Shock Therapy for Clean Innovation*

Esther Ann Bøler† Katinka Holtsmark‡ Karen Helene Ulltveit-Moe§

August 2025

Abstract

We study how a negative profitability shock in the fossil energy supply chain affects firms' direction of innovation. We develop a stylized model to show that adjustment costs in R&D create incentives for exposed firms to reallocate innovation toward clean technologies. Next, we propose a novel method to identify firms' exposure to the 2014 oil price collapse, and find that more exposed firms significantly increased clean R&D relative to less exposed peers. The results suggest that firms in the fossil energy supply chain possess transferable capabilities for clean innovation, and that declining fossil profitability—e.g., via carbon pricing—can accelerate the clean transition along the fossil energy supply chain.

JEL codes: D22, F14, F18, O31, Q55, Q56.

Keywords: R&D, clean innovation, supply chains, green transition.

*We thank Lassi Ahlvik, Johannes Boehm, Eugenie Dugoua, Rikard Forslid, Bård Harstad, Linnea Lorentzen, Christoph Meinerding, Natalie Rickard, Sjak Smulders, and Andre Trindade, along with seminar audiences at University of Oxford, University of Cambridge, Geneva Graduate Institute, University of Surrey, NHH Norwegian School of Economics, and OsloMET and participants at the 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, CCA Workshop on International Economics, Villars Workshop on International Trade, QMUL/Bristol/KCL/Bank of England Workshop on Firm Dynamics, Market Structures and Productivity in the Macroeconomy, and ERWIT 2025 for valuable comments and discussion. This project has received funding from the Norwegian Ministry of Trade, Industry and Fisheries (INNOPROD). Data from Statistics Norway have been crucial for this project.

†Imperial 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. Legacy firms in the fossil energy sector are increasingly keen to position themselves as central actors in this clean transition, highlighting their technological capabilities.[1] Policymakers, facing the challenge of designing incentives to stimulate investment in clean technologies, often introduce policies that *reduce* the profitability of fossil-related activities. These policies will shift resources away from the fossil energy sector. However, it remains poorly understood whether a decline in profitability will encourage firms in this sector to reallocate resources toward clean R&D, or instead constrain their ability to do so by eroding the resources needed to leverage their existing capabilities.[2]

In this paper, we investigate to what extent firms in the fossil energy supply chain react to a decline in profitability, induced by a negative oil price shock, by reallocating resources towards clean R&D.[3] While there is an extensive literature investigating how clean R&D responds to changes in energy prices, the main focus has been on how incentives for clean innovation are affected by changes in demand for renewable energy and clean technologies (e.g., Aghion et al. (2016), Acemoglu et al. (2023)). What has received less attention in the empirical literature to date is whether changes to the fossil energy price also affect investments in clean R&D for firms in the fossil energy supply chain by altering their expected future profitability. We fill this gap in the literature and identify a novel channel through which lower fossil energy profitability triggers a reallocation of R&D resources towards clean innovation.

Our analysis is motivated by three stylized facts based on the rich dataset at hand: First, in the aftermath of the 2014 oil price shock, the share of input suppliers in the fossil energy supply chain conducting clean R&D increased sharply. Second, there is strong

[1]For example, BP has prominently marketed its "And, not or" campaign, which argues that the future of energy depends on pursuing both oil and renewables simultaneously rather than choosing between them. Similarly, Shell, ExxonMobile and TotalEnergies have launched campaigns showcasing their investments in electric vehicle charging infrastructure, hydrogen projects, and biofuels, often framing these initiatives as core to their corporate identity. These campaigns have attracted the attention of regulators in several countries, including the US, the UK and France.

[2]Examples of green technology areas where capabilities and knowledge originating in the fossil energy supply chain are particularly relevant include wind and geothermal energy. For example, in September 2024, the US Department of Energy launched an initiative 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). For examples from our study context, see Section 2.5.

[3]Throughout the paper, we will refer to firms directly involved in oil extraction as *oil-extracting firms*, and their suppliers of machines and other inputs used in the extraction process as *input suppliers.* When we refer to the *fossil energy supply chain,* we include both oil-extracting firms and their input suppliers. Clean R&D refers to R&D expenditures directed towards renewable energy and energy efficiency. More details about the data can be found in Section 2.3.

persistence in R&D at the firm level over time.[4] Third, the majority of the increase in clean R&D after the shock came from firms that also had dirty R&D investments. In short, the data suggest that there is a nontrivial 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.

Building on these facts, we develop a stylized theoretical model of firms potentially investing in both clean and dirty innovation. We introduce two key features to account for within-firm dynamics. First, we follow Bloom et al. (2013b), and assume that there are adjustment costs related to the rescaling of R&D.[5] The presence of adjustment costs implies that in the aftermath of a negative shock to expected returns to dirty innovation, firms more exposed to the shock will face a relatively stronger incentive to divert their R&D towards clean innovation. Second, we assume that there are technological complementarities between dirty and clean R&D. This extension works in the opposite direction: firms more exposed to the shock will respond by scaling down clean R&D relative to other firms. In the standard model of directed technical change with clean and dirty innovation (see Acemoglu et al. (2012)), a fall in the fossil energy price increases the relative profitability of clean innovation, which gives an incentive for all firms to increase their clean R&D, regardless of what type of production and innovation they currently undertake. Our model takes seriously the notion that firms will react differently to a fall in the fossil energy price depending on how much they are involved in dirty production. However, because the adjustment costs and technological complementarities in R&D pull in opposite directions, we turn to the data to determine which of these forces dominates.

In our empirical analysis, we study the impact of a negative oil price shock on clean R&D in firms that are more exposed to the shock because they produce inputs used by oil-extracting firms. We use detailed micro data from Norway, where the fossil energy supply chain holds a central role. We exploit the large negative oil price shock in 2014, when the crude oil price dropped by approximately 60 percent. This quasi-natural experiment allows us to study resource reallocation dynamics in Norwegian manufacturing.

Our empirical analysis estimates the causal effect of the oil price shock on input suppliers' clean R&D investments relative to less exposed firms. We propose a novel method for identification in which we rely on firm-level trade data to compute a measure of firm exposure to the oil price shock. This firm-level exposure measure is based on the share of products in the firm's export basket that are used as inputs by oil-extracting firms, and allows us to compare firms in the same industry that otherwise have similar characteristics at baseline but are differentially exposed to the shock because of their

[4] This stylized fact is in line with existing empirical evidence; see, e.g., Bloom (2007).

[5] Adjustment costs can arise for many reasons, for example, in connection with research team dynamics or from search and hiring costs.

product mix. Hence, we isolate the effects of within-firm dynamics from the general effect of changes in relative prices, and we can control for potential confounding factors at the industry level.

We find that more exposed firms increased clean R&D significantly more than less exposed peers. This is true whether we look at the firm's likelihood of investing in clean R&D, the share of clean R&D, or the value of its clean R&D expenditures. These results are robust to using an alternative exposure measure based on firm-level production data. In light of our theoretical model, our results suggest that adjustment costs encourage exposed firms to shift their R&D towards clean innovation in the aftermath of a negative shock. If there is path dependency in technological development, this shift may turn out to have long-lasting effects on the speed of the transition. More broadly, our findings demonstrate that the substantial technological capabilities in the fossil energy supply chain can be leveraged in the clean transition, and that a fall in the profitability of the fossil energy supply chain will accelerate this process.

The paper makes two main contributions. First, we show that reduced expected future profitability induces firms in the fossil energy supply chain to shift their R&D investments towards clean R&D. Second, we propose a novel method for identifying firm-specific exposure to shocks. Our method relies on firm–product trade data, which are increasingly available to researchers, and is therefore a method that we believe can be applied to many other settings.

Our results have implications for our understanding of the effects of climate policy in general and carbon pricing specifically. A carbon price introduces a wedge between the consumer price and the producer price of dirty goods: The consumer price increases, and the producer price decreases. It is well understood that the increase in consumer prices incentivizes clean innovation by increasing demand for clean alternatives (Aghion et al., 2016). Moreover, the change in relative prices in favor of clean energy will shift innovation towards the more profitable sector (Acemoglu et al., 2012). Our results point to an additional channel through which carbon pricing may shift resources from dirty to clean R&D: By lowering profitability, a carbon price will lead firms in the fossil energy supply chain to further increase their clean R&D to avoid adjustment costs.

Although the focus of our empirical analysis is input suppliers in the Norwegian fossil energy supply chain, we conjecture that our results can extend to other settings. First, we expect the same mechanism to be salient for the oil-extracting firms themselves. Our focus on input suppliers is motivated by choice of empirical strategy and identification. Second, while it is clear that our main finding of increased clean innovation as a consequence of a drop in the oil price would not extend to firms globally in a naive way, our results may well apply to firms in the fossil energy supply chain in other countries facing a drop in expected future profitability.

The paper relates to several strands of literature. The theoretical framework builds 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.[6] We add to this literature by modeling firms investing in both clean and dirty R&D and including within-firm dynamics which give rise to additional responses to energy price changes.[7]

We also contribute to the literature on comparative advantage of firms and multi-product firms. Resource-based theories of the firm, such as Penrose (1955), suggest that firms grow by diversifying into products that require common capabilities. Bernard et al. (2010) show that firms are more likely to produce in certain pairs of industries, and Boehm et al. (2022) focus on the role of input similarities in driving these production choices. We focus on the role of firm capabilities in innovation and show that firms in the fossil energy supply chain can use their capabilities towards clean innovation when faced with the right incentives.

On the empirical side, the paper adds 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 found by Crabb and Johnson (2010), Calel and Dechezleprêtre (2016), Ley et al. (2016) and Hu et al. (2023), while Acemoglu et al. (2023) show how lower prices are generally negatively

[6]Acemoglu et al. (2012) underline the key role in climate policy of R&D subsidies, in addition to 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) uses 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 appropriate carbon pricing. Acemoglu et al. (2023) apply the framework to study the short versus long-term effects of the boom in shale gas production in the US. Aghion et al. (2025) analyze a model of green technological transition along a supply chain.

[7]The feature of parallel innovation in both dirty and clean is also present in Fried (2018), who presents a dynamic general equilibrium model with endogenous innovation, calibrated with historical oil shocks. Our paper complements Fried's in several ways. First, we model a specific channel through which a negative oil price shock induces within-firm reallocation of R&D investments. Second, we provide a well-identified empirical investigation of the effects of a negative shock to the expected future profits from oil production.

related to clean innovation. 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.[8] The main focus of this literature has been on how higher fossil energy prices drives demand for clean alternatives up. We show how declining expected profitability can reorient R&D within fossil energy supply chain firms, emphasizing the role of within-firm dynamics.

The paper also contributes to the broader literature on the role of carbon taxation in the green transition.[9] We show that carbon pricing may have a positive impact on the development of clean technologies not only through increased demand for these technologies but also through increased supply as profitability falls in the fossil energy supply chain.

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. 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 Background, 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 among manufacturing firms in Norway.[10] 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 a 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 for this decline was a combination of oversupply in the market and low global demand for oil. Shale oil production in the US increased substantially, while the Organization of the Petroleum Exporting Countries (OPEC) contributed to the

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

[9] 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. Furthermore, important recent contributions to the literature on optimal climate policy are, among others, Gerlagh et al. (2009), Stern and Valero (2021) and Blanchard et al. (2023).

[10] Lorentzen (2024) studies the impact of the same shock on reallocation of workers across industries and subsequent effects on earnings in destination industries.

oversupply by maintaining high production levels despite the falling price. At the same time, a global economic decline reduced global demand for oil, further exacerbating the supply–demand imbalance.[11]

Figure 1: Oil price: Brent blend

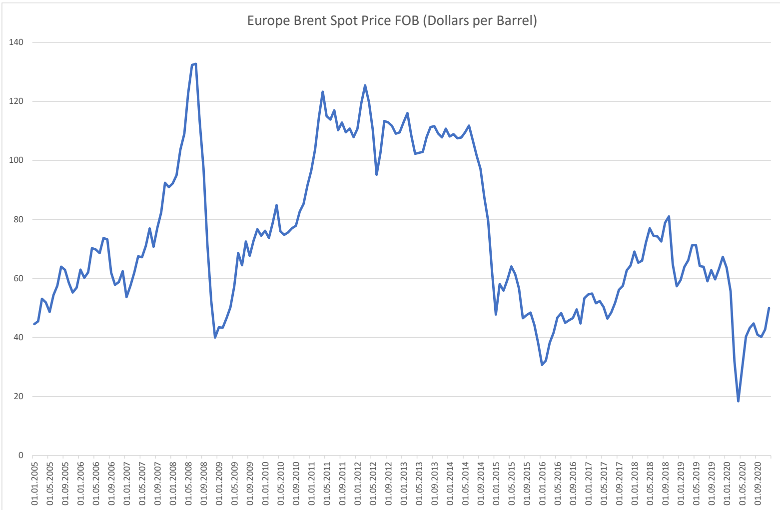

Note: The figure plots the European Brent spot price for oil per quarter for the period 2005–2020. Source: US Energy Information, https://www.eia.gov/dnav/pet/hist/LeafHandler.ashx?n=PET&s=RBRTE&f=D.

The result was a large negative shock to the current and future price of oil. Figure 1 shows the development of the Brent oil price over the period we study. It plummeted starting in June 2014. The dramatic decline was not expected. In Figure 2, the dotted line shows the future price of oil in Q2 2014, just before the price drop started: The price was expected to remain high for the coming years. The dashed line, on the other hand, shows the future price 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 important for our analysis, as the decisions firms make about R&D investments are forward-looking in nature and firms are unlikely to respond much to temporary price fluctuations.

[11]The relative importance of supply and demand factors in driving the shock is somewhat contested. While Baumeister and Kilian (2016) attribute the fall to factors on both sides of the market, Prest (2018) concludes that the global economic decline was the main driver. Kilian (2016) further discusses the role of the US shale oil revolution. Although we believe the mechanisms on which we shed light in this paper are relevant for understanding the consequences of a downturn in fossil energy supply chains outside Norway as well, the US shale oil industry faced a very different situation in 2014 from that faced by the Norwegian fossil energy supply chain.

2.2 The Norwegian Fossil Energy Sector

This oil price shock has appealing methodological properties in the setting we study. The fossil energy supply chain holds a central role in the Norwegian economy, accounting for approximately 60 percent of exports and 14 percent of private sector employment in 2014.[12] However, the oil price is set in the world market and is completely exogenous from the perspective of firms in a small open economy such as Norway.[13] Moreover, the Norwegian fossil energy sector is exclusively based on offshore activity. The technology, intermediates and capital goods used for offshore oil extraction differs substantially from the technology used for standard onshore as well as shale oil extraction. It is therefore highly unlikely that Norway's input suppliers experienced increased demand for their products as a consequence of the positive shale oil shock in the US.

Figure 2: Crude oil price and oil future prices

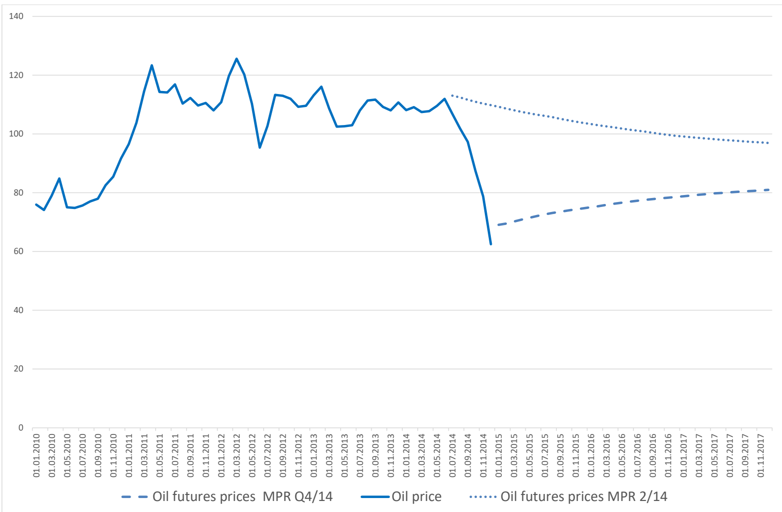

Note: The figure plots the crude oil price until end 2014 and oil future prices from 2014 to 2017. Source: The Norwegian Central Bank Monetary Policy Reports 2/2014 and 4/2014.

Instead, the fall in the oil price in 2014 had a substantial negative impact on expected output growth for input suppliers in Norway. Figure 3 reports the results from the Business Sector Survey conducted by the Norwegian Central Bank for January–August 2015. According to the survey, expectations varied significantly across industries, with

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

[13] Norway's global market share of crude oil is approximately 2 percent. See https://www.norskpetroleum.no/en/production-and-exports/exports-of-oil-and-gas/.

a positive outlook in most industries, while the industries serving and supplying the oil industry stand out, with a pessimistic outlook for both domestic sales and exports.

Hence, the 2014 oil price drop serves a useful quasi-natural experiment that allows us to study resource reallocation in manufacturing firms differentially exposed to the shock through their supply linkages to the oil-extracting firms.[14]

Figure 3: Expected output growth

Note: The figure shows expected annual output growth the coming six months by industry in the Norwegian economy based on the Norwegian Central Bank's regular Survey of the Business Sector. Source: The Norwegian Central Bank Monetary Policy Report 3/2015.

2.3 Data

Our empirical analysis is based on five datasets which we link using a unique firm identifier. The first is administrative firm register data from Statistics Norway, which cover the universe of firms across all sectors. The firm register provides information on the date of the entry and exit of each individual firm, as well as the number of employees and the industry classification.

The second dataset 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 the data cover the near-universe of firms in Norway. The income statement and balance sheet data are based on

[14]One might be concerned that the global economic slowdown affected the Norwegian economy. We will discuss this further in Section (4) when we describe the empirical strategy, which is designed to isolate the effect of the fall in oil prices from other, potentially confounding effects.

data from annual accounting reports that, according to Norwegian law, must be filed with the public Register of Company Accounts. From this dataset we get information on tangible assets, and the share of energy costs in operational costs.

The third dataset is the R&D survey.[15] The survey provides information about the value of total R&D expenditures and the number and wage bill of R&D personnel. Importantly, the survey also gives information about the share of clean R&D expenditures. We provide details on the R&D survey and how we define clean R&D in Appendix Section A. Note that the survey does not provide a good measure for dirty R&D, so in our empirical analysis will divide total R&D expenditures into clean and non-clean R&D.

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

The fifth dataset provides complete information on all direct support from the government for R&D and innovation to firms.[16]

2.4 Sample Selection

For our analysis, we focus on all joint-stock firms in manufacturing (defined by Nomenclature of Economic Activities (NACE) industry codes 10 to 35) that are covered by the R&D survey. We do not include the firms directly hit by the fall in oil price, namely, oil producers (NACE code 6).[17] We focus on the manufacturing industries, 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 code 6) to construct our exposure measure as described below in Section 4.1. Our sample is constructed to cover 2007–2017. The years 2007–2013 define the pre-shock period, and 2014–2017 define the post-shock period. The final sample used for the stylized fact and empirical analysis is an unbalanced panel of approximately 1,300 firms per year.

[15] An alternative outcome variable we could have used is patents. However, the use of patent data has some disadvantages. First, there is a great deal of variation across industries in the propensity of an innovation to be patented. Second, patent data are produced with a time lag. Given the time span of our data, we would not have much of a post-shock period to look at had we used patent data.

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

[17] We also leave out firms in two other industries: NACE code 12 and 19. The first is tobacco products, which is a sparsely populated industry in Norway. The second is Manufacture of coke and refined petroleum products, i.e., the sector further downstream from the oil-extracting companies. Since our focus is on the upstream suppliers, we exclude these firms from the analysis.

2.5 Stylized Facts on R&D and Clean R&D

We present three stylized facts on the development of R&D in the manufacturing industries in the aftermath of the oil price shock. These facts will guide both our theoretical model and our empirical analysis.

Fact 1: After the oil price shock, the share of firms investing in clean R&D increased in both 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. By 2017, these numbers had increased to 46 percent for overall R&D, whereas the share of firms investing in clean R&D had more than doubled, to 11 percent. Figure 4 shows the increase in the share of firms with clean R&D expenditure relative to all firms engaged in R&D.

Figure 4: The oil price and the share of R&D firms with clean R&D

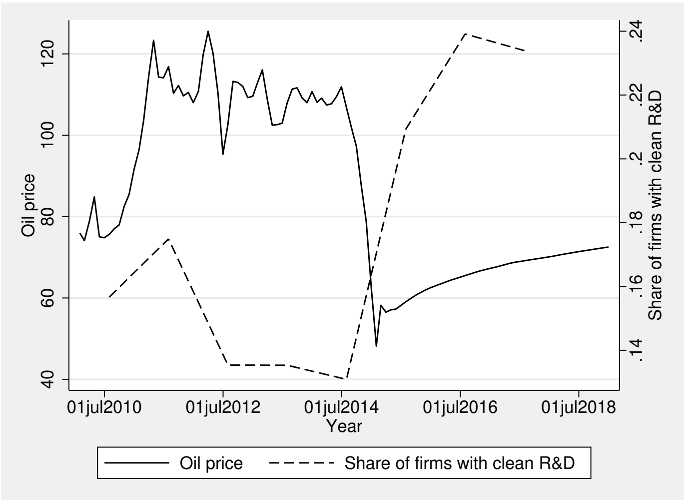

Note: The figure shows the share of firms in the baseline sample that reports investments in clean R&D relative to all firms engaged in R&D, the evolution of the oil price until Q1 of 2015, and the oil future price for the remaining time period. The sample covers 2010–2017.

Fact 2: Most of the increase in clean R&D came from firms with ongoing investments in non-clean R&D. Figure 5 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 that stop doing R&D (including exiting firms), and firms continuing their R&D. For each of these categories, we split firms into two groups by whether their R&D investments are 100 percent clean or a combination of clean and non-clean. As the graph shows, most of the increase in clean R&D after the oil price shock came from firms with both clean and non-clean R&D.

Figure 5: Decomposition of clean R&D

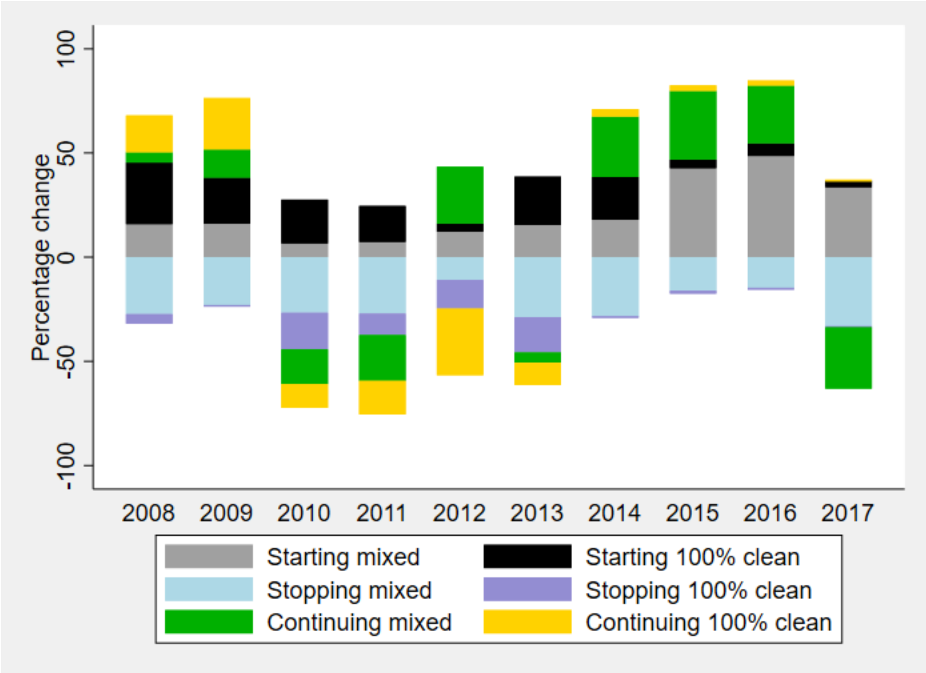

Note: The figure shows the annual percentage change in clean R&D – decomposed by firm R&D status, i.e., whether the firm starts, stops or continues investment in both clean and dirty R&D ("mixed"), or starts, stops or continues investment in only clean R&D. The sample covers 2007–2017.

Anecdotal evidence suggests that capabilities in the fossil energy supply chain can be useful for clean technologies.[18]

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

[18]For example, IEA (2024) underlines the role of gas fracking technology development for the technology frontier for geothermal energy. For examples from the Norwegian context, see, e.g., https://www.norskpetroleum.no/en/developments-and-operations/service-and-supply-industry/, https://www-innovasjonnorge-no.translate.goog/artikkel/teknologioverforing-i-praksis?_x_tr_sl=auto&_x_tr_tl=en&_x_tr_hl=en-US&_x_tr_pto=wapp&_x_tr_hist=true.

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 2007–2017.

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

3 A Theoretical Model of Clean and Dirty Innovation

In this section, we develop a stylized model of clean and dirty innovation in firms, drawing on the literature of directed technical change and the environment (Acemoglu et al., 2012). The framework allows us to illustrate how and why dirty input suppliers may react differently to an oil price drop compared to other firms. We use the model to guide our empirical analysis. The basic model presented in this section is kept simple, in order to illustrate a few key mechanisms. In Appendix B, we extend this basic model in several directions, to illustrate that the conclusions regarding the effect of the within-firm dynamics on the response in clean R&D to a drop in the dirty energy price do not hinge on the simplicity of the model presented here.

Informed by the three stylized facts presented above, we introduce two types of within-firm dynamics in a model with heterogenous firms. First, we allow for technological complementarity between clean and dirty R&D. We aim to capture the existence of economies of scope in R&D, i.e. that dirty R&D focused towards fossil energy production may also be valuable for developing technologies for renewable energy production or energy saving.[19] Second, in line with Bloom et al. (2013b), we assume that there are adjustment costs related to the scaling of a firm's overall R&D. We view these adjustment costs as capturing both the direct costs of firing scientists and the broader economic and

[19]Between-firm spillovers are well documented in the literature; see specifically Bloom et al. (2013a). Many of the mechanisms proposed in this literature are, however, also likely to be present within firms.

behavioral responses within a scientific team facing change.[20]

3.1 Model

Consider a firm, denoted k, that produces inputs for the use in clean and/or dirty energy production, denoted by $j = c, d$, respectively. The time horizon is infinite, with discrete time periods and a discount factor $\beta \in (0, 1)$.

When there has not been a new innovation in the previous period, there are zero profits for the input suppliers.[21] However, the firm can invest in R&D of each type, $j = c, d$. If the R&D leads to an innovation in period t, the firm obtains a patent for the input and is allocated monopoly rights for the next time period, resulting in a profit of $\pi(p_{jt+1})$,where p_{jt+1} is the price of the final energy good of type j, in period $t+1$. Assume $\pi' > 0$, reflecting that a higher price of dirty energy increases demand for dirty-energy inputs, while a higher price of clean energy increases the demand for clean-energy inputs. We study a small, open economy, and treat current and future dirty and clean energy prices as exogenous.

Firms are heterogeneous, and the probability of a successful innovation of type j in firm k depends on a firm- and type-specific innovation productivity and on a type-specific innovation production function. The innovation productivity of firm k for innovation of type j is given by $\eta_{kj} \in \{0, \eta_j\}$, with $\eta_j \in (0, 1)$ while the innovation production function for innovation of type j is concave and given by $g^j(\cdot)$. We assume that a firm with $\eta_{kj} = 0$ will not be able to innovate for type j and will therefore not produce type-j inputs.

To innovate, the input suppliers must hire scientists. Scientists have skills that can be used for both types of R&D.[22] Let w_t be the wage for scientists in period t, and let $\tilde{s}_{kdt}$ denote the number of scientists engaged in type-j R&D in firm k in period t. The probability of a dirty innovation in firm k in period t is given by $\eta_{kd} g_d(\tilde{s}_{kdt})$, with $g_d(0) = 0$, $g_d'(\cdot) > 0$, $g_d'' < 0$, while the probability of a clean innovation is given by $\eta_{kc} g_c(\tilde{s}_{kct}, \tilde{s}_{kdt})$, with $g_c(0, \tilde{s}_{kdt}) = 0$, $\partial g_c / \partial s_{kct} > \partial g_c / \partial s_{kdt} \geq 0$, $\partial^2 g_c / \partial s_{kct}^2 < 0$, $\partial^2 g_c / \partial s_{kdt}^2 \leq 0$ and $\partial^2 g_c / \partial s_{kct} \partial s_{kdt} \geq 0$. The probability of clean innovation depends not only on the number

[20] For instance, managers may be reluctant to let go of scientists with hard-to-replace specialized skills or to disrupt the dynamics of a well-functioning team. Team cohesion means that a team's collective performance can exceed the sum of individual contributions, so losing one or more members may reduce the productivity of those who remain. Relatedly, the literature on worker co-mobility and "aqui-hiring" (e.g., Marx and Timmermans, 2017) suggests that managers may be hesitant to break up teams, fearing that it could lower morale, harm productivity, or trigger further departures. Managers may also act in their own career interests, preferring to maintain the size of their teams even when downsizing would benefit the firm.

[21] We abstract from the process of production of inputs in this basic model. In Appendix B.4, we include the within-firm dynamics modeled here in the framework of Acemoglu et al. (2012) of directed technical change, including modeling of the production and profits of the firms.

[22] We here follow the standard assumption in the literature on directed technical change.

of scientists hired to undertake clean R&D, but also on the level of dirty R&D, i.e. there is a a technological complementarity between the more mature dirty R&D and the less mature clean R&D.[23]

Finally, we allow for adjustment costs to arise if the firm hires or fires new scientists within a given time period. Define s_{kt} as the total number of scientists hired in the firm at the beginning of period t. We include a convex cost of rescaling total R&D within the firm:[24]

$$R_{kt} = \frac{1}{2}c \cdot \left(\frac{\tilde{s}_{kt} - s_{kt}}{s_{kt}}\right)^2 \tag{1}$$

with $c \geq 0$.[25] We assume that the firm can costlessly change the number of scientists hired by the end of the time period, s_{kt+1}.

At the beginning of each time period, the firm makes three decisions: How many scientists to start off with in the next time period, s_{kt+1}, and how many scientists to engage in each type of R&D in the current time period, $\tilde{s}_{kct}$ and $\tilde{s}_{kdt}$. We assume throughout that the firm expects the price of energy of type j in any future period to equal the observed price in the current time period, and maximizes profits given these price expectations. Let $p^e_{jt+1} = p_{jt}$ denote the expected price of energy of type j in period $t+1$, and let w^e_{t+1} denote the expected wage for the next period.[26]

3.2 Equilibrium Firm Behavior

In period t, a firm k with $\eta_{kj} > 0$ for $j = c, d$ maximizes profits from its R&D efforts by solving:

$$\begin{aligned} V_{kt}(s_{kt}) = \max_{s_{kt+1}, \tilde{s}_{kct}, \tilde{s}_{kdt}} \Bigg\{ & - w_t \tilde{s}_{kt} - \frac{1}{2}c \cdot \left(\frac{\tilde{s}_{kt} - s_{kt}}{s_{kt}}\right)^2 \\ & \beta \Big[\eta_{kc} g_c(\tilde{s}_{kct}, \tilde{s}_{kdt}) \pi_c(p^e_{ct+1}) + \eta_{kd} g_d(\tilde{s}_{kdt}) \pi_d(p^e_{dt+1}) + V_{kt+1}(s_{kt+1}) \Big] \Bigg\} \end{aligned} \tag{2}$$

[23]We further assume that $lim_{\tilde{s}_{kdt}\to\infty} g_d(\tilde{s}_{kdt}) \leq 1$, $\lim_{\tilde{s}_{kct}\to\infty, \tilde{s}_{kdt}\to\infty} g_c(\tilde{s}_{kct}, \tilde{s}_{kdt}) \leq 1$, and that $\lim_{\tilde{s}_{kdt}\to 0} g'_d(\tilde{s}_{kdt})$ and $\partial \lim_{\tilde{s}_{kct}\to 0} g_c(\tilde{s}_{kct}, \tilde{s}_{kdt}) \partial \tilde{s}_{kct}$ are large enough to ensure interior solutions to the problems presented in the following for firms with strictly positive innovation productivities. To keep the basic model as simple as possible, we include complementarity only in the clean-innovation probability. In Appendix B.2, we allow for $\tilde{s}_{kjt}$ to enter in the same way into the dirty-innovation probability and show that the mechanisms we describe do not depend on this simplification.

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

[25]In Appendix B.3, we extend the model by allowing adjustment costs to arise also when the firm makes adjustments to the number of scientists applied for each type of R&D, and show that our conclusions do not depend on abstracting away from this type of adjustment costs.

[26]The firm has rational expectations with respect to the future wage, which is determined in equilibrium.

The solution to the problem determines $s_{kt+1}(p^e_{ct+1}, p^e_{dt+1}, w^e_{t+1})$, $\tilde{s}_{kct}(p^e_{ct+1}, p^e_{dt+1}, w_t)$ and $\tilde{s}_{kdt}(p^e_{ct+1}, p^e_{dt+1}, w_t)$.

Only firms with positive type-specific innovation productivity will innovate in that type of R&D. Let h or l denote a firm with $\eta_{hd} = 0$ or $\eta_{lc} = 0$, respectively. Firm h will be active only in clean R&D and production, while firm l will be active only in dirty R&D and production.[27] For these firms, the problem simplifies accordingly, and the solution to the firm's problem determines the pairs $s_{ht+1}(p^e_{ct+1}, w_{t+1})$, $\tilde{s}_{hct}(p^e_{ct+1}, w_t)$ or $s_{lt+1}(p^e_{dt+1}, w_{t+1})$, $\tilde{s}_{ldt}(p^e_{ct+1}, w_t)$, respectively.

We show in Appendix B.1 how the different types of firms respond to changes in the expected future energy prices and to a change in the scientist wage. The direct effect of a drop in the dirty energy price is a fall in the number of scientists engaged in dirty R&D. For a firm engaged only in clean R&D and production (or only in dirty), there is no direct effect on R&D of a drop in the dirty-energy price. For firms engaged in both types of R&D, the within-firm dynamics create a link between clean R&D and the dirty energy price. In Appendix B.1 we show that the two following propositions hold:

Proposition 1. *In a firm engaged in both dirty and clean production and R&D, the direct response to a drop in the dirty energy price will be positive if the effect of the adjustment costs dominates the effect of technological complementarity:*

$$-\frac{\partial \tilde{s}_{kct}(p^e_{ct+1}, p^e_{dt+1}, w_t)}{\partial p_{dt}} > 0 \ \textit{if} \ \frac{c}{s^2_{kt}} > \beta \eta_{kc} \pi \frac{\partial^2 g_c}{\partial \tilde{s}_{kct} \partial \tilde{s}_{kdt}}. \tag{3}$$

If, on the other hand, the technological complementarity dominates the effect of the adjustment costs, the direct response in clean R&D to the dirty energy price drop in the firm will be negative:

$$-\frac{\partial \tilde{s}_{kct}(p^e_{ct+1}, p^e_{dt+1}, w_t)}{\partial p^e_{dt+1}} < 0 \ \textit{if} \ \frac{c}{s^2_{kt}} < \beta \eta_{kc} \pi \frac{\partial^2 g_c}{\partial \tilde{s}_{kct} \partial \tilde{s}_{kdt}}. \tag{4}$$

The fossil energy price drop reduces the expected profits from investment in fossil R&D, inducing a firm engaged in this activity to reduce $\tilde{s}_{kdt}$. Because the marginal cost of adjusting the total number of scientists in the firm is increasing, the actual marginal cost of engaging researchers in clean R&D falls for a firm that is engaged in both types of R&D. Therefore, if these adjustment costs are dominating, the firm will increase its clean R&D as a direct response to the fossil energy price drop. On the contrary, the

[27] The case with $\eta_{kd} = \eta_{kc} = 0$ is not relevant, as a such a firm would not be active in any of the two sectors.

technological complementarity implies that the fall in $\tilde{s}_{kdt}$ induced by the fall in the fossil energy price instead decreases the marginal value of scientists in clean R&D in the firm. If the technological complementarity is dominating, the firm will thus decrease its clean R&D as the direct response to the shock.

The overall demand for scientists, for a given wage, decreases as a consequence of a dirty energy price fall, as we show in Appendix B.1. In a model where the supply of scientists is not infinitely elastic, the wage will fall. The general-equilibrium effect of a drop in the dirty energy price on a firms' clean R&D will be the sum of the direct response and the response to the drop in the scientist wage:

$$d\tilde{s}_{kct}/dp^e_{dt+1} = \partial\tilde{s}_{kct}/\partial p^e_{dt+1} + \partial\tilde{s}_{kct}/\partial w_t \cdot \partial w_t/\partial p^e_{dt+1} \tag{5}$$

$$d\tilde{s}_{hct}/dp^e_{dt+1} = \partial\tilde{s}_{hct}/\partial w_t \cdot \partial w_t/\partial p^e_{dt+1}. \tag{6}$$

Proposition 2 summarizes the total response in clean R&D of a firm that is engaged in both types of R&D, relative to a firm only engaged in clean R&D. Calculations are provided in Appendix B.1.[28]

Proposition 2. *The overall response in clean R&D to a drop in the dirty energy price will be larger in a firm engaged both in dirty and clean production and R&D than in a firm that is only engaged in clean R&D and production if the effect of the adjustment costs dominates the effect of technological complementarity:*

$$-\frac{d\tilde{s}_{kct}(p^e_{ct+1}, p^e_{dt+1}, w_t)}{dp^e_{dt+1}} > -\frac{d\tilde{s}_{hct}(p^e_{ct+1}, w_t)}{dp^e_{dt+1}} \;\; \textit{if} \;\; \frac{c}{s^2_{kt}} > \beta\eta_{kc}\pi\frac{\partial^2 g_c}{\partial\tilde{s}_{kct}\partial\tilde{s}_{kdt}}. \tag{7}$$

If, on the other hand, the technological complementarity dominates the effect of the adjustment costs, the overall response in clean R&D to the dirty energy price drop will be lower in a firm engaged in both dirty and clean R&D and production than in a firm engaged only in clean R&D and production:

$$-\frac{d\tilde{s}_{kct}(p^e_{ct+1}, p^e_{dt+1}, w_t)}{dp^e_{dt+1}} < -\frac{d\tilde{s}_{hct}(p^e_{ct+1}, w_t)}{dp^e_{dt+1}} \;\; \textit{if} \;\; \frac{c}{s^2_{kt}} < \beta\eta_{kc}\pi\frac{\partial^2 g_c}{\partial\tilde{s}_{kct}\partial\tilde{s}_{kdt}}. \tag{8}$$

[28] In Appendix B.1, we assume that there is a fixed pool of scientists available in the market and that there are three types of firms with the same size before the shock: A share α^l of the firms with $\eta_{lc} = 0, \eta_{ld} = \eta_d > 0$, a share α^h with $\eta_{hc} = \eta_c > 0, \eta_{hd} = 0$, and the remaining share $\alpha^k = 1 - \alpha^l - \alpha^h$ with $\eta_{kct} = \eta_c > 0, \eta_{kdt} = \eta_d > 0$. We further assume that $\alpha^h + \alpha^k$ is not close to zero. In Appendix B.1, we discuss a special case that can arise if the effect of the technological complementarity dominates the effect of the adjustment costs, while the market is strongly dominated by firms that are engaged only in dirty R&D and production, resulting in firms engaged in clean production being hit mainly by a wage shock rather than a price shock. This special case is unlikely to be relevant in our empirical setting.

In the next section, we take these propositions to the data and investigate how clean R&D investments in firms supplying input to oil-extracting firms respond to a drop in the price of oil, relative to similar firms that are not engaged in production of inputs for oil extraction. Note that in the theoretical framework, we have deliberately focused on the intensive margin to avoid corner solutions. We do this to highlight the key mechanisms that we propose, which would also arise in a more complicated setup allowing for corner solutions and extensive-margin responses. In the empirics that follow, we will also consider the extensive margin; that is, we do not limit the analysis to firms already conducting clean 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 firms engaged in dirty production, which we will refer to as exposed firms, relative to non-exposed firms is ambiguous, and depends on the strength of the within-firm dynamics. [29] It thus remains an empirical question. To investigate this question, we develop a novel measure of identification that lets us exploit firms' heterogeneous exposure to the shock in our analysis.

4.1 Identification: Measuring 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 industry-level input–output tables, which provide the share of production in an industry i 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, there is considerable variation in our exposure measure within industries. We would not be able to capture this heterogeneity if we relied on input–output tables. Second, if our treatment were defined at industry level, we would not be able to control for any other contemporaneous industry-level shocks.

Therefore, we instead develop a novel firm-specific exposure measure leveraging our

[29] Note that in the theoretical framework, a dirty input supplier will generally, by assumption, also be active in dirty R&D. This is not necessarily true in the data, where we identify exposed firms through their product mix. As long as a firm's R&D investments are related to its product mix, the theoretical framework is still informative about the interpretation of the empirical findings.

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 we construct our firm-specific exposure measure by following these steps:

1. Identify the HS8 products imported by the oil producers (NACE code 6) in the pre-shock period (2007–2013).[30] We denote these *oil-related* products ($j \in o$).[31]

2. Identify the input suppliers in manufacturing that export oil-related products.

3. Calculate firm-level exposure x_{ok} as the share of oil-related products 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 Appendix Section C.

Measuring exposure to the shock in this way has some clear advantages over the use of industry-level input–output tables. First, it captures the fact that the oil price shock was a global shock. Instead of measuring exposure only through domestic linkages, we capture the extent to which firms were exposed to the oil price drop through international trade. Second, and crucially for identification, this approach yields firm-level variation in exposure, allowing us to control for other, potentially time-varying factors affecting all firms in an industry.

Our using this exposure measure does come at a cost, as it reduces the sample to only goods-exporting firms in the manufacturing industries 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 oil-extracting firms is an appropriate measure of the share of the firm's activity devoted to these products.

Figure 6 highlights the importance of our using a firm-specific measure. The box plot describes the within-industry variation in the exposure measure for all 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 those in the rest of the distribution.

[30] 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.

[31] Alternatively, we could have picked only products disproportionately imported by the oil producers relative to firms in other industries. We decided against this for two reasons. First, we are not confident that the oil producers' share of imported inputs necessarily reflects the share of inputs used in production. Second, our preferred measure is more conservative and works against our finding an effect of the oil price shock. We discuss this further in Appendix C.

The box plot shows substantial variation in the exposure of firms within practically all of the manufacturing industries (NACE 2-digit level). An alternative approach would be to follow the commonly used classification from Statistics Norway, which defines industries directing at least 50 percent of the value of their production to the petroleum sector as oil-related.[32] This would, in practice, mean that we would classify all firms in industry codes 25, 26, 27, 28, 30 and 33 as exposed to the shock (and firms in all other industries as not exposed). We find it reassuring that these industries also stand out as being strongly affected on the basis of our exposure measure. However, as Figure 6 shows, there is also substantial exposure among firms in the other industries. In particular, NACE codes 22 (manufacture of rubber and plastic products), 29 (manufacture of motor vehicles, trailers and semi-trailers), 31 (manufacture of furniture) and 32 (other manufacturing) stand out as having medians close to 1.

Figure 6: Within-industry variation in exposure

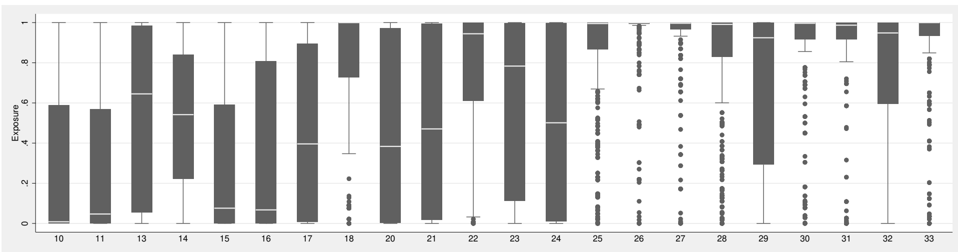

Note: The figure shows the variation in the exposure measure described in Section 4.1 across manufacturing industries (NACE 2-digit level). The median of the exposure measure is shown by the line within the box, and the box's boundaries represent the first (Q1) and third (Q3) quartiles, indicating the 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

To estimate the impact of the oil price shock, we specify a difference-in-difference (DID) 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 the outcomes before and after 2014 for firms differentially affected by the oil price shock. The baseline specification takes the form:

$$y_{kt} = \alpha_k + \beta x_{ok} \times Post_t + \gamma Z_{kt} + \delta_{st} + \varepsilon_{kt} \tag{9}$$

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

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 its export portfolio (as described in Section 4.1).[33] α_k is firm fixed effects, capturing all time-invariant differences between the more 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.[34] 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 for years in the the post-period ($t > 2013$). Z_{kt} is a vector of firm-level controls aimed at ensuring that the more and less exposed firms are otherwise 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 differently hit by the global oil price shock because of their exposure to trade, particularly the 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.[35]

As pointed out above, we limit our analysis to firms in the manufacturing industries. That is, we leave out the oil producers (NACE code 6), which were directly affected by the oil price shock, since we do not have a good control group for these firms. 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 because of their use of energy as an input in production. Manufacturing firms in Norway typically rely on two main sources of energy: hydroelectricity, and fossil fuels for transport and some types of machinery. In Appendix Section E, we show how firms' costs of energy have developed over time for each of these energy goods and in aggregate; see Figure 8. The main takeaway is that electricity prices developed smoothly over the relevant horizon. Nevertheless, to control for potential differences in exposure to the shock on the cost side, we add to the specification a control variable capturing firms' energy cost share. This is measured as cost of energy related to production and fuel for transportation, relative to operational costs. Finally, to alleviate concerns that firms might differ in the degree to which they receive financial support for R&D, we add a dummy for whether the firm has received government 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.

[33]To be precise, this is a DID specification with a continuous treatment variable. We discuss the implications of using a continuous treatment variable in Section 5.2 below. Our approach is conceptually similar to a shift-share design, but with a one-off common shock.

[34]The results are robust to instead using NACE 3-digit industry codes.

[35]The results are robust to our instead using the net export share (export share – import share) as a control variable.

5 Empirical Results on Clean Innovation

We start by presenting the main results on the effects of the oil price shock on clean R&D. We then proceed to discuss various threats to identification and robustness checks.

5.1 Main results

Our main outcome variables of interest are related to clean R&D. We start by analyzing the effect on the firm's 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 our estimating different versions of Equation (9) with and without controls and fixed effects can be found in Table 2.

In column (1), we include only industry–year fixed effects. We find that the shock leads to an increase in the likelihood of a firm's reporting any clean R&D investment of approximately 4.5 percentage points for exposed firms.[36] The baseline likelihood is 7.5 percent, meaning that this effect is substantial and of economic importance. In column (2), we add firm fixed effects, which attenuate 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 explanatory power. That the point estimates change when we include the energy share in the regression is not surprising, as firms that use substantial energy in production may clearly also be affected by the oil price drop through this direct channel. Our results suggest that firms more exposed to the oil price shock because of their delivery of inputs to oil-extracting firms were more likely to invest in clean R&D after the shock.

[36] We discuss potential issues arising from the fact that our treatment variable is a continuous variable below in Section 5.2.

Table 2: Probability of clean R&D

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

Note: Standard errors in parentheses are clustered by 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 a firm's 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 identify this, 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 column (1) in Table 3 below. We find a positive and significant effect on this measure, as well, albeit not as precisely estimated as the coefficient on the dummy. The baseline share of clean R&D is 6 percent, so again, this is an economically significant effect. These findings suggest that firms more exposed to the oil price shock increased their clean relative to their non-clean investments relatively more than less exposed firms. In column (2), we replace the share of clean R&D with the value of clean R&D. Here, we find support for the hypothesis that exposed firms also increased their clean R&D spending in absolute terms. The coefficient from column (2) translates into an increase in the value of clean R&D spending of approximately 35% for firms exposed to the shock.

Table 3: Clean R&D: Share and value

Variable:	Share (1)	Log Value (3)
$Post_t{*}x_{oi}$	1.575*	0.346**
	(0.899)	(0.140)
Controls	Yes	Yes
Firm FE	Yes	Yes
Ind.*year FE	Yes	Yes
Observations	11,695	11,695

Note: Standard errors in parentheses are clustered by 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 input suppliers reacted relatively more strongly to the price shock in terms of investments in clean R&D.

We investigate further where the increase in clean R&D investments originated: Was it due mainly to firms that had not previously invested in any R&D that started investing, or to firms with preexisting R&D investments that reallocated their R&D spending towards clean R&D? To this end, we create two dummy variables: First, *New Clean* equals 1 if the firm reports spending on clean R&D in period t but not in period $t-1$. Second, *R&D to Clean* equals 1 if the firm reports spending on clean R&D in period t but not in period $t-1$, conditional on the firm's having reported positive R&D spending (on any type of R&D) in period $t-1$.

The results from this exercise are found in Table 4. In column (1), we see a strong and positive coefficient, indicating that firms more exposed to the shock are significantly likelier to start new clean R&D projects than less exposed firms . In columns (2) and (3), we show the results for the *R&D to Clean* dummy. In column (2), all firms are included, regardless of whether they report any R&D spending, In column (3), we limit the sample to firms reporting positive R&D investments before the shock. We again find that firms 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 only R&D performers.

Table 4: Switching to 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 parentheses are clustered by firm. *New Clean* is a dummy that takes 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* but conditions on the firm's 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. $*p < 0.1, **p < 0.05, ***p < 0.01$.

In summary, we find that firms more exposed to the 2014 fall in the oil price through the oil producer supply chain increased their clean R&D more than firms less exposed to the shock. This is true whether we look at the likelihood of a firm's 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 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 how robust these findings are and what mechanisms are at play in- producing them.

5.2 Robustness

Exposure measure. The exposure measure is based on firms' export share. For identification, we rely on this export share being representative of the firms' production share as explained above. For a subset of manufacturing firms, we also have information on product-level output through the PRODCOM survey which we can employ as a robustness check.[37] First, we map the PRODCOM product classification to the HS 8 product classification using the concordance from Magerman (2022). We then calculate exposure shares based on the share of production at the firm level which is oil-related, using the same definition of oil related products as above. We rerun our main specifications (Tables 2 and 3) using this alternative exposure measure.[38] We find that the coefficients of interest are practically unchanged, which reassures us that the export-based exposure measure is reflecting relevant variation in the production shares across firms.

[37] For a description of this survey, see https://ec.europa.eu/eurostat/cache/metadata/EN/prom_esms_at.htm#shortd

[38] Results are available upon request.

Pre-trends. A potential concern is that more exposed and less exposed firms faced different pre-trends. It would be reassuring to find similar pre-trends for firms with different exposure levels, which would imply 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:

$$y_{kt} = \alpha_k + \beta x_{ok} \times \delta_t + \gamma Z_{kt} + \delta_{st} + \varepsilon_{kt} \tag{10}$$

For ease of exposition, we show the graph for a binary treatment variable here.[39] The result from our estimating this specification is illustrated in Figure 7. First, the lack of evidence of any sort of pattern in the pre-trends is reassuring and suggests that our treatment and control firms behaved similarly in our pre-period. We see a clear upward trajectory beginning in 2014 and turning significant from 2015.

Figure 7: Dynamic DID

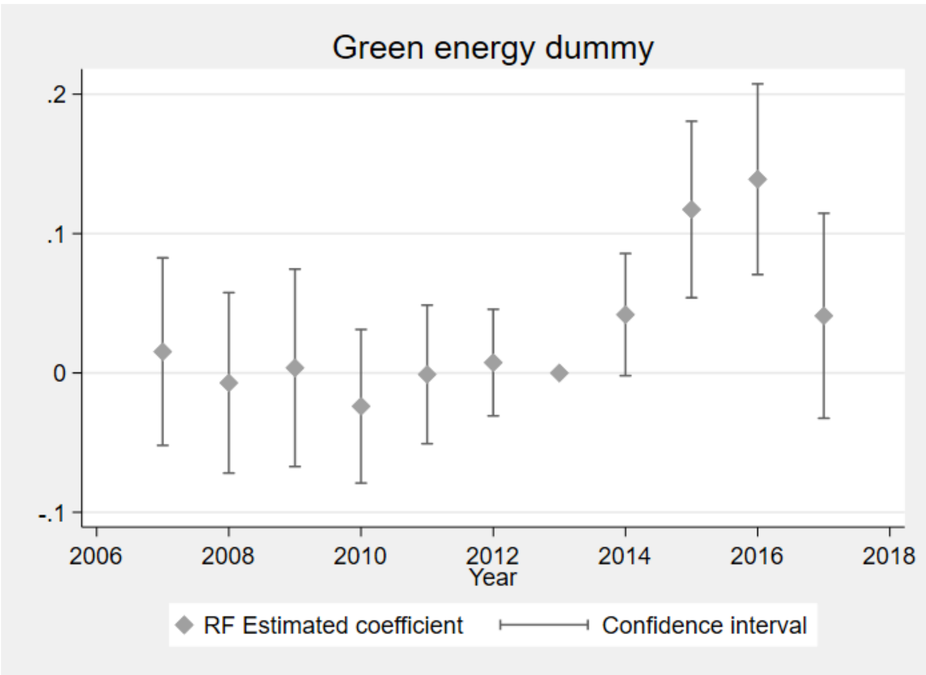

Note: This figure plots the yearly coefficient and 90 percent confidence intervals of the event study DID estimator in equation Equation (10).

Renewable energy prices. Although our dynamic DID estimation results are reassuring, we still cannot rule out the possibility that our results are driven by something other than the oil price. In particular, if clean energy prices increased substantially during the 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 affected by changes in

[39] The graph for the continuous treatment variable looks very similar but is harder to interpret.

clean energy prices differently from less exposed firms. However, Figure 9 in Appendix Section E shows how clean energy prices evolved over our sample period. If anything, they declined smoothly, making the above conjecture less of a concern. It is possible that the adoption of clean energy technology has increased in this time frame, however it is unlikely to drive the sharp increase in clean R&D for the exposed firms in the aftermath of the shock we see in Figure 7.

Continuous exposure. Our exposure measure is continuous, which has been shown by Callaway et al. (2024) to complicate inference in a two-way fixed effects setting. 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 the standard parallel trends assumption: namely, that the average evolution of the outcomes that firms with any exposure would have experienced without treatment is the same as the evolution of outcomes that firms in the unexposed group actually experienced. We reestimate all our results using a binary dummy variable that takes the value of 1 for firms with positive exposure and 0 for unexposed firms. The results (available upon request) are very similar to our baseline results, suggesting that our coefficient of interest can be interpreted as the average level treatment effect.

Government R&D support. In our main specification, we include a dummy for whether the firm has received government funding for R&D and innovation, measured at baseline and interacted with year dummies. We do this to control for the possibility that firms differ in the degree to which they receive financial support for R&D. One might be concerned that the shock triggered changes to the types or generosity of such funding schemes, or that firms increasingly make use of them. We have found no indications of relevant policy changes. However, to alleviate any remaining concerns, we replace the time-invariant control variable with a time-varying dummy. The results are robust.

Choice of estimator. In Appendix Table 8, we reestimate our main specifications using a Poisson maximum likelihood (PPML) estimator instead of ordinary least squares (OLS). We do this to mitigate two concerns. First, when we use a dummy variable as the dependent variable, one might be worried that a linear model is not the best choice. Second, when we use a log transformation of the dependent variable, observations with

Table 5: Placebo

Variable:	Biotech 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 parentheses are clustered by 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. $*p < 0.1, **p < 0.05, ***p < 0.01$.

zero values are an issue. PPML allows us to deal with both of these concerns.[40] The results in Table 8 show that our results from this alternative specification are robust, although the estimated coefficient on the value of clean R&D expenditures is just below the threshold for statistical significance at the 10% level.

Industry trends. One might worry that the NACE 2-digit industry classification is too aggregate and will group together firms that are in reality exposed to very different shocks. As a robustness check, we experiment with using NACE 3-digit industry trends instead. The results are robust to this. Unfortunately we cannot go more disaggregate than this without losing statistical power, although the coefficients remain similar in magnitude.

Reallocation towards other technology fields. 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 biotech and information and communication technology (ICT) R&D. We estimate Equation (9) with these two new variables as dependent variables. Table 5 reports the results. The estimated coefficients are negative and not significantly different from zero.

[40] 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, which 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 of the logistic and therefore the use of PPML can be justified.

Effects on sales and profits. 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 future profits, which are hard to measure in the data. Instead, we investigate whether more exposed firms experienced a decrease in sales and profits.

We create two new variables, sales per employee and an indicator variable for operating profits. The indicator takes 0, -1 or 1, depending on whether the firms makes zero, negative or positive profits. We estimate Equation (9) with these two new variables as dependent variables. The results from these regressions are presented in Table 6. We find that sales per employee decreased significantly, as did profits, although the latter result is not as robust as the first. Our results imply 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 parentheses are clustered by 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 0, -1 or 1, depending on whether the firms makes zero, negative or positive profits. $*p < 0.1, **p < 0.05, ***p < 0.01$.

Expansion or reallocation of R&D? Next, we explore whether the most exposed firms, relative to other firms, intensified their R&D investments across the board. As shown in Table 7, we find no effects of the shock on total R&D investments, number of R&D employees, or share of R&D employees in total employment. These results suggest that firms did not scale up all R&D investments in response to the shock. Rather, in line with our theoretical predictions, they appear to have reallocated 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 parentheses are clustered by 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 shed light on a novel channel through which a drop in the oil price may induce investments in clean technology in the fossil energy supply chain. We develop a theoretical framework that highlights the role of adjustment costs in R&D. Exploiting the rich data at hand, we develop a novel measure of firm-specific shock exposure. We show that reduced expected future profitability following the oil price shock induced firms in the fossil energy supply chain to shift their R&D investments towards clean R&D. While it is well known that a shift in relative prices in favor of clean energy sources will induce a shift in R&D in the same direction, our findings indicate that lower profitability in exposed firms creates an additional channel for reallocation of R&D investments. Furthermore, our results highlight the role of within-firm dynamics in the transition from dirty to clean energy sources.

Although our empirical analysis is based on the specific Norwegian context, we believe our results are also informative about innovation in the energy sector in other countries. For example, the onshore US fossil energy supply chain followed a very different path and was experiencing a boom around 2014. However, it is possible that a downturn in this industry would induce a similar shift towards clean innovation.

Our findings imply that policies reducing fossil-energy profitability, with carbon pricing as an important example, will induce clean innovation not only by increasing demand for clean technologies and shifting relative prices, but also by triggering a shift of resources from dirty to clean R&D within firms in the fossil energy supply chain. Our results indicate that fossil energy supply chain firms can repurpose their capabilities for the clean transition, and that reduced fossil-energy profitability encourages this shift—highlighting an additional advantage of policies like carbon pricing.

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, 6, 3, 21, B.4

_ , _ , 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, 6

_ , Ufuk Akcigit, Douglas Hanley, and William Kerr, "Transition to Clean Technology," *Journal of Political Economy*, 2016, *124* (1), 52–104. 6

Aghion, P., L. Barrage, E. Donald, D. Hémous, and E. Liu, "Transition to Green Technology along the Supply Chain," Discussion paper 20378, CEPR 2025. 6

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

Baumeister, Christiane and Lutz Kilian, "Understanding the Decline in the Price of Oil since June 2014," *Journal of the Association of Environmental and Resource Economists*, 2016, *3* (1), 131–158. 11

Bernard, Andrew B., Stephen J. Redding, and Peter K. Schott, "Multiple-Product Firms and Product Switching," *American Economic Review*, 2010, *100*, 70–97. 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. 9

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 2013, *81* (4), 1347–1393. 19

_ , 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. 4, 2.5

Boehm, Johannes, Swati Dhingra, and John Morrow, "The Comparative Advantage of Firms," *Journal of Political Economy*, 2022, *130*, 3025–3100. 1

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. 5.2

Casey, Gregory, "Energy Efficiency and Directed Technical Change: Implications for Climate Change Mitigation," *Review of Economic Studies*, 2024, *91* (1), 192–228. 6

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

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. 7

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. 9

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski, "Optimal Taxes on Fossil Fuel in General Equilibrium," *Econometrica*, 2014, *82* (1), 41–88. 6, 9

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. 6

Hamermesh, Daniel S., "Labor Adjustment and the Source of Adjustment Costs," *The Economic Journal*, 1995, *105*, 620–634. 24

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. 6

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

IEA, "The Future of Geothermal Energy," Technical Report, International Energy Agency 2024. 18

Kilian, Lutz, "The impact of the shale oil revolution on US oil and gasoline prices," *Review of Environmental Economics and Policy*, 2016, *10* (2). 11

Köppl, Angela and Margit Schratzenstaller, "Carbon Taxation: A Review of the Empirical Literature," *Journal of Economic Surveys*, 2023, *37* (4), 1353–1388. 9

Lemoine, Derek, "Innovation-Led Transitions in Energy Supply," *American Economic Journal: Macroeconomics*, 2024, *16* (1), 29–65. 6

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. 10

Magerman, Glenn, "Correspondences of EU Product Classifications," Discussion paper 2022. 5.2

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

Penrose, Edith, "Limits to the Growth and Size of Firms," *The American Economic Review*, 1955, *45* (2), 531–543. 1

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. 8

— **, Richard G Newell, and Adam B Jaffe**, “Energy, the Environment, and Technological Change,” *Handbook of the Economics of Innovation*, 2010, *2*, 873–937. 8

Prest, Brian C, “Explanations for the 2014 oil price decline: Supply or demand?,” *Energy Economics*, 2018, *74*, 63–75. 11

Prestmo, Joakim Blix, Birger Strøm, and Hilde Karoline Midsem, “Ringvirkninger av Petroleumsnæringen i Norsk Økonomi,” Technical Report 8 2015. 12

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

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. 5.2

Timilsina, Govinda R, “Carbon Taxes,” *Journal of Economic Literature*, 2022, *60* (4), 1456–1502. 9

Appendix

A Data

A.1 R&D survey

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 Organization for Economic Co-operation and Development (OECD) and Eurostat. As a general rule, the R&D survey is, in even years, based on firms with at least 10 employees, while in odd years, the survey is based on firms with at least 5 employees. The R&D survey includes three parts: (i) all firms with at least 50 employees; (ii) all firms with 5 (10)–49 employees and with substantially reported R&D in the previous survey (of more than NOK 1 million intramural R&D or more than NOK 3 million of extramural R&D); (iii) among other firms with 5 (10)–49 employees, a random sample selected within each stratum (NACE 2-digit industry and size class).

The R&D survey provides details on the share of R&D spending 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 in the R&D survey for the period 2007–2014, *Renewable energy* and *Other environment-related energy.* The questionnaire for the R&D survey changed slightly in 2015, and *Other environment-related energy* was replaced by *Energy efficiency.* Based on the description provided for the survey, it is our understanding that these thematic areas encompass the same type of clean R&D. We combine these categories both to get a time-consistent measure of clean R&D, and to increase statistical power.

B Calculations and extensions for the analytical model

B.1 Equilibrium firm behavior in the basic model

The firm problem given in Equation 2 is solved by three first-order conditions:

$$0 = -w_t - c\frac{\tilde{s}_{kct} - s_{kt}}{s_{kt}^2} + \beta\eta_{kc}\pi_c\frac{\partial g_c}{\partial \tilde{s}_{kct}} \;\forall t \tag{11}$$

$$0 = -w_t - c\frac{\tilde{s}_{kct} - s_{kt}}{s_{kt}^2} + \beta\eta_{kc}\pi_c\frac{\partial g_c}{\partial \tilde{s}_{kdt}} + \beta\eta_{kd}\pi_d g_d' \;\forall t \tag{12}$$

$$0 = \beta V_{kt+1}' = c\frac{\tilde{s}_{kt+1} - s_{kt+1}}{s_{kt+1}^2}\frac{\tilde{s}_{kt+1}}{s_{t+1}}, \;\forall t, \tag{13}$$

where $\tilde{s}_{kt+1}$ is the expected optimal total number of scientists engaged in R&D in the firm in the next period, determined by the expected price for the next period and the resulting expected equilibrium wage. Equation 13 implies that $s_{kt+1} = \tilde{s}_{kt+1}$ is optimal. Equations 11 and 12 jointly determine the optimal number of scientists to engage in each type of R&D in the current period, $\tilde{s}_{kct}(w_t, p^e_{ct+1}, p^e_{dt+1})$ and $\tilde{s}_{kdt}(w_t, p^e_{ct+1}, p^e_{dt+1})$. To simplify notation in the following, define:

$$A \equiv -\frac{c}{s^2_{kt}} + \beta\eta_{kc}\pi_c\frac{\partial^2 g_c}{\partial\tilde{s}^2_{kct}} \tag{14}$$

$$B \equiv -\frac{c}{s^2_{kt}} + \beta\eta_{kc}\pi_c\frac{\partial^2 g_c}{\partial\tilde{s}^2_{kdt}} + \beta\eta_{kd}\pi_d g''_d \tag{15}$$

$$C \equiv -\frac{c}{s^2_{kt}} + \beta\eta_{kc}\pi_c\frac{\partial^2 g_c}{\partial\tilde{s}_{kct}\partial\tilde{s}_{kdt}} \tag{16}$$

$$E \equiv \beta\eta_{kd}\pi'_d g'_d \tag{17}$$

The second-order conditions for the problem require that $A < 0$, $B < 0$ and $AB - C^2 > 0$, and we focus on interior solutions and assume that these hold.

Differentiating the first-order conditions in Equations 11 and 12 gives the direct response in clean and dirty R&D to a change in p^e_{dt+1} for a firm engaged in both types, as:

$$\frac{\partial\tilde{s}_{kdt}}{\partial p^e_{dt+1}} = -\frac{AE}{AB - C^2} > 0 \tag{18}$$

$$\frac{\partial\tilde{s}_{kct}}{\partial p^e_{dt+1}} = \frac{CE}{AB - C^2} \begin{cases} > 0 \text{ if technological complementarity dominates} \\ = 0 \text{ if there are no within-firm dynamics} \\ < 0 \text{ if adjustment costs dominate} \end{cases} \tag{19}$$

The signs of the partial derivative of of $\tilde{s}_{kct}$ for the three different cases imply the results presented in Proposition 1. Note that the direct response in clean R&D to the price drop for a firm h with $\eta_{hd} = 0$ will be zero, while the response in dirty R&D will be $\partial\tilde{s}_{ldt}/\partial p^e_{dt+1} = -E/F$ with $F \equiv -c/s^2_{lt} + \beta\eta_{ld}g''_d\pi_d$ for a firm l with $\eta_{lc} = 0$.

To calculate the full general-equilibrium response to the price shock for the firms, assume now that the firms are equally sized at the beginning of period t, that the share of firms with $\eta_{hd} = 0$ is given by α^h, the share with $\eta_{hc} = 0$ is α^l and the share with both innovation productivities larger than zero is α^k, with $\alpha^l + \alpha^h + \alpha^k = 1$. Moreover, assume that the firms have rational expectations regarding the future scientist wage and optimal hirings. Differentiating Equations 11 and 12 with respect to w_t gives the following partial responses to a wage change:

$$\frac{\partial \tilde{s}_{kdt}}{\partial w_t} = \frac{A - C}{AB - C^2} < 0 \tag{20}$$

$$\frac{\partial \tilde{s}_{kct}}{\partial w_t} = \frac{B - C}{AB - C^2} < 0 \tag{21}$$

$$\frac{\partial \tilde{s}_{ldt}}{\partial w_t} = \frac{1}{F} < 0 \tag{22}$$

$$\frac{\partial \tilde{s}_{hct}}{\partial w_t} = \frac{1}{A} < 0 \tag{23}$$

In equilibrium, we must have:

$$\alpha^l \tilde{s}_{ldt}(w_t, p^e_{dt+1}) + \alpha^h \tilde{s}_{hct}(w_t, p^e_{ct+1}) + \alpha^k \left(\tilde{s}_{kct}(w_t, p^e_{ct+1}, p^e_{dt+1}) + \tilde{s}_{kdt}(w_t, p^e_{ct+1}, p^e_{dt+1}) \right) = S \tag{24}$$

with S as the total supply of scientists in the market, for simplicity assumed to be given exogenously. The market equilibrium condition in 24 defines the equilibrium wage as a function of the expected future prices: $w_t(p^e_{ct+1}, p^e_{dt+1})$. Differentiating the condition with respect to the dirty energy price and inserting for the derivatives given in above give the equilibrium response in the wage to a change in the dirty energy price:

$$\frac{\partial w_t}{\partial p^e_{dt+1}} = -\frac{AEG}{H} > 0 \tag{25}$$

with

$$G \equiv \alpha^k F(C - A) - \alpha^l (AB - C^2) \tag{26}$$

$$H \equiv \alpha^h F(AB - C^2) + \alpha^l A(AB - C) + \alpha^k AF(A - C + B - C) \tag{27}$$

The total response in clean R&D to the dirty energy price drop for a firm k that is active in both clean and dirty R&D and production, and the total response for a firm h that is active only in clean R&D and production are given by:

$$
\begin{aligned}
\frac{d\tilde{s}_{kct}}{dp^e_{dt+1}} &\equiv \frac{\partial \tilde{s}_{kct}}{\partial p^e_{dt+1}} + \frac{\partial \tilde{s}_{kct}}{\partial w_t}\frac{\partial w_t}{\partial p^e_{dt+1}} && (28)\\
\frac{d\tilde{s}_{hct}}{dp^e_{dt+1}} &\equiv \frac{\partial \tilde{s}_{hct}}{\partial w_t}\frac{\partial w_t}{\partial p^e_{dt+1}}, && (29)
\end{aligned}
$$

respectively. Finally, we can now calculate the response in clean R&D in firm k relative to that of firm h to a dirty energy price drop:

$$
\begin{aligned}
\Delta \equiv & -\frac{d\tilde{s}_{kct}}{dp^e_{dt+1}} - \left(-\frac{d\tilde{s}_{hct}}{dp^e_{dt+1}}\right)\\
= & -\frac{1}{AB-C^2}\frac{E}{H}C\Bigg((\alpha^k+\alpha^h)F(AB-C^2)+\alpha^l(AB-C^2)C\Bigg). \qquad (30)
\end{aligned}
$$

Except for a special case that can arise if α^l is very large relative to $\alpha^k + \alpha^h$ (see discussion below), we have:

$$
\Delta \quad \begin{cases} > 0 \text{ if technological complementarity dominates (if C<0)} \\ = 0 \text{ if there are no within-firm dynamics (if C=0)} \\ < 0 \text{ if adjustment costs dominate (if C>0)} \end{cases} \qquad (31)
$$

which is the result presented in Proposition 2.

The exception that can render the inequalities in Equation 31 false arises if the last term in brackets in 30 is large enough to outweigh the rest of the expression in the brackets, i.e., if α^l is large relative to $\alpha^k + \alpha^h$. If this is the case, Δ will be positive if C is positive (complementarity dominates), while it will be negative if C is negative (adjustment costs dominate). The reason this relationship between the within-firm dynamics and the total firm reaction in clean R&D to a dirty-energy price drop can arise, is that the firms react differently to a change in the scientist wage when there are within-firm dynamics. This happens because the wage fall induces the k-firm to hire workers also for dirty R&D. If technological complementarity is important, more workers in dirty R&D counteracts the fall in marginal productivity of the scientists applied to the clean activity induced by increasing the number. In this case, the wage drop will lead the k-firm to increase its clean R&D more than an h-firm. On the other hand, if adjustment costs are important, the increase in hiring of scientists for dirty R&D in the k-firm increases the cost of hiring more scientists for the clean activity, and thus leads the k firm to increase its clean R&D less than an h-firm. Therefore, if the firms are hit by a drop in the scientist wage, the relative change in clean R&D will be different from that described in in Equation 31. If the sector is dominated by firms that do not engage in clean R&D or production (α^l very is large), and these firms react strongly to the dirty energy price drop, the firms in group k and h will in practice be hit mostly by a drop in the wage, while the direct response to the drop in the dirty energy price will be very small. In our empirical setting, with the large drop in the oil price in 2014 as our starting point, and with a labor market where centralized wage bargaining play a key role, it is very unlikely that the firms react almost only to an equilibrium wage change. We therefore disregard this special case in the rest of our analysis.

B.2 Technological complementarity: Generalization

In the basic model, we assume that the technological complementarity arises only through the production technology for clean innovation. The model can easily be extended to allow for somewhat more general technological complementarity, replacing $g_d(\cdot)$ from the basic model with $\hat{g}_d(\tilde{s}_{kdt}, \tilde{s}_{kct})$, with both first derivatives positive, the second derivatives negative and $\partial^2 \hat{g}_d / \partial \tilde{s}_{kct} \partial \tilde{s}_{kdt} > 0$. Firm k's problem is still given by 2, replacing $g_d(\cdot)$ with $\hat{g}_d(\cdot)$, and the first-order conditions are now given by:

$$0 = -w_t - c\frac{\tilde{s}_{kct} - s_{kt}}{sk_t^2} + \beta\eta_{kc}\pi_c\frac{\partial g_c}{\partial \tilde{s}_{kct}} + \beta\eta_{kd}\pi_d\frac{\partial \hat{g}_d}{\partial \tilde{s}_{kct}} \ \forall t \tag{32}$$

$$0 = -w_t - c\frac{\tilde{s}_{kct} - s_{kt}}{s_{kt}^2} + \beta\eta_{kc}\pi_c\frac{\partial g_c}{\partial \tilde{s}_{kdt}} + \beta\eta_{kd}\pi_d\frac{\partial \hat{g}_d}{\partial \tilde{s}_{kdt}} \ \forall t \tag{33}$$

$$0 = \beta V'_{kt+1} = c\frac{\tilde{s}_{kt+1} - s_{kt+1}}{s_{kt+1}^2}\frac{\tilde{s}_{kt+1}}{s_{kt+1}} \ \forall t. \tag{34}$$

Define:

$$\hat{A} \equiv -\frac{c}{s_t^2} + \beta\eta_{kc}\pi_c\frac{\partial^2 g_c}{\partial \tilde{s}_{kct}^2} + \beta\eta_{kd}\pi_d\frac{\partial^2 \hat{g}_d}{\partial \tilde{s}_{kct}^2} \tag{35}$$

$$\hat{B} \equiv -\frac{c}{s_t^2} + \beta\eta_{kc}\pi_c\frac{\partial^2 g_c}{\partial \tilde{s}_{kdt}^2} + \beta\eta_{kd}\pi_d\frac{\partial^2 \hat{g}_d}{\partial \tilde{s}_{kdt}^2} \tag{36}$$

$$\hat{C} \equiv -\frac{c}{s_t^2} + \beta\eta_{kc}\pi_c\frac{\partial^2 g_c}{\partial \tilde{s}_{kct}\partial \tilde{s}_{kdt}} + \beta\eta_{kd}\pi_d\frac{\partial^2 \hat{g}_d}{\partial \tilde{s}_{kct}\partial \tilde{s}_{kdt}} \tag{37}$$

$$\hat{E}_d \equiv \beta\eta_{kd}\pi'_d\frac{\partial \hat{g}_d}{\partial \tilde{s}_{kdt}} \tag{38}$$

$$\hat{E}_c \equiv \beta\eta_{kd}\pi'_d\frac{\partial \hat{g}_d}{\partial \tilde{s}_{kct}}. \tag{39}$$

As before, differentiating, the first-order conditions gives the partial derivative of $\tilde{s}_{kct}$ with respect to the dirty energy price for a firm engaged in both types of R&D and production:

$$\frac{\partial \tilde{s}_{kct}}{\partial p^e_{dt+1}} = \frac{\hat{C}\hat{E}_d - \hat{B}\hat{E}_c}{\hat{A}\hat{B} - \hat{C}^2}. \tag{40}$$

This expression takes a negative value if the effect of adjustment costs dominate, while it takes a positive value if the effects of technological complementarity dominates, as in the basic model (Proposition 1). As in the basic model, there will be no direct effect of the price drop in a a firm engaged only in clean R&D and production. The two-directional technological complementarity does not alter the qualitative effects of the within-firm dynamics, as the added complementarity works in the same direction as the complementarity already included in the basic model. As a consequence, the general-equilibrium effects derived in Section B.1 play out in the same way in this extended model, again leading the total effect of a drop in the dirty energy price to increase clean investments more (less) in a firm engaged in both types of R&D if adjustment costs (technological complementarity) dominates.

B.3 Costs of adjusting type-specific R&D

In the basic model, we followed the standard assumption in the literature on directed technical change that scientists can be applied to both types of R&D. Moreover, when including adjustment costs, we only included costs of adjusting the overall R&D, while allowing firms to costlessly shift workers from one type of R&D to the other. Our empirical findings support the hypothesis that there exist overall R&D adjustment costs and, moreover, that it is possible for the firms to shift resources from dirty to clean R&D. However, there may very well also be costs to changing the size of each of the two types of R&D within a firm that is engaged in both types. In this section, we show that the theory model can easily be extended to allow for such more general adjustment costs. We show that the consequence of the overall adjustment cost and of the technological complementarity for the response in clean R&D to a drop in the dirty energy price does not change qualitatively when such type-specific adjustment costs are included in the framework, as long as the firms have some possibility of shifting scientists between the two types of R&D.

In the same model setup as presented in Section 3.1, assume that the adjustment cost in Equation 1 is replaced by the following cost:

$$\hat{R}_{kt} = \frac{1}{2}c \cdot \left(\frac{\tilde{s}_{kt} - s_{kt}}{s_{kt}}\right)^2 + \frac{1}{2}c_c \cdot \left(\frac{\tilde{s}_{kct} - s_{kct}}{s_{kct}}\right)^2 + \frac{1}{2}c_d \cdot \left(\frac{\tilde{s}_{kdt} - s_{kdt}}{s_{kdt}}\right)^2 \tag{41}$$

The firm solves the same problem as in the basic model, with the adjustment costs replaced by this expression. As before, differentiating the first-order conditions gives the response in clean R&D to a change in p^e_{dt+1}, and the expression is now given by:

$$\frac{\partial \tilde{s}_{kct}}{\partial p^e_{dt+1}} = \frac{CE}{\hat{\hat{A}}\hat{\hat{B}} - C^2}, \tag{42}$$

with $\hat{\hat{A}} \equiv A - c/s^2_{kct}$ and $\hat{\hat{B}} \equiv B - c/s^2_{kdt}$. This expression takes a positive value if the effect of adjustment costs dominate, while it takes a negative value if the effects of technological complementarity dominates, as in the basic model (Proposition 1). The within-type adjustment costs will weaken the effect of both technological complementarity and the overall adjustment cost on the response in clean innovation to a change in the dirty energy price, as it makes it more costly for the firm to shift scientists from one type of R&D to the other. However, as long as the marginal cost is sufficiently low for the first scientist that is shifted, the qualitative effect of the overall adjustment cost and of the technological complementarity on the response is the same as in the basic model. And as in the previous section, the conclusions from the basic model regarding the general-equilibrium effect of the price drop is therefore also, qualitatively, the same.

B.4 Within-firm dynamics in a general model of directed technical change

In this section, we add the within-firm dynamics from our basic model to a model of directed technical change, to illustrate that the mechanisms play out in the same way as described in Section 3 when the dynamics are included in this well-known framework in the literature. We follow Acemoglu et al. (2012), and extend their model by including heterogeneous firms conducting R&D that face adjustment costs in their total R&D and technological complementarity between clean and dirty R&D.

As in Section 3, we consider a discrete-time, infinite-horizon economy. Now, let Y_{jt}, with $j = c, d$, be the quantity produced in period t of clean and dirty energy. Because we consider a small open economy, the final energy prices are exogenously given, denoted p_{jt}. The production technology for the energy goods is given by:

$$Y_{jt} = \int_0^1 A^{1-\alpha}_{jit} x^{\alpha}_{jit} di, j = c, d \tag{43}$$

with $\alpha \in (0, 1)$, where x_{jit} is the quantity of input ji. A_{jit} denotes the productivity of the input at time t. We abstract from the use of other inputs, such as labor, in the final good sector.

There is a mass of 1 input suppliers 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, let $\phi = \alpha^2$. Innovation takes place within the input-supplying firms. Define A_{jt} as the aggregate state of the technology for type j at time t, with:

$$A_{jt} \equiv \int_0^1 A_{jit} di \tag{44}$$

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.[41] For all inputs with no successful innovation, the monopoly rights are allocated to a random firm.

To innovate, the input suppliers must hire scientists to conduct R&D. 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 scientist hired for type-j R&D is randomly assigned to an input of that type. The input suppliers act competitively in the market for scientists. As in Section 3, $\tilde{s}_{kjt}$ is the number of scientists engaged in type-j R&D in firm k in period t,while s_{kjt} is the number of scientists that the firm has in their type-j work force at the beginning of the period, with $s_{kt} = s_{kdt} + s_{kct}$. Each firm has an innovation productivity for each type, given by $\eta_{kj} \in \{0, \eta_j\}$, with $\eta_j \in (0, 1)$, and the innovation production functions are defined as in Section 3, as $g^c(\tilde{s}_{kct}, \tilde{s}_{kdt})$, and $g^d(\tilde{s}_{kdt})$.

Firms will not engage in R&D of type j if their innovation productivity for that type is zero. Let α^l, α^h,and α^k represent the share of firms with only dirty R&D, with only clean R&D and with both types of R&D, respectively. The adjustment cost is modeled as in Equation 1.

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

$$A_{ct} = \left(1 + \gamma \int_0^1 \eta_{kc} g^c(\tilde{s}_{kct}, \tilde{s}_{kdt}) dk\right) A_{ct-1} \tag{45}$$

$$A_{dt} = \left(1 + \gamma \int_0^1 \eta_{kd} g^d(\tilde{s}_{kdt}) dk\right) A_{dt-1} \tag{46}$$

with $\int_0^1 s_{kt} dk = 1$ and $\int_0^1 \tilde{s}_{kt} dk = 1$ in equilibrium. Allocation of scientists will therefore drive the technological state of the economy in the long run. Because there is path dependency in the technological development in the model, a temporary change in the

[41] Note that, formally, the timing is adjusted here relative to the basic model, to follow the standard in this literature. The effective timing is the same, as the investment decision for R&D happens before the profits from innovation is realized.

allocation of scientists, for example induced by within-firm dynamics, may have long-lasting effects in the economy.

B.4.1 Market equilibrium

Let p_{jit} denote the price of input ji. The producer of energy of type j will choose inputs to maximize profits:

$$\max_{x_{jit}} p_{jt} \int_0^1 A_{jit}^{1-\alpha} x_{jit}^{\alpha} di - \int_0^1 p_{jit} x_{jit} di \tag{47}$$

Solving the maximization problem gives the downward-sloping demand for each input of $p_{jit}(x_{jit}) = \alpha p_{jt} A_{jit}^{1-\alpha} x_{jit}^{\alpha-1}$. The monopolistic producer of good x_{jit} will maximize profits subject to the demand for inputs. In equilibrium, production of and profits related to input ji are given by $x_{jit} = A_{jit} p_{jt}^{\frac{1}{1-\alpha}}$ and $\pi_{jit} = \alpha(1-\alpha) A_{jit} p_{jt}^{\frac{1}{1-\alpha}}$, respectively. Finally, the expected within-period revenue from engaging scientists in type-j R&D for firm k in period t, denoted I_{kjt}, is given by:

$$\begin{aligned} I_{kct} =& \eta_{kc} g^c(\tilde{s}_{kct}, \tilde{s}_{kdt}) \int_0^1 \pi_{cit} di = \eta_{kc} g^c(\tilde{s}_{kct}, \tilde{s}_{kdt}) \alpha(1-\alpha) p_{ct}^{\frac{1}{1-\alpha}} (1+\gamma) A_{ct-1} &(48)\\ I_{kdt} =& \eta_{kd} g^d(\tilde{s}_{kdt}) \int_0^1 \pi_{dit} di = \eta_{kd} g^d(\tilde{s}_{kdt}) \alpha(1-\alpha) p_{dt}^{\frac{1}{1-\alpha}} (1+\gamma) A_{dt-1} &(49)\end{aligned}$$

The number of scientists hired for the beginning of the next period is, as in the basic model, chosen to equal the expected optimal number for the next period. Because this choice is independent of the choice of the number of scientists to engage in the current period, we limit the analysis here to the choice of $\tilde{s}_{kct}$ and $\tilde{s}_{kdt}$. Scientists are paid wage w_t, and firm k thus engages scientists to solve the following problem in period t:

$$\max_{\tilde{s}_{kdt}, \tilde{s}_{kct}} \left\{ \alpha(1-\alpha)(1+\gamma) \left(\eta_{kc} p_{ct}^{\frac{1}{1-\alpha}} A_{ct-1} g^c(\tilde{s}_{kct}, \tilde{s}_{kdt}) + \eta_{kd} p_{dt}^{\frac{1}{1-\alpha}} A_{dt-1} g^d(\tilde{s}_{kdt}) \right) - w_t \tilde{s}_{kt} - \frac{1}{2} c \cdot \left(\frac{\tilde{s}_{kt} - s_{kt}}{s_{kt}} \right)^2 \right\} \tag{50}$$

subject to $\tilde{s}_{kt} = \tilde{s}_{kct} + \tilde{s}_{kdt}$. Note that the problem for firm k is completely parallel to that in the basic model (Equation 2) The first-order conditions are given by:

$$\begin{aligned} 0 =& \alpha(1-\alpha)(1+\gamma) \eta_c A_{ct-1} p_{ct}^{\frac{1}{1-\alpha}} \frac{\partial g^c(\tilde{s}_{kct}, \tilde{s}_{kdt})}{\partial \tilde{s}_{kct}} - w_t - c \frac{\tilde{s}_{kt} - s_{kt}}{s_{kt}^2} &(51)\\ 0 =& \alpha(1-\alpha)(1+\gamma) \eta_d A_{dt-1} p_{dt}^{\frac{1}{1-\alpha}} g^{d\prime}(\tilde{s}_{kdt}) - w_t - c \frac{\tilde{s}_{kt} - s_{kt}}{s_{kt}^2}. &(52)\end{aligned}$$

We let the resulting demand for scientists from a firm h that conducts only clean R&D, be given by $\tilde{s}_{hct}(p_{ct}, w_t)$. Demand for scientists from a firm that conducts only dirty R&D, is given by $\tilde{s}_{ldt}(p_{dt}, w_t)$, and demand for scientists from firm k that conducts both types of R&D, is given by $\tilde{s}_{kct}(p_{ct}, p_{dt}, w_t)$ and $\tilde{s}_{kdt}(p_{ct}, p_{dt}, w_t)$. Equilibrium in the market for scientists requires:

$$\alpha^h \tilde{s}_{hct}(p_{ct}, w_t) + \alpha^l \tilde{s}_{ldt}(p_{dt}, w_t) + \alpha^k \tilde{s}_{kct}(p_{ct}, p_{dt}, w_t) + \alpha^k \tilde{s}_{kdt}(p_{ct}, p_{dt}, w_t) = 1 \tag{53}$$

Differentiating the first-order conditions and the equilibrium condition gives the direct response in $\tilde{s}_{kct}$ to a drop in p_{dt},which depends on the within-firm dynamics in the same way as described in the basic model: Adjustment costs lead the firm to increase its clean R&D to avoid adjustment costs if the dirty energy price falls, while the technological complementarity results in lower clean R&D as when p_{dt} falls, because lower dirty R&D lowers the marginal productivity of scientists in the clean R&D. In equilibrium, the wage will fall as a consequence of lower overall demand for scientists, leading all firms to increase their overall R&D as a response to the price drop. Defining $\hat{\Delta}$ as the change in clean R&D in exposed firms (firms engaged in both clean and dirty R&D) relative to the change in non-exposed firms (only engaged in clean R&D):

$$\hat{\Delta} \equiv -\frac{d\tilde{s}_{kct}}{dp_{dt}} - \left(-\frac{d\tilde{s}_{kct}}{dp_{dt}} \right) \tag{54}$$

with

$$\frac{d\tilde{s}_{kct}}{dp_{dt}} = \frac{\partial \tilde{s}_{kct}}{\partial p_{dt}} + \frac{\partial \tilde{s}_{kct}}{\partial w_t}\frac{\partial w_t}{\partial p_{dt}} \tag{55}$$

$$\frac{d\tilde{s}_{hct}}{dp_{dt}} = \frac{\partial \tilde{s}_{hct}}{\partial w_t}\frac{\partial w_t}{\partial p_{dt}}, \tag{56}$$

we have:

$$\hat{\Delta} \begin{cases} > 0 \text{ if adjustment costs dominate} \\ = 0 \text{ if there are no within-firm dynamics} \\ < 0 \text{ if technological complimentarity dominates} \end{cases} \tag{57}$$

C Firm-Level Measure of Exposure to Oil Price Drop

We use a two-step process to identify the level of exposure of each firm in our dataset to the 2014 oil price drop as explained in Section 4.1.

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

When determining which products to define as inputs to oil production, we take into account the potentially strong home bias of the oil industry operating in Norway regarding 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 achieve this goal. If the home bias differs across industries and inputs, the share of imports of good that is directed towards the oil-extracting firms will not be informative about the relative importance of good as an input to that industry. Therefore, we do not weight the goods categorized as inputs to the oil-extracting firms using import shares.

D Additional Tables

Table 8: Clean R&D: Dummy, share and value with PPML

Variable:	Dummy (1)	Share (2)	Value (3)
$Post_t$*x_{oi}	0.499**	0.358*	0.716
	(0.199)	(0.218)	(0.470)
Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Ind.*year FE	Yes	Yes	Yes
Observations	3,024	3,024	3,024

Note: Standard errors in parentheses are clustered by 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. $*p < 0.1, **p < 0.05, ***p < 0.01$.

E Additional Figures

Figure 8: Producer price index – Norway (2021=100)

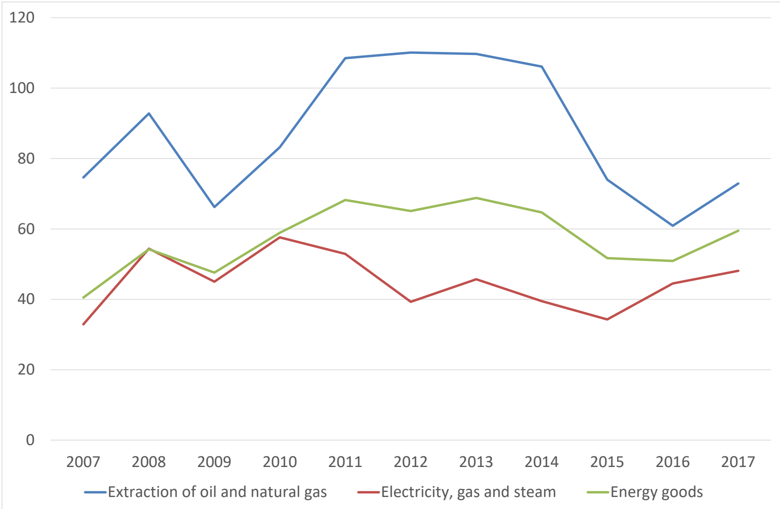

The figure shows the development of the producer price index for the domestic market for the commodities/industries (i) oil and gas extraction, (ii) electricity, gas and steam and (iii) the aggregate of energy goods, with 2021 as the reference year (2021=100). Source: Statistics Norway.

Figure 9: Global renewable costs (IRENA)

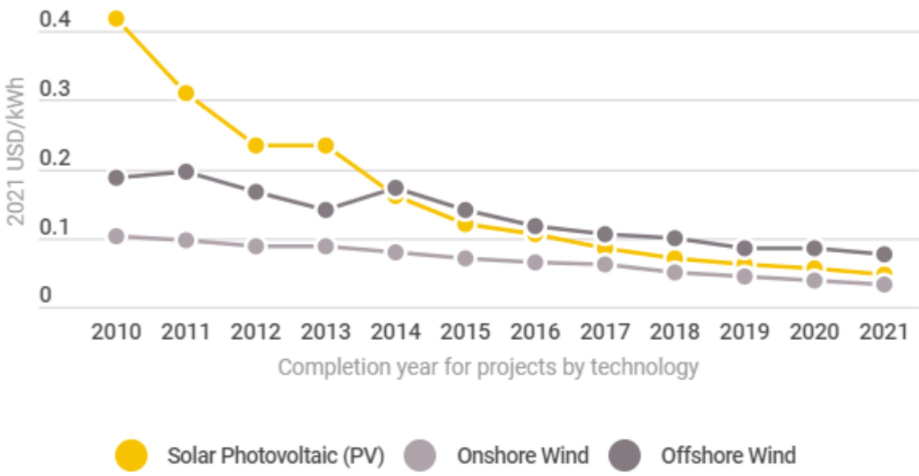

The figure shows the development of the cost of renewable energy by technology. Source: IRENA (www.irena.org).