh, Reyer, Snorre Kverndokk, and Knut Einar Rosendahl, "Optimal Timing of Climate Change Policy: Interaction Between Carbon Taxes and Innovation Externalities," *Environmental and Resource Economics*, 2009, *43*, 369–390. 7

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski, "Optimal Taxes on Fossil Fuel in General Equilibrium," *Econometrica*, 2014, *82* (1), 41–88. 5, 7

Greaker, Mads, Tom-Reiel Heggedal, and Knut Einar Rosendahl, "Environmental Policy and the Direction of Technical Change," *The Scandinavian Journal of Economics*, 2018, *120* (4), 1100–1138. 5

Hamermesh, Daniel S., "Labor Adjustment and the Source of Adjustment Costs," *The Economic Journal*, 1995, *105*, 620–634. 15

Hart, Rob, "To Everything there is a Season: Carbon Pricing, Research Subsidies, and the Transition to Fossil-Free Energy," *Journal of the Association of Environmental and Resource Economists*, 2019, *6* (2), 349–389. 5

Hu, Xiaolu, Jing Yu, and Angel Zhong, "The Asymmetric Effects of Oil Price Shocks on Green Innovation," *Energy Economics*, 2023, *125*, 106890. 1

Köppl, Angela and Margit Schratzenstaller, "Carbon Taxation: A Review of the Empirical Literature," *Journal of Economic Surveys*, 2023, *37* (4), 1353–1388. 7

Lemoine, Derek, "Innovation-Led Transitions in Energy Supply," *American Economic Journal: Macroeconomics*, 2024, *16* (1), 29–65. 5

Ley, Marius, Tobias Stucki, and Martin Woerter, "The Impact of Energy Prices on Green Innovation," *The Energy Journal*, 2016, *37* (1), 41–76. 1

Lorentzen, Linnea, "Domino Effects: Understanding Sectoral Reallocation and its Wage Implications," Technical Report 2024. 8

Marx, Matt and Bram Timmermans, "Hiring Molecules, Not Atoms: Comobility and Wages," *Organization Science*, 2017, *28* (6), 1115–1133. 14

Newell, Richard G, Adam B Jaffe, and Robert N Stavins, "The Induced Innovation Hypothesis and Energy-Saving Technological Change," *Quarterly Journal of Economics*, 1999, *114* (3). 1

Nix, Emily, "Learning spillovers in the firm," Technical Report, IFAU Working Paper No. 2020:14 2020. 12

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

_ , "Environmental Policy and Innovation: A Decade of Research," *National Bureau of Economic Research Working Paper*, 2019. 6

_ , Richard G Newell, and Adam B Jaffe, "Energy, the Environment, and Technological Change," *Handbook of the Economics of Innovation*, 2010, *2*, 873–937. 6

Prestmo, Joakim Blix, Birger Strøm, and Hilde Karoline Midsem, "Ringvirkninger av Petroleumsnæringen i Norsk Økonomi," Technical Report 8 2015. 4

Stern, Nicholas and Anna Valero, “Innovation, Growth and the Transition to Net-Zero Emissions,” *Research Policy*, 2021, *50* (9), 104293. 7

Sun, Liyang and Jesse M. Shapiro, “A Linear Panel Model with Heterogeneous Coefficients and Variation in Exposure,” *Journal of Economic Perspectives*, 2022, *36* (4), 193–204. 5

Timilsina, Govinda R, “Carbon Taxes,” *Journal of Economic Literature*, 2022, *60* (4), 1456–1502. 7

Appendix

A Data

The R&D survey measures R&D activity in the Norwegian business enterprise sector. The statistics are comparable to statistics for other countries and are reported to the OECD and EUROSTAT. The R&D survey includes: (i) all firms with at least 50 employees; (ii) all firms with less than 50 employees and with reported intramural R&D activity in the previous survey of more than NOK 1 million or extramural R&D of more than NOK 3 million; (iii) among other firms with 10-49 employees a random sample was selected within each strata (NACE 2-digit and size class).

The R&D survey provides details on the share of R&D spend in certain thematic areas. We use information on the share of R&D spent on what we define as clean energy, which encompasses the two areas: Renewable energy and Energy efficiency.

B Calculations for the analytical model

In the following, we provide the details for calculation of the signs of the first- and second-order effects of the drop in the price of the dirty final good on clean R&D activity, in the analytical model presented in Section 3,The calculations provided in the following makes us able to compare the total effects of the price drop on clean R&D in firms engaged in both clean and dirty R&D (group III) relative to firms only engaged in clean R&D (group I), and provide proofs for Proposition 1 and 2.

For notational simplicity, let us define the function $H^l_{kjt}(\cdot)$ as the derivative of the objective function (Equation 12) of firm k, belonging to group l, with respect to s_{kjt}. $H^l_{kjt}(\cdot)$ represents the net value of a marginal increase in s_{kjt} for firm k in group l. Each firm will hire scientists until the relevant marginal value is zero, i.e. until $H^l_{kjt}(\cdot) = 0$.

For firms in group I and II, the first-order conditions are given by:

$$H^I_{kct}(s_{kct}, p_{ct}, w_t) = \alpha(1-\alpha)(1+\gamma)\eta_c A_{ct-1} p_{ct}^{\frac{1}{1-\alpha}} \frac{\partial g^c(s_{kct}, 0)}{\partial s_{kct}} - w_t - 2c\frac{s_{kct} - s_{kct-1}}{s^2_{kct-1}} = 0 \tag{19}$$

$$H^{II}_{kdt}(s_{kdt}, p_{dt}, w_t) = \alpha(1-\alpha)(1+\gamma)\eta_d A_{dt-1} p_{dt}^{\frac{1}{1-\alpha}} g^{d\prime}(s_{kdt}) - w_t - 2c\frac{s_{kdt} - s_{kdt-1}}{s^2_{kdt-1}} = 0. \tag{20}$$

For firms in group III, the two first-order conditions are given by:

$$H^{III}_{kct}(s_{kct}, s_{kdt}, p_{ct}, w_t) = \alpha(1-\alpha)(1+\gamma)\eta_c A_{ct-1} p_{ct}^{\frac{1}{1-\alpha}} \frac{\partial g^c(s_{kct}, s_{kdt})}{\partial s_{kct}} - w_t - 2c\frac{s_{kt} - s_{kt-1}}{s^2_{kt-1}} = 0 \tag{21}$$

$$\begin{aligned} H^{III}_{kdt}(s_{kct}, s_{kdt}, p_{ct}, p_{dt}, w_t) =& \alpha(1-\alpha)(1+\gamma)\Bigg(\eta_c A_{ct-1} p_{ct}^{\frac{1}{1-\alpha}} \frac{\partial g^c(s_{kct}, s_{kdt})}{\partial s_{kdt}} \\ &+ \eta_d A_{dt-1} p_{dt}^{\frac{1}{1-\alpha}} g^{d'}(s_{kdt})\Bigg) - w_t - 2c\frac{s_{kt} - s_{kt-1}}{s^2_{kt-1}} = 0 \end{aligned} \tag{22}$$

with $s_{kt} = s_{kct} + s_{kdt}$.

The first-order conditions in equations 19-22 give the demand for scientists from all three groups of firms.

In the following, the group-specific superscripts are dropped for expressions that take the same value for all relevant groups. For group I and II, respectively, the following second-order conditions ensure that the first-order conditions characterize a maximum:

$$\frac{\partial H_{kct}}{\partial s_{kct}} = \alpha(1-\alpha)(1+\gamma)\eta_c A_{ct-1} p_{ct}^{\frac{1}{1-\alpha}} \frac{\partial^2 g^c}{\partial s^2_c} - \frac{2c}{s^2_{kt-1}} \leq 0 \tag{23}$$

$$\frac{\partial H_{kdt}}{\partial s_{kdt}} = \alpha(1-\alpha)(1+\gamma)\eta_d A_{dt-1} p_{dt}^{\frac{1}{1-\alpha}} g^{d\prime\prime} - \frac{2c}{s^2_{kt-1}} \leq 0. \tag{24}$$

These inequalities follow directly from our initial assumptions. For group III, we also require that:

$$\frac{\partial H_{kct}}{\partial s_{kct}} \frac{\partial H_{kdt}}{\partial s_{kdt}} - \left(\frac{\partial H^{III}_{kct}}{\partial s_{kdt}}\right)^2 \geq 0 \tag{25}$$

with

$$\frac{\partial H^{III}_{kct}}{\partial s_{kdt}} = \frac{\partial H^{III}_{kdt}}{\partial s_{kct}} = \alpha(1-\alpha)(1+\gamma)\eta_c A_{ct-1} p_{ct}^{\frac{1}{1-\alpha}} \frac{\partial^2 g^c}{\partial s_c \partial s_d} - 2c\frac{1}{s^2_{kt-1}} \lesseqgtr 0. \tag{26}$$

The cross-derivative of the objective function of a firm in group III is positive if the spillover effect dominates the effect of the adjustment cost, while it is negative if the adjustment cost dominates. In either case, we assume that Equation (**??**) holds , i.e. that the effect of dirty-scientist hiring on the marginal value and cost of hiring scientists for clean innovation, independent of its direction, is not large enough to prevent an interior solution to the problem of any firm in group III.

First-order effects of the oil price drop.

To get at the effects of the price drop on each firm's R&D activity (demand for scientists), we use the first-order conditions . First, differentiating the first-order conditions with respect to p_{dt} and w gives:

$$\frac{\partial H^l_{kdt}}{\partial p_{dt}} =\alpha(1+\gamma)\eta_d A_{dt-1} p_{dt}^{\frac{\alpha}{1-\alpha}} g^{d\prime}(s^l_{kdt}) > 0, \quad \text{for } l = II, III. \tag{27}$$

$$\frac{\partial H^l_{kjt}}{\partial w_t} =-1, \quad \text{for } l = I, II, III \ og j = d, t \tag{28}$$

Next, it follows directly from differentiating the first-order conditions and using the inequalities in equations 23-26 that:

$$\frac{\partial s^{II}_{kdt}}{\partial p_{dt}} = -\frac{\frac{\partial H^{II}_{kdt}}{\partial p_{dt}}}{\frac{\partial H_{kdt}}{\partial s_{kdt}}} > 0 \tag{29}$$

$$\frac{\partial s^{III}_{kdt}}{\partial p_{dt}} = -\frac{\frac{\partial H_{kct}}{\partial s_{kct}}\frac{\partial H^{III}_{kdt}}{\partial p_{dt}}}{\frac{\partial H_{kct}}{\partial s_{kct}}\frac{\partial H_{kdt}}{\partial s_{kdt}} - \left(\frac{\partial H^{III}_{kct}}{\partial s_{kdt}}\right)^2} > 0 \tag{30}$$

$$\frac{\partial s^{III}_{kct}}{\partial p_{dt}} =\frac{\frac{\partial H^{\mathrm{III}}_{kct}}{\partial s_{kdt}}\frac{\partial H^{III}_{kdt}}{\partial p_{dt}}}{\frac{\partial H_{kct}}{\partial s_{kct}}\frac{\partial H_{kdt}}{\partial s_{kdt}} - \left(\frac{\partial H^{III}_{kct}}{\partial s_{kdt}}\right)^2} \lesseqgtr 0 \tag{31}$$

$$\frac{\partial s^{III}_{kct}}{\partial p_{dt}} + \frac{\partial s^{III}_{kdt}}{\partial p_{dt}} =\frac{\alpha(1-\alpha)(1+\gamma)\eta_c A_{ct-1} p_{ct}^{\frac{1}{1-\alpha}}\left(\frac{\partial^2 g^c}{\partial s_c \partial s_d} - \frac{\partial^2 g^c}{\partial s_c^2}\right)\frac{\partial H^{III}_{kdt}}{\partial p_{dt}}}{\frac{\partial H_{kct}}{\partial s_{kct}}\frac{\partial H_{kdt}}{\partial s_{kdt}} - \left(\frac{\partial H^{III}_{kct}}{\partial s_{kdt}}\right)^2} > 0 \tag{32}$$

Effects of wage changes.

Similarly, it follows, again from differentiating the first-order conditions and using 23-26, that:

$$\frac{\partial s^I_{kct}}{\partial w_t} = 1 \Bigg/ \frac{\partial H_{kct}}{\partial s_{kct}} < 0, \tag{33}$$

$$\frac{\partial s^{II}_{kdt}}{\partial w_t} = 1 \Bigg/ \frac{\partial H_{kdt}}{\partial s_{kdt}} < 0, \tag{34}$$

$$\frac{\partial s^{III}_{kct}}{\partial w_t} = \left(-\frac{\partial H^{III}_{kct}}{\partial s_{kdt}} + \frac{\partial H_{kdt}}{\partial s_{kdt}} \right) \Bigg/ \left(\frac{\partial H_{kct}}{\partial s_{kct}} \frac{\partial H_{kdt}}{\partial s_{kdt}} - \left(\frac{\partial H^{III}_{kct}}{\partial s_{kdt}} \right)^2 \right) < 0 \tag{35}$$

$$\frac{\partial s^{III}_{kdt}}{\partial w_t} = \left(-\frac{\partial H^{III}_{kct}}{\partial s_{kdt}} + \frac{\partial H_{kct}}{\partial s_{kct}} \right) \Bigg/ \left(\frac{\partial H_{kct}}{\partial s_{kct}} \frac{\partial H_{kdt}}{\partial s_{kdt}} - \left(\frac{\partial H^{III}_{kct}}{\partial s_{kdt}} \right)^2 \right) < 0. \tag{36}$$

The market equilibrium condition in Equation 13 gives the equilibrium wage as a function of final good prices: $w_t(p_{ct}, p_{dt})$. Differentiating shows that the wage will drop as a consequence of a fall in the fossil energy price:

$$\frac{\partial w_t}{\partial p_{dt}} = -\frac{\lambda^{II} \frac{\partial s^{II}_{kdt}}{\partial p_{dt}} + \lambda^{III} \left(\frac{\partial s^{III}_{kct}}{\partial p_{dt}} + \frac{\partial s^{III}_{kdt}}{\partial p_{dt}} \right)}{\lambda^I \frac{\partial s^I_{kct}}{\partial w_t} + \lambda^{II} \frac{\partial s^{II}_{kdt}}{\partial w_t} + \lambda^{III} \left(\frac{\partial s^{III}_{kct}}{\partial w_t} + \frac{\partial s^{III}_{kdt}}{\partial w_t} \right)} > 0. \tag{37}$$

We can now turn to the proof of Proposition 1 and 2.

Proof of Proposition 1 and 2.

Define

$$\Delta \equiv -\frac{ds^{III}_{kct}}{dp_{dt}} - \left(-\frac{ds^I_{kct}}{dp_{dt}} \right), \tag{38}$$

with

$$\frac{ds^{III}_{kct}}{dp_{dt}} = \frac{\partial s^{III}_{kct}}{\partial p_{dt}} + \frac{\partial s^{III}_{kct}}{\partial w_t} \frac{\partial w_t}{\partial p_{dt}} \tag{39}$$

$$\frac{ds^I_{kct}}{dp_{dt}} = \frac{\partial s^I_{kct}}{\partial w_t} \frac{\partial w_t}{\partial p_{dt}} \tag{40}$$

Insert from Equation (**??**) and (**??**) and reorganize, to get:

$$\Delta = \lambda^I X_1 + \lambda^{II} X_2 + \lambda^{III} X_3, \tag{41}$$

with:

$$X_1 = \frac{\partial s^{III}_{kct}}{\partial p_{dt}} \frac{\partial s^{I}_{kct}}{\partial w_t} \begin{cases} > 0 \text{ if } \frac{\partial s^{III}_{kct}}{\partial p_{dt}} < 0 & \text{(adjustment costs dominate)} \\ < 0 \text{ if } \frac{\partial s^{III}_{kct}}{\partial p_{dt}} > 0 & \text{(spillovers dominate)} \end{cases} \tag{42}$$

$$X_2 = \frac{\frac{\partial H^{II}_{kdt}}{\partial p_{dt}}}{\frac{\partial H_{kct}}{\partial s_{kct}} \frac{\partial H_{kdt}}{\partial s_{kdt}} - \left(\frac{\partial H^{III}_{kct}}{\partial s_{kdt}}\right)^2} \frac{\frac{\partial H^{III}_{kct}}{\partial s_{kdt}}}{\frac{\partial H_{kct}}{\partial s_{kct}} \frac{\partial H_{kdt}}{\partial s_{kdt}}} \left(\frac{\partial H^{III}_{kct}}{\partial s_{kdt}} - \frac{\partial H^{III}_{kct}}{\partial s_{kct}} \left(1 - \frac{g^{d\prime}(s^{III}_{kdt})}{g^{d\prime}(s^{II}_{kdt})}\right)\right), \tag{43}$$

$$X_3 = \frac{\frac{\partial H^{III}_{kdt}}{\partial p_{dt}}}{\frac{\partial H_{kct}}{\partial s_{kct}} \frac{\partial H_{kdt}}{\partial s_{kdt}} - \left(\frac{\partial H^{III}_{kct}}{\partial s_{kdt}}\right)^2} \frac{\frac{\partial H^{III}_{kct}}{\partial s_{kdt}}}{\frac{\partial H_{kct}}{\partial s_{kct}}} \begin{cases} > 0 \text{ if } \frac{\partial s^{III}_{kct}}{\partial p_{dt}} < 0 & \text{(adjustment costs dominate))} \\ < 0 \text{ if } \frac{\partial s^{III}_{kct}}{\partial p_{dt}} > 0 & \text{(spillovers dominate)} \end{cases} \tag{44}$$

The first fraction in X_2 is always non-negative. The second fraction takes the same sign as the fraction $\partial H^{III}_{kct}/\partial s_{kdt}$, positive if the spillover effect dominates, negative if the adjustment cost effect dominates. Next, the expression inside the inner brackets will take a negative value as long as $s^{III}_{kdt} \leq s^{II}_{kdt}$. With $s^{III}_{kt-1} = s^{II}_{kt-1}$,where the group-$III$ firm uses only a share of their scientists for dirty R&D , we assume in the following that this is the case. The second term in the outer brackets will then take a positive value.

If the adjustment costs dominates, the full expression in the outer brackets will be negative, and X_2 will take the opposite value of $\partial H^{III}_{kct}/\partial s_{kdt}$. The same will be the case if the spillover effect dominates, as long as $\partial H^{III}_{kct}/\partial s_{kdt}$ is not large too large. In the main body of the paper, we assume this is the case. However, it is worth noting that for the special case of very large spillover effects, X_2 may take the same sign as $\partial H^{III}_{kct}/\partial s_{kdt}$. As a consequence, if λ^{II} is sufficiently large, the same can become true for the full expression for Δ in Equation (**??**). The firms in group II will react to the price drop by decreasing their demand for scientists, which will push the wage down. If the share of firms belonging to this group is very large, the firms in group III may face a wage drop that is large relative to the drop in prices.

Given the signs of X_1,X_2 and X_3, disregarding the special case just described, we have:

$$\Delta \begin{cases} > 0 \text{ if } \frac{\partial s_{kct}^{III}}{\partial p_{dt}} < 0 \\ = 0 \text{ if } \frac{\partial s_{kct}^{III}}{\partial p_{dt}} = 0 \\ < 0 \text{ if } \frac{\partial s_{kct}^{III}}{\partial p_{dt}} > 0 \end{cases}$$

i.e., exposed-firm clean R&D increase more than clean R&D in non-exposed firms if the rescaling cost dominates, while one would see the largest increase in the non-exposed firms (and possibly even a decrease in exposed firms) if the spillover effect dominates. This concludes the proof.

Note that the analysis presented above assumes interior solution to the maximization problem of the firms with respect to s_{kjt} in for all firms with $\eta_{kjt} = \eta_j$. As a consequence, a drop in the dirty-good price will only induce intensive-margin shifts in clean R&D, as all firms with the ability to engage in clean R&D will have $s_{kct} > 0$ for any set of prices, by assumption. However, allowing the marginal productivity of innovation to be finite as s_{kct} approaches 0: $\lim_{s_{kct} \to 0} g^c(s_{kct}, s_{kdt}) \neq \infty$, would induce firms in group I and III to engage in clean R&D activity only if the expected revenue is sufficiently high. As a result, a drop in the dirty-good price would potentially induce increases in clean R&D in both groups both on the extensive and on the intensive margin.

C The firm-level measure of exposure to the oil price drop

We use a two-step process to identify the level of exposure of each firm in our data set to the 2014 oil price drop.

First, we categorize all products imported to Norway as oil related or not oil related. In this first step, we rely on data on imports to oil producing firms in Norway. All products that are imported by these firms are defined as oil related. In the period before the 2014 oil price drop, around 7000 products are imported to Norway, and about 1800 of these are imported (in any quantity) by the oil producing sector, and thereby defined as oil related in our measure.

When determining which products to define as inputs to the oil production , we take into account the potentially strong home bias of the oil extraction sector operating in Norway regarding the sourcing of inputs. Since the start of Norwegian oil extraction, it has been a stated goal to build a Norwegian supply sector, and strong policy measures have been in place to obtain this goal.

Building on Armington (1969) we assume that the production function of the oil producers encompass different preferences for domestically produced and foreign inputs :

$$Y_{dt} = \beta_{di} \int_0^1 A_{Ddit}^{1-\alpha} x_{Ddit}^{\alpha} di + (1 - \beta_{di}) \int_0^1 A_{Mdit}^{1-\alpha} x_{Mdit}^{\alpha} di, \tag{45}$$

with subscript D denoting domestically produced inputs and subscript M denoting foreign inputs, while β_{di} is the value of the imported good in production relative to the domestically produced good. The size of β_{di} determines the home bias for input di. For higher values of β_{di}, input demand would be directed more towards domestically produced goods. If the home bias differs across sectors and inputs, the share of imports of good x_{di} that is directed towards the oil producing sector will not be informative of the relative importance of good x_{di} as an input in the oil producing sector.Therefore, we do not weight the goods categorized as inputs to the oil producing sector using import shares.

D Additional Figures

Figure 9: Producer Price Index - Norway (2021=100)

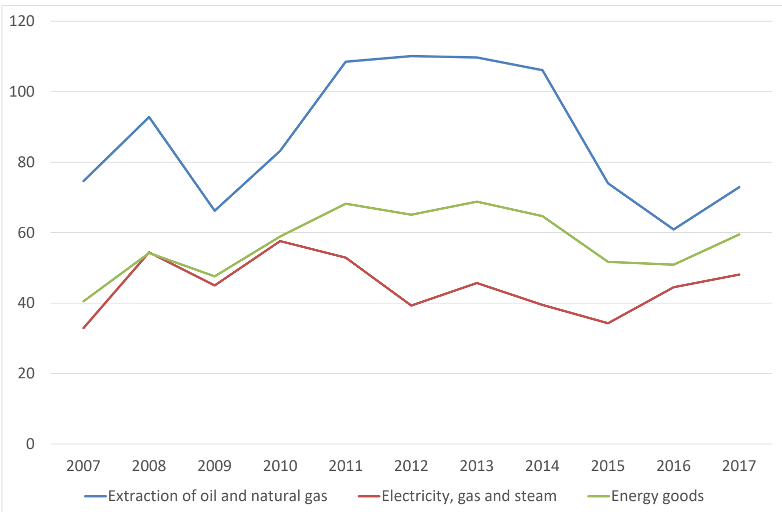

The figure shows the development of the producer price index for the domestic market for the commodities/industries (i) oil- and gas extraction, (ii) electricity, gas and steam and (iii) the aggregate of energy goods with 2021 as the reference year (2021=100).

Figure 10: Global renewable costs (IRENA)

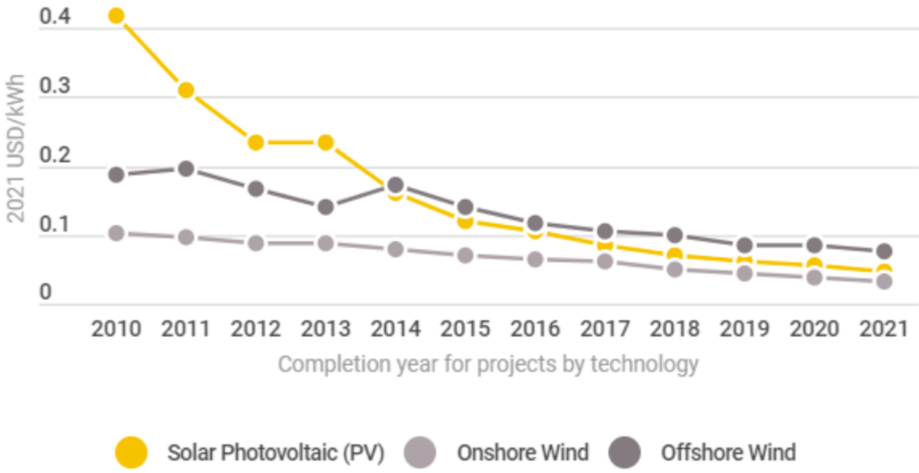

The figure shows the development of cost of renewable energy according to technology. Source: IRENA (www.irena.org)